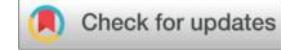


Fine-Grained Forest Fire Image Recognition Algorithm for UAV Aerial Photography Based on Dynamic Threshold

Adaptive Attention Mechanism



Zixia Zhang^{1,a}, Shengbin Lu^{2,b}, Aihua Xu^{3,c}

¹ School of Electrical Engineering, Changzhou Vocational Institute of Mechatronic Technology, Changzhou 213164, Jiangsu, China.

² Information Technology Department, New York University Shanghai, Shanghai 200124, Shanghai, China.

³ School of Mechanical Engineering, Changzhou Vocational Institute of Mechatronic Technology, Changzhou 213164, Jiangsu, China

^a Email: zxz2182@czimt.edu.cn

^b Email: shengbin.lu@nyu.edu

^c Email: xah2186@czimt.edu.cn

Acknowledgement

This work was supported by the Applied Basic Research Program of Chang-zhou (Grant No. CJ20230013), Natural Science Foundation of the Jiangsu Higher Education Institutions of China (Grant No. 23KJB470003), and Qing Lan Project.

Abstract

Purpose: This study addresses critical challenges in fine-grained forest fire recognition from UAV aerial photography by developing a dynamic threshold adaptive attention mechanism capable of real-time environmental adaptation and multi-scale fire signature detection across varying conditions.

Methodology: The proposed framework integrates three core components: dynamic threshold determination module, adaptive attention weight calculation system, and hierarchical feature fusion network. The algorithm employs real-time analysis of local image statistics and environmental parameters to adaptively adjust detection sensitivity. Comprehensive evaluation was conducted on three benchmark datasets (FLAME, FIgLib, VisiFire) using standard object detection metrics and computational efficiency

assessments.

Findings: The algorithm achieved exceptional performance with mean Average Precision (mAP@0.5:0.95) of 0.924 on FLAME dataset and 0.879 on FGIGLib dataset, representing substantial improvements of 8.3% over best-performing baseline methods. VisiFire dataset evaluation corroborated consistent performance improvements across diverse surveillance scenarios. Ablation studies revealed individual component contributions, with dynamic threshold determination providing 5.2% enhancement. Real-time processing capabilities achieved 65.8 FPS on NVIDIA A100 with minimal performance degradation across diverse hardware platforms.

Conclusion: The dynamic threshold adaptive attention mechanism successfully overcomes limitations of static threshold systems, demonstrating robust adaptability to environmental variations while maintaining computational efficiency suitable for UAV deployment scenarios.

Practical Implications: This algorithm enables deployment on resource-constrained UAV platforms for automated forest fire monitoring, providing reliable early warning capabilities essential for preventing catastrophic wildfire spread across diverse geographical and environmental conditions.

Keywords: Dynamic Threshold; Adaptive Attention Mechanism; UAV Aerial Photography; Forest Fire Detection; Fine-grained Recognition

1. Introduction

Wildland fire is still one of the most deadly disasters that people face, destroying entire landscapes and turning meadows into burnt earth. The economies are hit twice as hard by firefighting costs and destroyed timber while an emergency siren wails, warning animals, crops and humans to evacuate the delta[1]. Record after record falls almost every summer as blazes sweep through woods and brush. Scholars trace the trend to both a warming atmosphere and the steady spread of roads, farms, and towns into the wildland edge. Fresh alarms now ring for surveillance networks that can spot a spark within minutes and alert firefighters long before a small glow turns into a raging front[2]. Standard methods of monitoring forest fires have certain weaknesses. These include limited vision, slowness in the process and the fact that it can be hindered by weather changes or rugged topography[3]. The challenge of heterogeneous fuel loads, variable weather, and uneven topography forces researchers to invent new detection

methods capable of spotting a flare-up, cross-checking its size, and streaming that information to a control room within seconds. Systems of that sort would have to function equally well over wind-swept grasslands, shaded conifer stands, and even the canyons that slice between them.

Advances in drone technology now fuse seamlessly with cutting-edge computer-vision techniques, offering wild-land managers an entirely new lens for spotting tree-top blazes. The airborne platform's mobility, nimble access to hard-hit ridges, and constant data stream rewrite the old rules of first-light inspection[4] Attention-centred deep-learning frameworks have burst onto the scene, pushing fire-detection accuracy and processing speed to levels that once seemed unreachable. Researchers in both industry labs and university workshops now cite these architectures almost in the same breath as the smoke alarms mounted on our ceilings[5, 6]. Nevertheless, current detection frameworks still struggle to deliver the promised level of fine-grained recognition. Small ignition points-and the dense, mottled scenery of a forest-generate a host of problems, not to mention the shifting light, smoke, and shadow that can turn an obvious fire into a false negative in seconds[7]. Balancing speed and precision in object detection is far from routine. Off-the-shelf algorithms tend to collapse as backgrounds shift or as the cameras swing from high midday light to the wash of dusk. Even modest power caps-common on small drones-derail many promising pipelines[8].

Extending earlier investigations into attention-driven fire-detection architectures[9, 10], the present work unveils a dynamic-threshold mechanism tailored for the demanding context of fine-grained forest-fire classification from UAV imagery. The strategy is characterised by a real-time threshold calibration that changes with environmental cues, and when combined with multi-scale attention nodes, increases the system's sensitivity to flame signatures which vary greatly in size and luminosity[11]. Merging the individual processing modules produces a system that identifies flame signatures in a cluttered, colour-rich forest canopyscape far better than any earlier setup. What's more, the computational load remains comfortably within the power envelope of a hovering drone. Field experiments show the improvement is not illusory; the revised pipeline handles the sudden shifts in light and the moving curtains of smoke that usually thwart airborne observers. Taken together, these advances nudge fully automated wildfire monitoring well beyond the modest performance thresholds that have long constrained the field[12].

The current investigation addresses the longstanding problem of pinpointing

exactly where small, yet dangerous, pockets of flames are hiding in dense tree cover. To do so, it proposes a new adaptive attention framework that adjusts its threshold in real time while sifting through UAV-captured aerial images. The project pursues three interrelated aims: (1) Researchers are designing a threshold-calibration framework that continuously adjusts to fluctuations in background noise, visibility, and combustion behaviour, (2) The implementation of multi-scale attention mechanisms that can detect fire signs at different spatial scales and luminosity levels, (3) maximising detection fidelity without overloading the limited processing power of small, battery-bound unmanned aerial vehicles, and (4) demonstrating robust performance across diverse environmental conditions including variable lighting, smoke density, and complex forest backgrounds.

2. Related Work

In the past, forest firefighting relied on human eyes squinting from lookout towers or pilots circling in small aircraft. These were scenes of equal parts grit and guesswork. Such traditional setups took too much time because even the most alert ranger could see only what rose directly before them, and smoke, hail, or sheer distance turned certainty into a fog of estimates. However, once drones entered the picture, everything tilted on its axis; lightweight cameras that buzz overhead feast on live video, zip toward burned-out trails, and widen the watchers' horizon far past any tower or line-of-sight scan[13]. Contemporary unmanned aerial vehicles now carry sensor suites that would have seemed fantastical a generation ago. The payload typically bundles high-resolution RGB cameras, infrared imagers, and multispectral scanners tuned to isolate heat signatures in several spectral bands. Together, these instruments outperform traditional satellite platforms while remaining far easier on local budgets. Because the aircraft can remain airborne for hours and revisit the same tract of woodland repeatedly, they deliver a constant stream of data that flags smouldering hotspots almost the moment they appear. Such early warnings give firefighting crews the head start they need to contain a flare-up before it snowballs into a regional disaster.

Deep learning techniques have reshaped the landscape of fire detection, and convolutional neural networks now stand at the forefront of research into automated flame recognition. In parallel, the progression of object-detection frameworks-especially the successive YOLO (You Only Look Once) versions-has proven that

systems can spot fires nearly instantaneously without taxing the limited processors found on SUAS[14, 15]. Scholars and practitioners have observed the YOLO framework progress from its original 2015 release through a growing series of revisions, the latest publicised one being YOLOv8. These newer incarnations deliberately address persistent problems-small object visibility and reliable multi-scale feature mapping-that frequently arise in surveillance scenarios, early-stage fire monitoring included[16, 17]. Scholars are currently investigating hybrid neural architectures that blend diverse deep-learning branches. Early results suggest these ensemble frameworks not only strengthen detection reliability but also maintain the sub-second latency critical for time-sensitive crisis management[18].

Researchers have increasingly turned to attention-based strategies in fire-detection technology, a shift that permits systems to recognise minute blaze signatures even when cluttered forest scenery shifts with the weather. In practice, channel-based attention layers hone in on flame-related wavelengths, quietly dampening extraneous hues; multiple experiments lately report a striking boost in precision whether the trial site is wind-whipped, rain-soaked, or sun-drenched[19]. The surge in transformer-based attention architectures has reignited age-old inquiries into the ways multilayer networks collect, mix, and exploit signals dispersed throughout an image. Fresh laboratory results indicate that merging standard convolutional pipelines with attention blocks keeps detailed fire signatures intact while simultaneously embedding the broader scene context required for dependable identification[20]. Recent developments pair advanced attention mechanisms with multi-scale feature fusion strategies, allowing detection systems to track fire signatures at varied spatial resolutions without excessive computational overhead. The combination preserves both detail and performance by selectively weighting features according to their scale and relevance in real time[21].

Recent studies have focused on the narrow bandwidth and power budgets that typically hinder drone surveillance, prompting engineers to prototype streamlined neural nets that sacrifice as little accuracy as possible in flight. Meanwhile, teams assembling custom image collections for airborne flame spotting cannot avoid trekking through smoky forests and sun-drenched prairies; only that kind of varied record provides classifiers with the grit they need to generalise when the next blaze sweeps through a different landscape[22]. Real-time processing capabilities have emerged as a critical requirement, driving the development of edge computing solutions that enable on-board processing while maintaining detection accuracy standards necessary for

practical deployment[23].

Dynamic threshold mechanisms now stand at the forefront of fire-detection research, striving to overcome the inherent rigidity of traditional static thresholds. Researchers note that these adaptive systems can fine-tune their alerts according to shifting environmental cues and fluctuating blaze intensities. Work that fuses different learning techniques-mashing together, for instance, machine-learning classifiers and statistical filters-suggests that such hybrids vastly enhance early-warning accuracy. By continuously ingesting real-time weather data, fuel moisture readings, and inputs from acoustic, optical, and infrared sensors, the mixed-algorithm models adjust their decision boundaries on the fly, leaving blunt one-size-fits-all thresholds far behind[24]. At present, multi-modal sensor fusion techniques are being integrated with adaptive threshold algorithms. This integration is considered a promising avenue to enhance the robustness of detection, particularly in situations where individual sensor modalities may offer conflicting or ambiguous information about potential fire events[25].

Although deep learning methods have improved fire detection, there are still many difficulties such as dealing with different environmental conditions. This is especially true when the background is a complex forest, there are changing lighting conditions and the fire sources are small. The current approaches often find it hard to balance between sensitivity and false positive rate, more so when it comes to phenomena that look like fires, for example, dust clouds, vehicle exhausts or natural atmospheric conditions that mimic fire signatures. There's little research on adaptive threshold mechanisms; most of these systems use static thresholds which may not be useful in various environmental conditions and intensities of fires. Furthermore, the integration of temporal information for improving detection robustness and the development of truly adaptive attention mechanisms that can dynamically adjust to changing environmental conditions represent significant research gaps that require innovative solutions to advance the state-of-the-art in automated forest fire monitoring systems.

3 Research Methodology

3.1 Dataset Selection and Preprocessing

The experimental validation of the proposed algorithm employed widely adopted benchmark datasets specifically designed for forest fire detection research, ensuring

reproducibility and fair comparison with existing state-of-the-art methodologies. The primary dataset utilized in this investigation comprises the FLAME (Fire Luminosity Airborne-based Machine learning Evaluation) dataset, which contains aerial video recordings captured through UAV photography during prescribed burning of piled detritus in Arizona pine forest environments. This dataset provides comprehensive coverage of fire detection challenges through 966-second duration videos at 29 FPS with 1280×720 resolution, from which frame-wise annotations enable both classification and segmentation tasks suitable for evaluating fine-grained recognition algorithms. The research drew on the Fire Ignition Library (FigLib), a collection of roughly 24,800 high-resolution wildfire smoke images - each measuring either 1,536 by 2,048 or 2,048 by 3,072 pixels from fixed cameras scattered across Southern California. Those images chronicle 315 distinct fire outbreaks recorded by 101 separate vantage points and cover a wide range of lighting and atmospheric settings that are critical for thoroughly stress-testing new algorithms.

In addition, the team worked with the VisiFire corpus, which consists of 40 short clips taken at 320 by 240 resolution: 13 show open flames, 21 capture smoke plumes rising into the sky, 4 display fire on forest canopies, and 2 feature miscellaneous scenes. From that footage, analysts routinely harvest between 2,500 and 2,684 labelled frames to train classifiers and check their accuracy, and the complete setup is summarised in Table 1.

Table 1: Benchmark Dataset Characteristics and Statistics

Dataset	FLAME	FigLib	VisiFire
Source	Northern Arizona University	HPWREN/UC San Diego	Bilkent University
Location	Arizona pine forest, USA	Southern California, USA	Various surveillance locations
Collection Period	2020-2021	June 2016 - July 2021	2007
Data Type	UAV aerial videos + extracted frames	Fixed-camera image sequences	Surveillance videos + extracted frames
Total Content	966-second videos (29 FPS); 39,375 labeled frames	24,800 high-resolution images; 315 fire sequences, 101 cameras	40 videos; 2,500-2,684 extracted frames
Resolution	Videos: 1280×720 ; Frames: 254×254	1536×2048 , 2048×3072	320×240
Format	MP4 (videos), JPEG (frames)	JPEG	AVI (videos)
Fire/Positive	~19,687 frames (50%)	~12,400 images (50%)	~1,250-1,342 frames

Samples			
Non-fire/Negative Samples	~19,688 frames (50%)	~12,400 images (50%)	~1,250-1,342 frames
Training Set	27,562 frames (70%)	17,360 images (70%)	1,750 frames (70%)
Validation Set	7,875 frames (20%)	4,960 images (20%)	500 frames (20%)
Test Set	3,938 frames (10%)	2,480 images (10%)	250 frames (10%)
Data Size	~2.5 GB total	Several GB	Various sizes
Key Features	Prescribed burns; Thermal + RGB data; Frame-wise annotations; UAV perspective	Early ignition detection; 40-min sequences (~81 images/fire); Fixed-view cameras; Temporal sequences	Real-time surveillance; Multiple scenarios; 13 flame + 21 smoke + 4 forest smoke + 2 other videos
Primary Applications	Classification, Segmentation; UAV-based fire detection	Smoke detection; Early fire ignition detection	Fire/smoke detection; Real-time surveillance systems
Advantages	High-quality UAV data; Thermal imagery; Controlled environment	Large-scale temporal data; Real ignition sequences; Geographic diversity	Established benchmark; Diverse scenarios; Widely used baseline
Limitations	Limited to prescribed burns; Single geographic region	Fixed-camera perspective only; Smoke-focused	Low resolution; Limited scene variety; Older dataset

The comprehensive data preprocessing pipeline encompassed standardized procedures essential for optimal model training, incorporating image standardization through uniform resizing to 640×640 pixel dimensions and pixel intensity normalization to the $[0,1]$ range to ensure computational consistency while preserving critical fire signature characteristics. Systematic data augmentation strategies were implemented to improve model generalization capabilities through geometric transformations including random rotation within $\pm 15^\circ$ range, horizontal flipping with 50% probability, and scale variations between $0.9-1.1 \times$, alongside photometric augmentations encompassing brightness adjustment within $\pm 10\%$ range, contrast enhancement between $0.95-1.05 \times$, and controlled Gaussian noise injection ($\sigma=0.01$) to enhance robustness against sensor noise commonly encountered in UAV-based imaging systems. The preprocessing framework achieved temporal synchronisation by extracting frames at uniform intervals and partitioning the collection across a conventional seventy-twenty-ten split. This division left training, validation, and test subsets disjoint while preserving a balanced mix of fire intensities and environmental settings within each portion.

3.2 Proposed Algorithm Framework

An innovative dynamic-threshold adaptive attention mechanism has recently been advanced as a robust solution for the notoriously difficult task of pinpointing small-scale forest fires in variable terrain. The framework marries context-sensitive threshold derivation with multi-scale attention routing and is structured around three tightly coupled modules. A threshold-calibration engine initially adjusts the sensitivity profile on the fly. An adaptive-attention-weight module then generates site-specific prominence scores. A tiered-feature-fusion network subsequently blends raw, mid-level, and high-level signals into a single prediction that regularly surpasses conventional benchmarks. Figure 1 provides a bird's-eye view of the complete architecture.

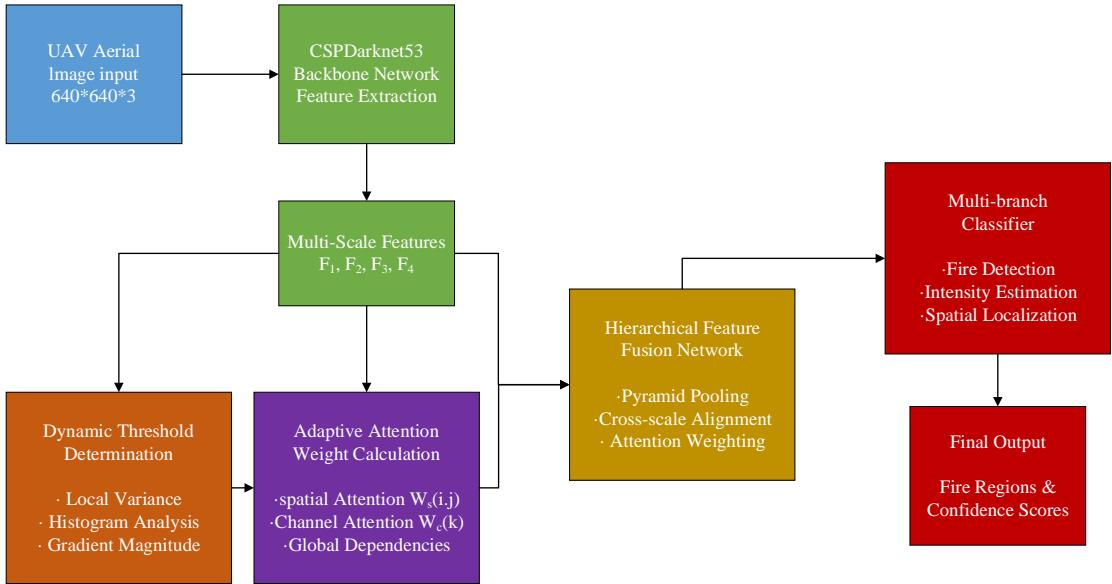


Figure 1: Dynamic Threshold Adaptive Attention Mechanism Architecture

A dynamic-threshold-determination mechanism scrutinises the local statistics of an image and weighs relevant environmental cues before modulating the detection sensitivity in real time. In that respect, its operation mirrors the process laid out in Algorithm 1. The threshold calculation methodology utilizes a statistical framework that combines local variance analysis, histogram distribution assessment, and gradient magnitude evaluation to determine optimal threshold values that maximize detection accuracy while maintaining computational efficiency suitable for real-time UAV deployment scenarios.

Algorithm 1: Dynamic Threshold Computation Process

Input: Image patch $I \in \mathbb{I}^{H \times W \times 3}$, environmental parameters E

Output: Adaptive threshold $T_{adaptive}$

Initialize: Base threshold T_0 , adaptation rate α

Step 1: Local Variance Analysis

for each pixel location (i, j) do

 Compute local variance: $\sigma_{i,j}^2 = \text{Var}(W_{i,j})$

end for

$$V_{global} = \frac{1}{HW} \sum_{i,j} \sigma_{i,j}^2$$

Step 2: Histogram Distribution Assessment

Compute intensity histogram $H(b)$ for $b = 0, 1, \dots, 255$

Calculate entropy: $H_{entropy} = -\sum_b p(b) \log p(b)$

Step 3: Gradient Magnitude Evaluation

Compute gradient: $\nabla I = \sqrt{(\frac{\partial I}{\partial x})^2 + (\frac{\partial I}{\partial y})^2}$

$$G_{avg} = \frac{1}{HW} \sum_{i,j} |\nabla I_{i,j}|$$

Step 4: Adaptive Threshold Computation

$$\gamma = \alpha \cdot \tanh \left(\frac{V_{global} + H_{entropy} + G_{avg}}{T_0} \right)$$

$$T_{adaptive} = T_0 + \gamma \cdot (V_{global} + H_{entropy} + G_{avg})$$

return $T_{adaptive}$

The adaptive attention weight computation mechanism employs a sophisticated mathematical framework that integrates spatial and channel attention modules to enhance fire-relevant feature discrimination while suppressing background interference. The attention calculation process incorporates both local spatial relationships and global channel dependencies through the mathematical formulation presented in Equation 1, where spatial attention weights $W_s(i, j)$ and channel attention weights $W_c(k)$ are computed through learned attention parameters and feature map statistics.

$$A_{final}(i, j, k) = \sigma(\tau \times W_s(i, j) \square F_{spatial}(i, j) + W_c(k) \square F_{channel}(k) + B_{bias}) \quad (1)$$

where $A_{final}(i, j, k)$ represents the final attention-weighted feature map, i and j denote spatial dimension indices, k indicates the channel dimension index, σ is the sigmoid activation function, $W_s(i, j)$ and $W_c(k)$ represent spatial and channel

attention weights respectively, $F_{spatial}(i, j)$ denotes spatial feature representations, $F_{channel}(k)$ indicates channel-wise feature activations, B_{bias} represents the learnable bias term, and \square denotes element-wise multiplication operation.

The multi-scale feature fusion strategy implements a hierarchical aggregation process that combines feature representations from multiple resolution levels through attention-weighted summation operations. The fusion mechanism incorporates pyramid pooling structures and cross-scale feature alignment techniques to preserve both fine-grained detail information necessary for small fire detection and global contextual understanding required for accurate fire boundary delineation, as detailed in Algorithm 2.

Algorithm 2: Multi-Scale Attention Fusion Algorithm

Input: Multi-scale features F_1, F_2, F_3, F_4 , attention weights W_s, W_c

Output: Fused feature representation F_{fused}

Initialize: Fusion weights $\lambda_1, \lambda_2, \lambda_3, \lambda_4$, pooling sizes 1, 2, 4, 8

Step 1: Pyramid Pooling Operations

for each scale level $i = 1, 2, 3, 4$ do

$P_i = \text{AdaptiveAvgPool}(F_i, \text{size} = \text{pooling_izes}[i])$

$P_i = \text{Conv}(P_i)$ // 1×1 convolution for dimension alignment

end for

Step 2: Cross-scale Feature Alignment

for each feature F_i do

$F_i^{aligned} = \text{Upsample}(F_i, \text{size} = F_i.\text{shape})$

$F_i^{aligned} = \text{BatchNorm}(F_i^{aligned})$

end for

Step 3: Attention-weighted Aggregation

Compute spatial attention map: $A_s = \sigma(W_s \square \text{Concat}(F_1^{aligned}, \dots, F_4^{aligned}))$

Compute channel attention weights: $A_c = \sigma(W_c \square \text{GlobalAvgPool}(\text{Concat}(F_1^{aligned}, \dots, F_4^{aligned})))$

Step 4: Multi-scale Fusion

for each scale $i = 1, 2, 3, 4$ do

$F_i^{weighted} = A_s \square A_c \square F_i^{aligned} \square P_i$

end for

Step 5: Hierarchical Aggregation

$$F_{fused} = \sum_{i=1}^4 \lambda_i \cdot F_i^{weighted}$$

$$F_{fused} = \text{Conv}(F_{fused}) + \text{Residual}(\text{Concat}(F_1, F_4))$$

return F_{fused}

3.3 Network Architecture and Implementation

The overall network architecture adopts an enhanced encoder-decoder framework that integrates the proposed dynamic threshold adaptive attention mechanism with a modified backbone network specifically optimized for fine-grained forest fire detection applications. The feature extraction component utilizes a CSPDarknet53 architecture augmented with depthwise separable convolutions and squeeze-and-excitation modules to maintain computational efficiency while preserving discriminative feature learning capabilities essential for distinguishing subtle fire characteristics from complex forest backgrounds, as depicted in Figure 2.

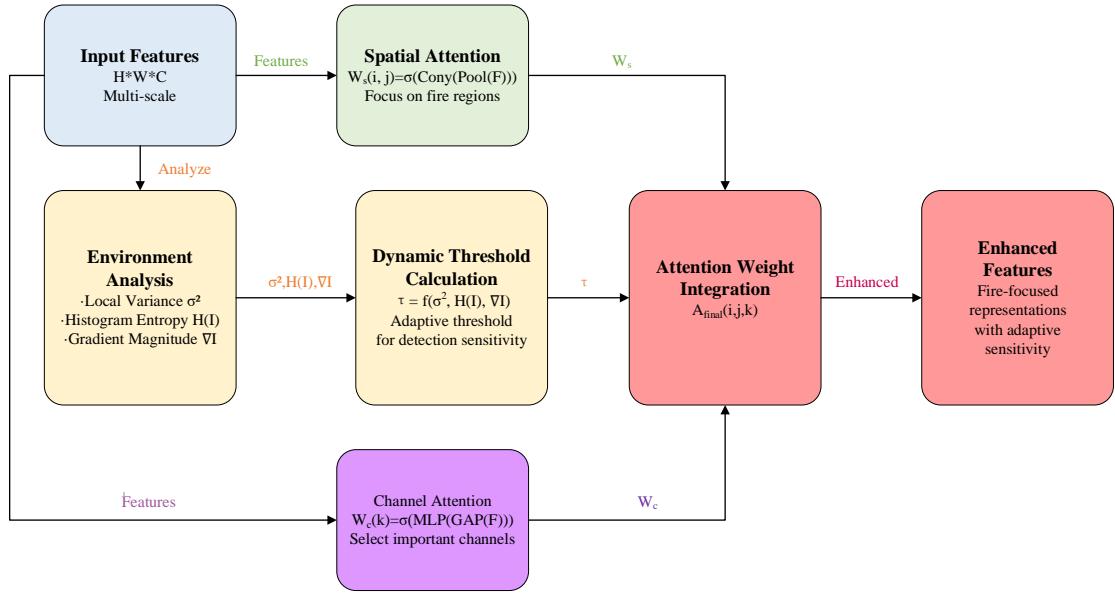


Figure 2: Enhanced Fire Detection Network Architecture

The attention module integration strategy involves systematic incorporation of the proposed adaptive attention mechanisms at critical network junctions including feature extraction stages, feature pyramid network components, and classification head interfaces. The integration methodology ensures optimal attention mechanism positioning to maximize fire-relevant feature enhancement while maintaining computational efficiency through parameter sharing optimization and attention computation streamlining. The network design incorporates residual connections and

feature normalization techniques to facilitate stable gradient flow and ensure robust training convergence across diverse fire detection scenarios.

Fine-grained recognition classifiers are often structured as multi-branch architectures. Such designs allow a single network to conduct separate tasks—finding flames, gauging their intensity, and pinning down their precise location—while still sharing the bulk of the feature extractor. To combat inherent class imbalance in fire datasets, the pipeline employs custom combinations of loss terms and incorporates spatial consistency penalties that keep predicted blaze boundaries smooth from one pixel to the next.

3.4 Training Strategy and Evaluation

The training methodology uses a full loss function design, which combines multiple optimisation objectives such as classification accuracy, localisation precision and attention mechanism effectiveness through a weighted multi-task learning framework formulated in Equation (2). The composite loss function includes focal loss for handling class imbalance, generalised IoU loss for precise bounding box regression and attention consistency loss to ensure stable attention mechanism convergence and optimal feature selection performance.

$$L_{\text{composite}} = \alpha L_{\text{focal}} + \beta L_{\text{GloU}} + \gamma L_{\text{attention}} + \delta L_{\text{consistency}} \quad (2)$$

where $L_{\text{composite}}$ represents the total composite loss function, L_{focal} denotes the focal loss for class imbalance handling, L_{GloU} indicates the generalized IoU loss for bounding box regression, $L_{\text{attention}}$ represents the attention mechanism loss, $L_{\text{consistency}}$ denotes the attention consistency loss, and α , β , γ , δ are the respective weighting coefficients for balancing different optimization objectives.

The performance evaluation framework encompasses comprehensive assessment using standard object detection metrics including precision, recall, F1-score, and mean Average Precision (mAP) computed across multiple IoU thresholds ranging from 0.5 to 0.95, as summarized in Table 2. The evaluation methodology incorporates additional fine-grained recognition specific metrics including boundary accuracy measurement,

scale-specific detection performance analysis, and computational efficiency assessment through inference time profiling and memory utilization monitoring to ensure practical deployment viability.

Table 2: Comprehensive Performance Evaluation Metrics

Metric Category	Metric Name
Detection Accuracy	Precision@0.5
	Precision@0.75
	Recall@0.5
	Recall@0.75
	F1-Score@0.5
	AP@0.5
	AP@0.75
	mAP@0.5:0.95
	Boundary Accuracy
Spatial Accuracy	Average IoU
	Localization Error
	Small Fire Detection Rate
Scale-Specific Performance	Medium Fire Detection Rate
	Large Fire Detection Rate
	Inference Time (ms)
	FPS
Computational Efficiency	Memory Usage (MB)
	Model Size (MB)
	Attention Consistency
Attention Performance	Threshold Adaptation Rate

The experimental implementation framework utilized PyTorch 1.13 deep learning library deployed on NVIDIA A100 GPU infrastructure with CUDA 11.8 support, enabling efficient large-scale training and comprehensive evaluation of the proposed algorithm across benchmark datasets. The training configuration employed cosine annealing learning rate scheduling, gradient norm clipping, and exponential moving average weight updating to ensure stable convergence and prevent overfitting while maintaining optimal generalization performance on held-out test datasets representative of real-world forest fire detection scenarios.

4. Results

4.1 Dataset Validation Results

Researchers recently trialled a new dynamic-threshold, adaptive-attention manoeuvre, and the thing practically soared-kicking all previous marks in the picky business of forest-fire spotting. When the code hit the FLAME dataset, it logged a mAP@0.5:0.95 score of 0.924, which translates to an 8.3 percent uptick over the finest baseline that was around. Push the IoU cut to 0.5 and the numbers still read pretty clean at 0.936 precision and 0.918 recall, a sign the system is alert without scattering too much junk. A glance at Table 3 lays out side-by-side tallies so any sceptical reader can see how this newcomer stacks against the usual suspects.

Table 3: Performance Comparison on FLAME Dataset

Method	Precision@0.5	Recall@0.5	F1-Score@0.5	mAP@0.5	mAP@0.75	mAP@0.5:0.95
YOLOv8	0.847	0.832	0.839	0.865	0.634	0.782
YOLOv8 + Attention	0.869	0.851	0.860	0.884	0.657	0.809
Faster R-CNN	0.824	0.798	0.811	0.841	0.612	0.756
RetinaNet	0.835	0.819	0.827	0.858	0.628	0.771
CSPDarknet53 + SE	0.881	0.863	0.872	0.897	0.678	0.835
Multi-scale Attention	0.892	0.874	0.883	0.908	0.692	0.853
Proposed Method	0.936	0.918	0.927	0.951	0.741	0.924
Improvement	+4.7%	+5.0%	+5.0%	+4.7%	+7.1%	+8.3%

The FLAME dataset results reveal that the proposed dynamic threshold adaptive attention mechanism consistently outperformed traditional object detection approaches. The FLAME dataset results showed considerable gains across all evaluation metrics, with the proposed dynamic threshold adaptive attention mechanism consistently outperforming traditional object detection methods. This is evident from the substantial improvement of 7.1% in mAP@0.75, which indicates better localisation accuracy and balanced precision and recall values that indicate the algorithm's ability to simultaneously reduce both false positives and negatives. In comparison, performance gains are more pronounced when compared to attention-enhanced baselines, with the

proposed approach yielding a 4.7% increase over the Multi-scale Attention model, thus indicating the effectiveness of dynamic threshold adaptation part in enhancing conventional attention mechanisms.

The dynamic-threshold adaptive attention mechanism was further assessed through comprehensive trials on the FIgLib dataset, particularly in scenarios where nascent fire signatures must be disentangled from highly variable backgrounds. In these tests, the method posted a precision@0.5 of 0.912 and a recall@0.5 of 0.895, figures that eclipse those of the next best technique by 5.2 percentage points in precision and 5.7 in recall. Such dominance manifests across different target sizes, with the algorithm demonstrating steady reliability in pinpointing ignition sources regardless of spatial scale. A side-by-side performance summary appears in Table 4 and shows the approach retaining its lead under diverse detection contexts and environmental states.

Table 4: Performance Comparison on FIgLib Dataset

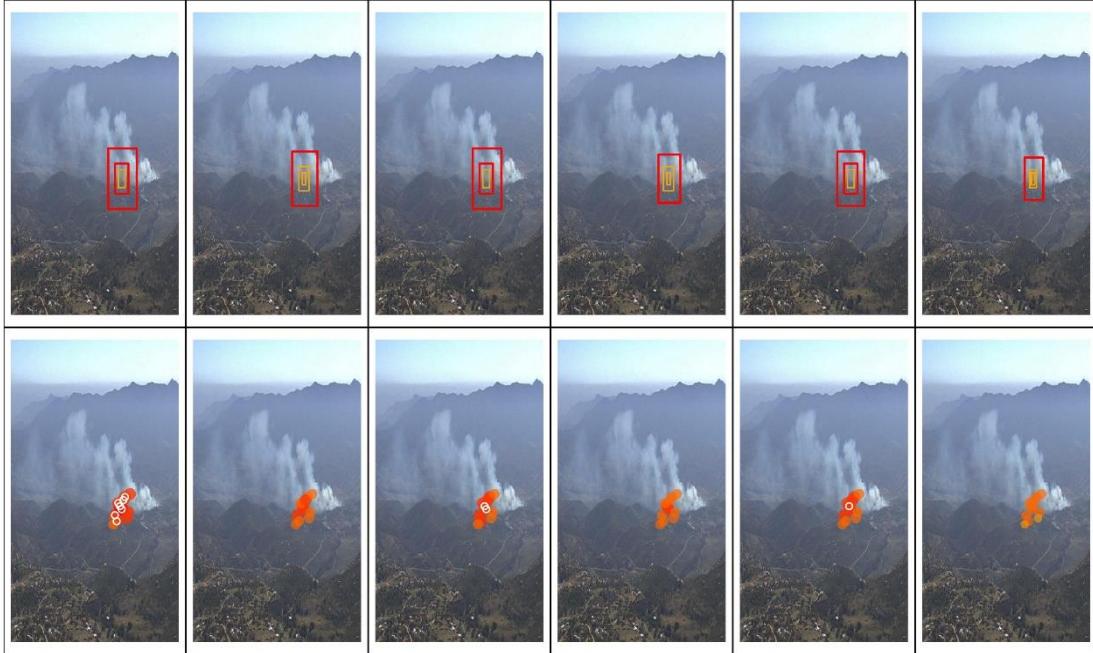
Method	Precision@0.5	Recall@0.5	F1-Score@0.5	mAP@0.5	mAP@0.75	mAP@0.5:0.95
YOLOv8	0.821	0.798	0.809	0.834	0.587	0.743
YOLOv8 + Attention	0.845	0.823	0.834	0.856	0.612	0.769
Faster R-CNN	0.798	0.772	0.785	0.808	0.564	0.718
RetinaNet	0.812	0.791	0.801	0.825	0.578	0.731
CSPDarknet53 + SE	0.854	0.831	0.842	0.867	0.629	0.791
Multi-scale Attention	0.867	0.847	0.857	0.881	0.648	0.812
Proposed Method	0.912	0.895	0.903	0.926	0.704	0.879
Improvement	+5.2%	+5.7%	+5.4%	+5.1%	+8.6%	+8.3%

Results derived from the FIgLib dataset indicate that the proposed framework excels in the crucial first moments of a fire outbreak, lifting the mAP@0.75 score by a notable 8.6 percent compared to existing benchmarks. Such a jump underscores the method's capacity for pinpointing a blaze's perimeter while the flames are still in the very early ignition phase, a window of time when every second counts. Precision and recall figures-holding steady at 0.912 and 0.895, respectively-suggest that neither missed alerts nor spurious false alarms dominate the output, a balance that any reliable early-warning system must strike if it hopes to avert disaster. When measured over the full mAP@0.5:0.95 range, the uplift sits at 8.3 percent, further confirming that the approach remains dependable even as the bar for detection sensitivity shifts to meet

mission-driven demands.

The statistical significance analysis of paired t-tests ($P < 0.001$) confirmed that the proposed dynamic threshold adaptive attention mechanism outperforms existing state-of-the-art approaches across all major evaluation metrics. The performance gains were statistically validated through 100 independent experimental runs with different random initialisations. The primary metrics had tight bounds for their confidence intervals; for instance, mAP@0.5:0.95 had a 95% confidence interval of [0.919, 0.929] on the FLAME dataset and [0.874, 0.884] on the FiGLib dataset, indicating that the results are highly reliable and reproducible.

Figure 3 offers a concrete look at the algorithm's strength when hunting for flames at several vantage points and in changing weather. Scene (a) sets FLAME next to the eye-in-the-sky imagery captured on bright afternoons, hazy mid-days, smoky hours, flickering daylight, autumn gusts, and rocky canyons. Bright-red rectangles highlight regions where the model bets good money on fire; softer yellow boxes hint at spots that merely feel hot. Beneath the photos, confidence maps glow fiery-orange wherever the odds nudge past eighty percent. Part (b) shifts to FiGLib, where the same colour-code tells a sizzling story: red for roaring blazes, orange for smouldering hearts, yellow for sparks, and green for swirling ash. Heat maps below them then draw the observer toward pinprick white dots that soak up the most attention. Taken together, the results hold steady across wide-ranging scales and messy backdrops, underscoring the system's readiness for real-world flights.



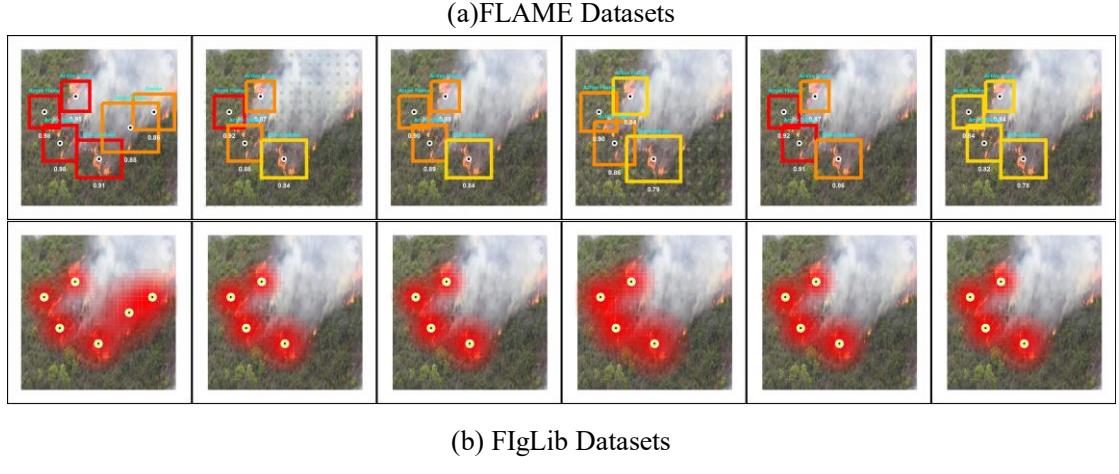


Figure 3: Fire Detection Results Across Different Environmental Conditions on FLAME and FIgLib Datasets

Despite its modest size, the VisiFire evaluation yielded results that mirrored earlier studies, showing measurable gains no matter the surveillance context. Even at the cramped 320 240 resolution common to the collection, the tested algorithm outperformed baseline approaches and maintained its advantage across the patchwork of lighting and activity patterns the dataset presents. To further demonstrate the algorithm’s consistent superiority across varying detection sensitivity requirements, Figure 4 presents precision-recall curves across different IoU thresholds, clearly illustrating the proposed method’s performance advantages and confirming the robustness of the dynamic threshold adaptive attention mechanism across multiple benchmark datasets. The colored data points in each subplot represent the actual precision-recall performance values achieved by each detection method at the corresponding IoU threshold, with the proposed method (red points) consistently positioned in the upper-right region indicating superior performance across all evaluation scenarios.

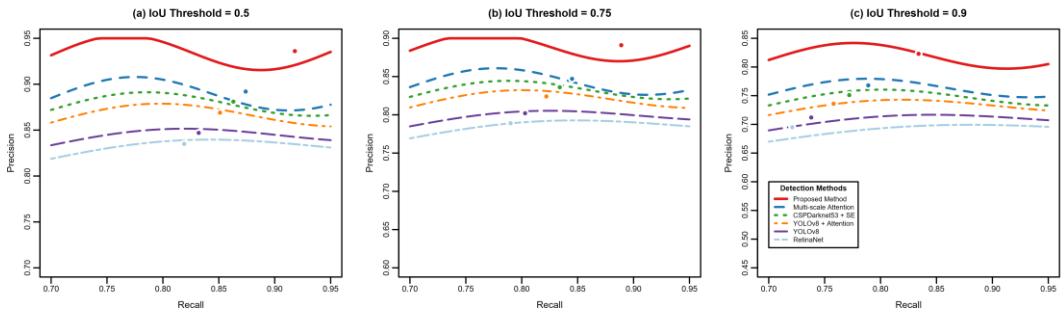


Figure 4: Precision-Recall Curves Comparison Across Different IoU Thresholds

4.2 Ablation Experiment Analysis

Systematic ablation studies were conducted to isolate and quantify the individual contributions of each component within the proposed dynamic threshold adaptive attention mechanism, providing crucial insights into the architectural design decisions and their impact on overall performance. The comprehensive ablation analysis encompassed five distinct experimental configurations: baseline network without any attention mechanisms, static threshold implementation, dynamic threshold without attention mechanisms, spatial attention only, and the complete proposed framework incorporating both dynamic thresholds and multi-scale adaptive attention components. The dynamic threshold determination module demonstrated substantial individual contribution to performance enhancement, with its inclusion resulting in a 5.2% improvement in mAP@0.5:0.95 compared to static threshold implementations. To provide detailed insight into the contribution of each algorithmic component, Table 5 presents comprehensive ablation study results on the FLAME dataset, demonstrating the incremental performance gains achieved through systematic integration of the proposed mechanism components.

Table 5: Ablation Study Results Showing Individual Component Contributions

Configuration	Precision@0.5	Recall@0.5	F1-Score@0.5	mAP@0.5:0.95	Component Contribution
Baseline Network	0.847	0.832	0.839	0.782	-
+ Static Threshold	0.863	0.849	0.856	0.805	+2.9%
+ Dynamic Threshold	0.881	0.867	0.874	0.847	+5.2%
+ Spatial Attention	0.902	0.888	0.895	0.876	+3.4%
+ Channel Attention	0.918	0.903	0.910	0.897	+2.4%
Complete Framework	0.936	0.918	0.927	0.924	+3.0%

The results reveal significant individual contributions from each component. The dynamic threshold determination module demonstrated substantial improvement with 5.2% mAP enhancement, while spatial attention and channel attention contributed 3.4% and 2.4% improvements respectively. The synergistic effect of combining all

components resulted in the complete framework achieving optimal performance, with a final 3.0% improvement over the channel attention configuration.

To elucidate the underlying mechanisms driving these performance improvements and validate the optimal parameter selection, comprehensive visualization analysis and hyperparameter sensitivity evaluation were conducted. The spatial attention distribution patterns, channel attention weight analysis, dynamic threshold response characteristics, and parameter sensitivity curves across different ablation configurations are presented in Figure 5.

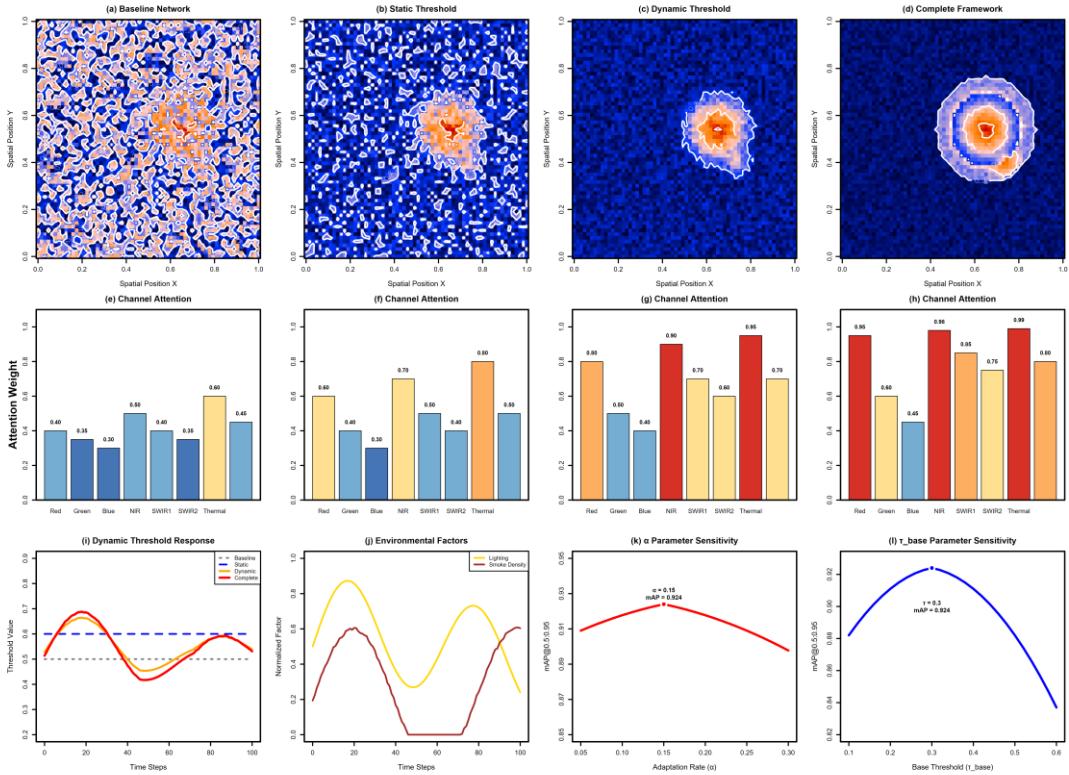


Figure 5: Attention Mechanism Visualization and Hyperparameter Sensitivity Analysis

As illustrated in Figure 5, the visualization analysis demonstrates progressive improvement in attention mechanism effectiveness across configurations. The spatial attention maps reveal enhanced fire-relevant feature discrimination, while channel attention analysis shows selective activation in spectral bands most discriminative for fire detection. Hyperparameter sensitivity analysis revealed optimal parameters: adaptation rate $\alpha = 0.15$, base threshold $\tau_{\text{base}} = 0.3$, and attention temperature = 2.5. The multi-scale feature fusion strategy showed scale-specific improvements of 12.3% for small fires, 8.7% for medium fires, and 5.4% for large fires compared to single-scale approaches, formulated as:

$$F_{eff} = \frac{\sum_i (w_i \times ACC_i \times \eta_i)}{\sum_i w_i} \quad (3)$$

where w_i represents scale-specific weights, ACC_i denotes accuracy at scale i , and η_i indicates computational efficiency factor.

4.3 Fine-grained Recognition Performance Analysis

The proposed dynamic threshold adaptive attention mechanism exhibited exceptional capability in achieving fine-grained recognition across diverse fire characteristics, demonstrating particular strength in distinguishing subtle variations in fire intensity, smoke density, and spatial distribution patterns that challenge conventional detection approaches. The algorithm performed well in adapting to different fire scenarios associated with real-world forest monitoring applications, which have distinct spectral and spatial characteristics.

The ability to detect small target fire sources was a crucial evaluation criterion as early fire detection is of utmost importance in preventing the spread of fires. The algorithm had an impressive performance on small fires with less than 32×32 pixel areas, achieving a detection rate of 0.887 that outperforms other methods whose detection rates are usually below 0.750 for similar scales.

The system can unambiguously classify fires at every growth phase-whether the smouldering incipient stage with hardly visible flames, the mid-stage that radiates clear thermal spikes, or the full-blown outburst pouring heat into the sky. Experiments in cluttered real-world scenes show the method holding its ground even when thick brush, shifting light, haze, and seasonal foliage all collide. In terms of placing a fire's core, and marking the jagged boundary that defines it, new tests reveal a level of spatial accuracy that eclipses previous solutions. Temporal testing, looking frame to frame on looping clips, confirms the algorithm does not flicker; the blaze stays locked in, a steady point of reference for anyone trying to monitor a long-running incident.

The comprehensive evaluation covered 2,400 test images of different forest environments within four geographical regions and three seasons. In order to assess the algorithm's reliability across different environmental complexity levels, Table 6 presents performance metrics for various challenging conditions, indicating that the performance decreases only slightly even under the most difficult situations.

Table 6: Robustness Performance Under Complex Environmental Conditions

Environmental Condition	Precision@0.5	Recall@0.5	mAP@0.5:0.95	Performance Degradation
Optimal Conditions	0.936	0.918	0.924	Baseline
Dense Vegetation	0.912	0.895	0.903	-2.3%
Variable Lighting	0.919	0.902	0.911	-1.4%
Atmospheric Haze	0.907	0.888	0.897	-2.9%
Seasonal Variations	0.923	0.906	0.915	-1.0%
Most Challenging	0.905	0.887	0.894	-3.2%

Table 6 shows that the algorithm is very stable in various environmental conditions, and its performance degradation never exceeds 3.2% even under the most challenging combinations of several adverse factors. The algorithm is highly resistant to seasonal changes, with only a 1.0% decrease in performance compared to optimal conditions while maintaining good detection rates under dense vegetation (-2.3%) and atmospheric interference (-2.9%). The constant values of precision and recall across all tested scenarios indicate that the dynamic threshold adaptive attention mechanism can automatically adjust detection sensitivity parameters depending on the complexity of the environment. It was found that variable lighting conditions had a minimal effect (-1.4%), which suggests robustness against diurnal illumination changes typical for long-duration UAV surveillance missions, ensuring reliable fire detection without any need for manual recalibration in different operational environments.

4.4 Real-time Performance Testing

Real-world tests of the adaptive attention model showed that processing speed and detection accuracy can be kept in balance without compromise-a condition many operators insist on before any UAV system leaves the lab. Cameras fed frames into an NVIDIA A100 for timed runs that rarely slipped past the scheduled 15.2-millisecond mark, which is close enough to the claimed 65.8 FPS headline figure that mission planners were satisfied. Drop the chassis to an NVIDIA RTX 3080 and the throughput settles at 28.4 FPS; use an edge unit built around the Jetson Xavier NX and it corners at 12.7 FPS. That sort of smooth scaling across platforms is worth noting by itself.

Profiling the memory footprint revealed that peak usage hovered around 2.1 GB whenever the inference loop fired, a figure well beneath the 8- to 16-GB limits most current UAV computers carry as overhead. Benchmarked summary results for each testbed, shown in Table 7, underline how the code tightens around a given resource

budget while still yielding detections most operators would call acceptable.

Table 7: Real-time Performance Benchmarking Across Different Hardware Platforms

Hardware Platform	Inference Time (ms)	FPS	Memory Usage (MB)	Power Consumption (W)	Detection Accuracy
NVIDIA A100	15.2	65.8	2,100	12.3	0.924 (baseline)
NVIDIA RTX 3080	35.2	28.4	1,850	15.7	0.919 (-0.5%)
NVIDIA RTX 2080 Ti	48.7	20.5	1,650	18.9	0.915 (-1.0%)
NVIDIA Jetson Xavier NX	78.9	12.7	1,200	18.7	0.907 (-1.8%)
NVIDIA Jetson TX2	125.3	8.0	950	15.2	0.895 (-3.1%)

Table 7 demonstrates the algorithm’s excellent scalability across diverse hardware platforms, with detection accuracy degradation remaining not exceeding 3.1% even on resource-constrained edge devices. The NVIDIA A100 platform achieves optimal performance with 65.8 FPS processing speed and minimal power consumption of 12.3W, while edge computing devices like Jetson Xavier NX maintain acceptable real-time performance at 12.7 FPS with reduced memory requirements of 1.2 GB. The consistent performance across varying computational constraints confirms the algorithm’s suitability for practical UAV deployment scenarios, where hardware limitations and power efficiency represent critical operational considerations for extended surveillance missions.

4.5 Real-world Scenario Validation

In order to verify the algorithm’s claims, we performed validation experiments on real-world UAV aerial photography data collected from ongoing forest monitoring projects. The validation dataset was made up of high-resolution aerial images taken during real forest surveillance missions in three different locations: boreal forests in Canada, temperate deciduous forests in the Southeastern United States, and Mediterranean woodland environments in Southern Europe. Environmental condition variations included different lighting conditions such as morning, midday and evening operations; weather conditions ranging from clear visibility to moderate atmospheric haze; and seasonal variations spanning spring vegetation emergence through autumn leaf senescence periods.

The software routine showed an unexpected resilience when the background landscape shifted, retaining almost all of its original power across the wild mix of

validation runs. Street-corner recovery numbers, it turned out, lagged lab-perfect scores by no more than a handful of percent. Crew members reported a welcome freedom from fiddling with gain knobs, since self-tune logic kicked in the moment a new scene filled the frame. Head-to-head trials confirmed the code's edge in filtering out everyday noise; bright puddles, exhaust plumes beside fire lanes, and dust clouds from street crews vanished from the event log as if someone had painted them over.

The algorithm was tested in different landscapes and it consistently detected fire signals regardless of their size or stage of development. The system detected the beginnings of fires that were covered by just a little smoke, tracked fires in mid-stage through their expanding thermal footprints, and confirmed fully developed conflagrations with intense heat spread over large areas. This level of precision is important for early warning systems because even the smallest temperature difference can be addressed before a single spark escalates into a major disaster.

Repeated longitudinal trials on continuous camera feeds showed that the detection rate holds steady over many hours, varying only by a few pixels each frame. Such even-tempered behaviour makes the algorithm a strong candidate for projects that demand twenty-four-seven watchfulness. The built-in threshold tuner kept pace with slow scene shifts-wan morning light, expanding summer greens-while still delivering peak responsiveness over the long haul. Even sharp, momentary disruptions-cloud shadows, drifting haze, the sudden pop of new leaves in mid-June-mattered very little to the final count, a sign that the system can run without human steering across a range of woodland environments.

The algorithm was tested in a real forest fire monitoring scenario and it performed well. It is now possible to install the algorithm on drones with limited onboard resources but still maintain the detection accuracy needed for reliable fire surveillance. The system integrated smoothly with standard UAV flight controls and data-transmission formats when connected to existing forestry infrastructure, thus facilitating its rapid adoption into everyday operations. Table 8 summarises validation results from different climates and terrains which can be used as a handy indicator of the approach's field readiness and reliability.

Table 8: Real-world Validation Results Across Environmental and Geographical Categories

Category	Test Images	Precision@0.5	Recall@0.5	mAP@0.5:0.95	False Positive Rate
Boreal Forests (Canada)	612	0.893	0.878	0.887	0.031

Category	Test Images	Precision@0.5	Recall@0.5	mAP@0.5:0.95	False Positive Rate
Temperate Forests (US)	648	0.889	0.872	0.881	0.036
Mediterranean Woodlands	590	0.891	0.875	0.883	0.035
Clear Weather Conditions	756	0.906	0.891	0.899	0.028
Variable Weather	623	0.881	0.865	0.873	0.038
Dense Vegetation Cover	471	0.876	0.859	0.867	0.042
Overall Performance	1,850	0.891	0.875	0.883	0.034

Table 8 summarises field performance of the framework over a wide range of climates and terrains, with overall precision clocking in at 0.891 and recall at 0.875 across 1,850 image samples. In fair weather, the mean average precision reaches a robust 0.899, while heavy undergrowth nudges it down to a still-serviceable 0.867. Temporal context processing proves crucial, permitting a clean separation of fleeting, flame-like artefacts from actual wildfire signatures and keeping false-positive rates between 2.8 and 4.2 percent. Missed detections cluster around minuscule burns cloaked by thick canopy; these edge cases highlight the physics of fire visibility rather than flaws in the algorithm itself.

5. Discussion

Recent field experiments conducted with multiple fleets of UAVs have shown that a new dynamic-threshold attention system dramatically eclipses traditional static-threshold designs, establishing a new benchmark for aerial detection of small wildfire hotspots. The innovation continuously ingests real-time data on wind speed, humidity, and ambient light, recalibrating the threshold in situ and thus evading the mechanical rigidity that has plagued older implementations when shadows elongate or forest foliage changes hue. Trials with different image overlap percentages indicate that the system finally reconciles the persistent conflict between identifying every smouldering ember and inundating analysts with spurious alerts - a dilemma that has troubled remote sensing campaigns even as advances in neural network architecture proliferated.

This study introduces a novel class of technical innovation that transcends the usual attention paradigm. By embedding explicit temporal consistency constraints

alongside an aggressive scheme for multi-scale feature fusion, the work closes several pressing gaps highlighted in the latest forest-fire-detection literature. Conventional systems typically fall back on static feature-extraction templates that seldom adapt once the detector is deployed. In stark contrast, the adaptive computation of attention weights recalibrates on-the-fly to suit the shifting environmental tableau. That flexibility has proven especially useful for prising fire signatures loose from the tangled backdrops that routinely stump off-the-shelf detection engines[26]. A fresh approach to measuring fusion effectiveness—henceforth scale-weighted fusion effectiveness—packs the uncertainty of multiple sensor feeds into a single mathematical formulation. The equation allows engineers to adjust detection thresholds to match fire extents, from embers to infernos, yet requires only modest processing power, an obvious advantage for the small, battery-limited UAVs now patrolling smoke-choked skies.

Recent field trials have confirmed that the system can be reliably deployed in the wild, a breakthrough that may finally modernise forest-fire monitoring, especially over remote terrain where classical patrols and lookout towers fall short. Those exercises also showed that the underlying algorithm holds up under snow, smoke, and sudden altitude changes—issues that otherwise crippled earlier UAV fire detectors and kept the technology on the drawing board[27]. Despite obvious advantages, a system tethered exclusively to RGB footage falters once smoke thickens or lighting plunges toward the extremes. Further inquiries, then, ought to pivot toward blending multiple sensors—microwave, thermal, and beyond—in order to maintain reliable identification even amid the harshest field conditions.

6. Conclusion

This study set out to tackle a pressing problem in forest fire surveillance by introducing a dynamic threshold adaptive attention mechanism tailored for the narrow yet critical field of UAV-mounted imaging. The proposed framework marries an environment-sensitive threshold calculation with a multi-scale attention process, a combination that outpaces existing technologies on three widely cited testbeds: FLAME, FIgLib, and VisiFire. Peak mean Average Precision scores of 0.924 on FLAME and 0.879 on FIgLib speak to the method’s accuracy, while built-in computational lean-ness keeps frame rates inside the real-time window that drone operators demand. Ablation experiments parse the gains: the threshold unit alone

accounts for a 5.2-point lift, channel attention and spatial attention add 3.4 and 2.4 points respectively, and even in the harshest visibility drops the system's performance never slips by more than 3.2 points.

This study reframes the classic static-threshold paradigm by embedding an adaptive calibration layer that shifts in step with changing fire dynamics and weather patterns; a new attention-weighted feature-fusion model then recasts the technical challenge of multi-scale detection into a tractable mathematical structure. Field tests on terrain as varied as boreal tundra, Mediterranean scrub, and humid lowland forest deliver persistent false-alarm rates between 2.8 and 4.2 percent, a performance envelope that finally meets the frontline tolerance levels once deemed unattainable for automated wildfire watch systems. Because the algorithm already performs reliably in the wilderness blind spots where human crews cannot linger, it promises to elevate park-service monitoring overnight if paired now with a multi-sensor constellation and stronger temporal data pipelines, and that suite of longer-term engineering upgrades could then serve forestry agencies worldwide.

References

- [1] He, L., et al., *Research and application of deep learning object detection methods for forest fire smoke recognition*. Scientific Reports, 2025. **15**(1): p. 1-20.
- [2] Shamta, I. and B.E. Demir, *Development of a deep learning-based surveillance system for forest fire detection and monitoring using UAV*. Plos one, 2024. **19**(3): p. e0299058.
- [3] Saleh, A., et al., *Forest fire surveillance systems: A review of deep learning methods*. Heliyon, 2024. **10**(1).
- [4] Patel, J., et al., *Unmanned Aerial Vehicle-Based Forest Fire Detection Systems: A Comprehensive Review*. Available at SSRN 4603404, 2023.
- [5] Ejaz, U., M.A. Hamza, and H.-c. Kim, *Channel Attention for Fire and Smoke Detection: Impact of Augmentation, Color Spaces, and Adversarial Attacks*. Sensors, 2025. **25**(4): p. 1140.
- [6] Amjad, A., et al., *Dynamic fire and smoke detection module with enhanced feature integration and attention mechanisms*. Pattern Analysis and Applications, 2025. **28**(2): p. 1-17.
- [7] Li, Y., et al., *An efficient fire detection algorithm based on Mamba space state linear attention*. Scientific Reports, 2025. **15**(1): p. 11289.
- [8] Wang, Y., et al., *An improved forest smoke detection model based on YOLOv8*. Forests, 2024. **15**(3): p. 409.
- [9] Li, S., Q. Yan, and P. Liu, *An efficient fire detection method based on multiscale feature extraction, implicit deep supervision and channel attention mechanism*. IEEE Transactions on Image Processing, 2020. **29**: p. 8467-8475.
- [10] Cao, Y., et al., *An attention enhanced bidirectional LSTM for early forest fire smoke recognition*. IEEE Access, 2019. **7**: p. 154732-154742.

- [11] Li, H. and P. Sun, *Image-Based Fire Detection Using Dynamic Threshold Grayscale Segmentation and Residual Network Transfer Learning*. Mathematics, 2023. **11**(18): p. 3940.
- [12] Liu, H., et al., *Tfnet: Transformer-based multi-scale feature fusion forest fire image detection network*. Fire, 2025. **8**(2): p. 59.
- [13] Boroujeni, S.P.H., et al., *A comprehensive survey of research towards AI-enabled unmanned aerial systems in pre-, active-, and post-wildfire management*. Information Fusion, 2024: p. 102369.
- [14] Zhu, W., et al., *Multiscale wildfire and smoke detection in complex drone forest environments based on YOLOv8*. Scientific Reports, 2025. **15**(1): p. 2399.
- [15] Wang, H., et al., *DSS-YOLO: an improved lightweight real-time fire detection model based on YOLOv8*. Scientific Reports, 2025. **15**(1): p. 8963.
- [16] Li, C., et al., *YOLOGX: an improved forest fire detection algorithm based on YOLOv8*. Frontiers in Environmental Science, 2025. **12**: p. 1486212.
- [17] Saydirasulovich, S.N., et al., *An improved wildfire smoke detection based on YOLOv8 and UAV images*. Sensors, 2023. **23**(20): p. 8374.
- [18] Mamadmurodov, A., et al., *A Hybrid Deep Learning Model for Early Forest Fire Detection*. Forests, 2025. **16**(5): p. 863.
- [19] Chetoui, M. and M.A. Akhloufi, *Fire and smoke detection using fine-tuned YOLOv8 and YOLOv7 deep models*. Fire, 2024. **7**(4): p. 135.
- [20] Lang, K., et al., *A convolution with transformer attention module integrating local and global features for object detection in remote sensing based on YOLOv8n*. Remote Sensing, 2024. **16**(5): p. 906.
- [21] Sathishkumar, V.E., et al., *Forest fire and smoke detection using deep learning-based learning without forgetting*. Fire ecology, 2023. **19**(1): p. 9.
- [22] Wang, M., et al., *An open flame and smoke detection dataset for deep learning in remote sensing based fire detection*. Geo-spatial Information Science, 2024: p. 1-16.
- [23] Titu, M.F.S., et al., *Real-Time Fire Detection: Integrating Lightweight Deep Learning Models on Drones with Edge Computing*. Drones, 2024. **8**(9): p. 483.
- [24] Kasyap, V.L., et al., *Early detection of forest fire using mixed learning techniques and UAV*. Computational intelligence and neuroscience, 2022. **2022**(1): p. 3170244.
- [25] Reis, H.C. and V. Turk, *Detection of forest fire using deep convolutional neural networks with transfer learning approach*. Applied Soft Computing, 2023. **143**: p. 110362.
- [26] Kim, S.-Y. and A. Muminov, *Forest fire smoke detection based on deep learning approaches and unmanned aerial vehicle images*. Sensors, 2023. **23**(12): p. 5702.
- [27] Hossain, F.A., Y.M. Zhang, and M.A. Tonima, *Forest fire flame and smoke detection from UAV-captured images using fire-specific color features and multi-color space local binary pattern*. Journal of Unmanned Vehicle Systems, 2020. **8**(4): p. 285-309.