WWW.IJRPR.COM

International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

Bayesian Hyperparameter Optimization of a CBAM–BiFPN-Enhanced YOLOv11-Nano Architecture for Fire and Smoke Detection

Gwamzhi Ponsah Emmanuel*, Prof.Jibril Danladi Jiya*, Prof. Ya'u Haruna Shuaibu*, Dr. Hassan Aliyu*

*Department of Electrical and Electronics Engineering, ATBU, Bauchi, Nigeria

ABSTRACT :

Vision-based fire and smoke detection systems must balance detection accuracy with the tight computational budgets of edge deployment. This paper presents an enhanced YOLOv11-Nano architecture in which a Convolutional Block Attention Module (CBAM) is embedded within every C2f block of the backbone and fused into a Bidirectional Feature Pyramid Network (BiFPN) neck, so that channel- and spatial-attention refinement operates at both the feature-extraction and multi-scale fusion stages. Because architectural additions of this kind introduce new sensitivity to training hyperparameters, a 19-parameter Bayesian hyperparameter optimization (HPO) search was conducted with Optuna across learning-rate, regularization, loss-weighting, and augmentation dimensions, using a resilient, checkpoint-resumable trial-management framework built around a persistent SQLite study store. On a curated corpus of 20,523 fire, smoke, and negative images (25,000 bounding boxes per class), the optimal configuration (AdamW, lr0 = 0.0006, box = 0.0658, cls = 0.809, mosaic = 0.948) achieved a validation mAP50 of 0.8118 across ten completed trials, an improvement over the weakest trial by more than 82% and over the median trial by roughly 8%. When the selected configuration was used to train the full CBAM–BiFPN architecture to convergence over 400 epochs at 416×416 resolution, the model reached a baseline mAP50 of 0.8076, precision of 0.8219, recall of 0.7348, and F1-score of 0.7759, with training and validation loss curves that tracked one another closely, indicating limited overfitting. These results establish the CBAM–BiFPN–YOLOv11n baseline and its optimized training recipe as a reusable starting point for downstream refinement stages such as iterative hard-example mining and post-training quantization.

Keywords*: YOLOv11; CBAM; BiFPN; hyperparameter optimization; Optuna; Bayesian optimization; fire detection; object detection*

1. Introduction

Fire outbreaks remain among the most destructive hazards faced by both urban and rural communities, and their consequences are magnified wherever early-detection infrastructure is inadequate (Ahmed et al., 2023; Khan et al., 2022). Conventional point sensors—smoke detectors, heat sensors, and flame detectors—respond only after physical changes reach the sensor, and their reliability degrades in large or open spaces and in the presence of steam, dust, or industrial fumes (Khan et al., 2022; Xu et al., 2021). Vision-based fire detection using deep object detectors addresses these limitations by monitoring wide fields of view directly from camera feeds and by recognizing visual fire and smoke signatures before they reach a physical sensor (Zhang et al., 2023).

Within the family of single-stage detectors, YOLO models are attractive for fire detection because they trade little accuracy for substantial gains in inference speed, and the Nano-scale variants are specifically sized for constrained edge hardware such as the NVIDIA Jetson Nano, Raspberry Pi, and Google Coral TPU (Li et al., 2022; Ahmed et al., 2023). YOLOv11-Nano, the architecture adopted in this study, introduces anchor-free detection and refined feature-extraction blocks (Wang et al., 2024), but its baseline configuration is not tuned for the specific visual characteristics of fire and smoke, diffuse, texture-poor, variably coloured regions that are easily confused with steam, fog, or artificial lighting.

Two complementary levers are available to close this gap: architectural enhancement and hyperparameter tuning. This paper addresses both jointly. First, a Convolutional Block Attention Module is integrated into every backbone stage and into each fusion node of a Bidirectional Feature Pyramid Network neck, so that channel and spatial attention refine features both before and during multi-scale fusion. Second, because attention modules alter the loss landscape and the relative importance of localization, classification, and augmentation hyperparameters, a systematic Bayesian hyperparameter search is conducted over the enhanced architecture rather than reusing hyperparameters tuned for a vanilla YOLOv11n. The objective of this paper is narrower than that of the wider dissertation from which it is drawn: it isolates and reports the architecture and the hyperparameter-optimization stage on their own terms, independent of the hard-negative-mining and quantization stages reported elsewhere, so that the architectural and tuning contribution can be

evaluated and reused independently.

2. Related Work

Attention mechanisms have been widely used to sharpen lightweight fire-detection backbones. Deng et al. (2025) integrated CBAM into a MobileNetV2 backbone (MV2-CBAM) for UAV-based forest-fire classification, reporting improved feature focus but restricting the task to binary classification rather than localization. Pan et al. (2024) proposed EFA-YOLO, combining efficient attention convolution and attention-guided downsampling to obtain a 1.4M-parameter fire/flame detector, but without a smoke class and on a modest 2,060-image evaluation set. He et al. (2024) combined a RepViT transformer backbone with Dynamic Snake Convolution in FireNet, reporting 80.2% mAP@0.5 at 26.7 ms inference on a custom dataset, while Hao et al. (2022) fused SepViT blocks with a lightweight BiFPN and a Global Attention Mechanism in Light-YOLOv5, obtaining a 3.3% mAP gain with 27.1% fewer parameters. Lan et al. (2025) replaced YOLOv8's C2f module with a FasterNet block in Light-YOLOv8-Flame, achieving parameter reduction alongside modest accuracy gains, though without disclosed baseline values.

Two further points from the wider literature sharpen this study's positioning. First, EFA-YOLO (Pan et al., 2024) shows how far parameter efficiency can be pushed when attention is fused directly into the convolution and downsampling operators (EAConv/EADown) rather than appended as a separate block: it reports 1.4M parameters, 4.6 GFLOPs, and 22.19 ms CPU inference, a 94.6% parameter reduction and 88× speed-up relative to mainstream YOLO baselines. FireNet (He et al., 2024) instead pursues accuracy through a simplified Vision Transformer backbone (RepViT) combined with dynamic snake convolution (DSConv) for curved flame/smoke boundaries, reaching 80.2% mAP@0.5 at 26.7 ms. Relative to both, the present architecture takes a middle path, CBAM plus BiFPN on a still-convolutional YOLOv11n backbone, favouring transfer-learning compatibility and training stability over the maximal compression of EFA-YOLO or the transformer-backbone accuracy ceiling of FireNet. Second, an independent comparative study of YOLOv9t, YOLOv10n, and YOLOv11n on a separate smoke/fire dataset (Alkhammash, 2025) reports YOLOv11n reaching a validation mAP50 of 0.825 (precision 0.803, recall 0.758) after 100 epochs, a very similar operating point to this study's 0.8076 mAP50 baseline after 400 epochs at lower resolution, which corroborates that the HPO-tuned CBAM–BiFPN configuration is competitive with, rather than behind, an unmodified YOLOv11n on comparable fire/smoke data.

These studies converge on two observations that motivate the present work. First, attention and multi-scale fusion consistently improve fire/flame feature discrimination, but most implementations apply attention either in the backbone or in the neck, rather than at both locations in a unified CBAM–BiFPN design that couples channel/spatial refinement with bidirectional cross-scale fusion. Second, few of the reviewed studies report a systematic hyperparameter search specific to the modified architecture; hyperparameters are typically inherited from the base detector's default recipe, even though attention and fusion modules alter gradient dynamics and the relative scale of loss terms. Bayesian hyperparameter optimization frameworks such as Optuna (Feurer & Hutter, 2019) are well established for standard object detectors, but their application to jointly search learning-rate, loss-weighting, and augmentation hyperparameters for an attention-augmented, BiFPN-fused YOLO variant applied to fire and smoke has not been reported in the literature reviewed for this study.

3. Materials and Methods

3.1 Architecture

The backbone follows the YOLOv11n topology (Conv stem, four Conv–C2f stages, and an SPPF block), with every C2f block replaced by a C2f_CBAM variant that appends a CBAM module, sequential channel attention (shared MLP over average- and max-pooled descriptors, reduction ratio 16) followed by spatial attention (7×7 convolution over channel-pooled maps), to the block output. The three backbone feature maps (P3, P4, P5) are each passed through a dedicated CBAM refinement layer before entering the neck. The neck replaces the default PANet-style fusion with a BiFPN block that performs top-down and bottom-up fusion across the three scales, applying a CBAM module after every fusion node (four in total: top-down P4, output P3, output P4, output P5). The refined multi-scale features feed a standard anchor-free Detect head configured for two classes (fire, smoke). The full network was subclassed from Ultralytics' DetectionModel and registered with the Ultralytics task registry so that the custom modules load correctly with pretrained YOLOv11n backbone weights via transfer learning.

3.2 Dataset

A corpus of 20,523 images was assembled from multiple public Kaggle fire/smoke sources (5.5 GB of positive imagery) supplemented with 45,000 negative background images drawn from COCO 2014/2017, from which negatives were subsampled to maintain class balance. Images below 500×400 resolution were discarded. The final corpus comprised 1,531 fire-only, 4,311 smoke-only, 9,681 fire-and-smoke, and 5,000 negative images, curated to

25,000 bounding boxes per class.

The corpus was split 70:15:15 into train, validation, and test partitions at the individual-image level using a fixed random seed, after which every file name was systematically randomized so that no residual naming pattern (e.g., a shared prefix from a particular Kaggle source or capture session) could act as a shortcut feature during training. Images were letterboxed rather than stretched during resizing, so that aspect ratio, and therefore object shape, was preserved identically across all three partitions.

Two caveats on data-leakage control are stated explicitly here, since they are not fully resolved by the measures above. First, the split was performed at the image level rather than at the source-video or capture-session level; where a Kaggle source contributed multiple frames from the same fire event or the same short video clip, near-duplicate frames could in principle appear on both sides of the train/validation or train/test boundary, which would inflate the reported validation and test metrics relative to a strictly held-out scene. Second, no perceptual-hash or embedding-based near-duplicate detection was run across the merged multi-source corpus prior to splitting. Neither omission is unusual in fire-detection datasets assembled from multiple public sources, several of the benchmark datasets used elsewhere in this literature share the same limitation, but it is flagged here as a specific, checkable threat to the validity of the reported mAP50 figures, and resolving it (via perceptual hashing followed by a source-aware, grouped split) is identified as a concrete next step in Section 6 rather than being presented as already solved.

3.3 Hyperparameter Search Space and Optimization Protocol

Hyperparameter optimization used Optuna's Bayesian (Tree-structured Parzen Estimator) sampler to maximize validation mAP50 over 10 completed trials of 50 epochs each at 640×640 resolution. The 19-dimensional search space spanned four categories: learning dynamics (lr0: 1e-5–1e-1 log-scale, lrf: 0.01–0.2, momentum: 0.85–0.98, weight_decay: 1e-5–1e-2 log-scale, warmup_epochs: 1–5, optimizer: {AdamW, SGD}); loss-term weighting (box: 0.02–0.2, cls: 0.2–0.8); colour augmentation (hsv_h: 0–0.1, hsv_s: 0–0.9, hsv_v: 0–0.9); and geometric/compositional augmentation (degrees: 0–20°, translate: 0–0.3, scale: 0.2–0.8, shear: 0–5°, fliplr: 0–0.5, mosaic: 0.5–1.0, mixup: 0–0.5, close_mosaic: 5–20 epochs). The optimization loop was made resilient to interruption: trial state was persisted to a SQLite-backed Optuna study with load-if-exists behaviour, checkpointed trials resumed from their last saved weights, and the study terminated only once exactly ten trials reached the COMPLETE state, rather than after ten attempts, so failed trials did not silently shrink the effective search budget.

Table 1. Hyperparameter search space (abridged).

Category	Parameter	Search range
Learning rate	lr0 / lrf	1e-5–1e-1 (log) / 0.01–0.2
Optimizer	momentum / weight_decay	0.85–0.98 / 1e-5–1e-2 (log)
Loss weights	box / cls	0.02–0.2 / 0.2–0.8
Colour augmentation	hsv_h / hsv_s / hsv_v	0–0.1 / 0–0.9 / 0–0.9
Geometric augmentation	degrees / translate / scale / shear / fliplr	0–20° / 0–0.3 / 0.2–0.8 / 0–5° / 0–0.5
Composition augmentation	mosaic / mixup / close_mosaic	0.5–1.0 / 0–0.5 / 5–20 epochs

3.4 Baseline Training Protocol

Following selection of the optimal trial configuration, the full CBAM–BiFPN architecture was trained from the transferred backbone for 400 epochs at 416×416 resolution with batch size 32, a 4-epoch warm-up, early-stopping patience of 50 epochs, mosaic disabled for the final 30 epochs, and an increased objectness weight (kobj = 8.5) to sharpen background/object discrimination. A checkpoint-based state-management layer preserved the optimizer state across interruptions, allowing the extended run to resume without loss of momentum.

Figure 1. CBAM–BiFPN-enhanced YOLOv11n architecture (data flow)

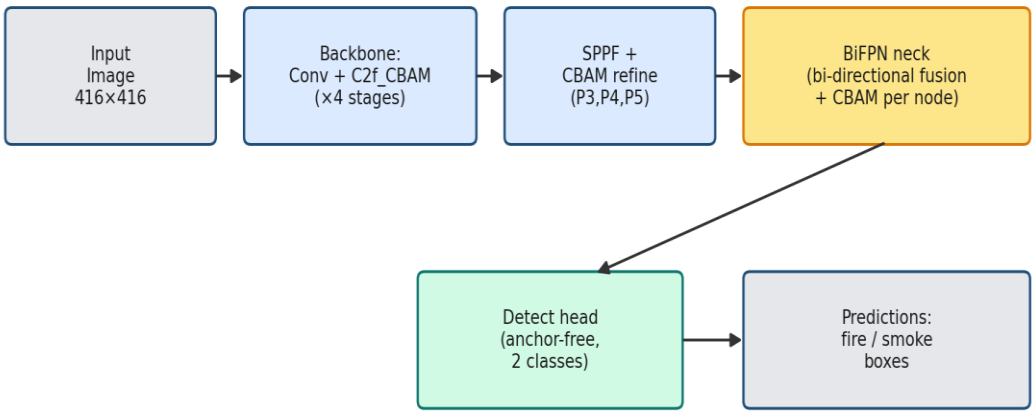


Figure 1. CBAM–BiFPN-enhanced YOLOv11n architecture: backbone, attention refinement, BiFPN neck, and detection head.

4. Results

Across the ten completed Optuna trials, validation mAP50 ranged from 0.4446 (Trial 0) to 0.8118 (Trial 3), a variation of more than 80% attributable entirely to hyperparameter configuration under an otherwise fixed 50-epoch budget. Figure 2 plots mAP50 for all ten trials; Table 2 summarizes the three strongest configurations.

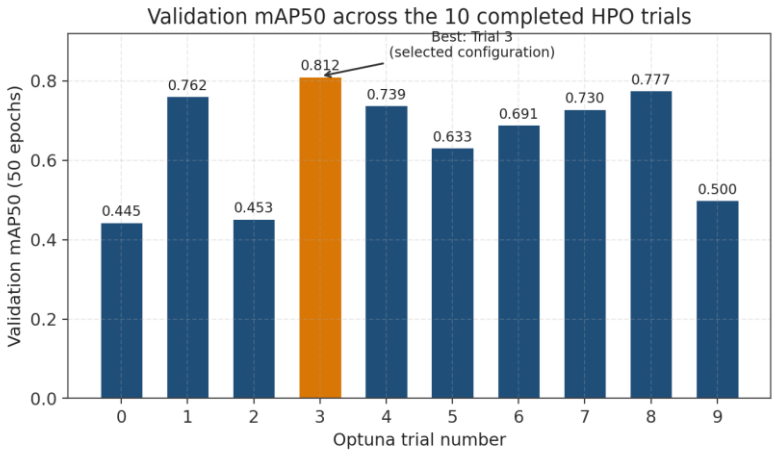


Figure 2. Validation mAP50 achieved by each of the 10 completed Optuna trials.

Table 2. Top hyperparameter optimization trials (validation set, 50 epochs).

Trial	Optimizer	lr0	mosaic	fliplr	mAP50	F1	Precision	Recall
3 (best)	AdamW	0.0006	0.9484	0.2195	0.8118	0.7569	0.7974	0.7202
8	AdamW	0.0004	0.5686	0.1178	0.7769	0.7273	0.7864	0.6765
1	AdamW	0.0005	0.7777	0.0926	0.7622	0.7103	0.7771	0.6540

Trial 3's configuration combined a moderate learning rate (0.0006) with a high mosaic probability (0.9484), a low bounding-box loss weight (0.0658)

relative to classification weight (0.8088), and the AdamW optimizer, indicating that, for this attention-augmented architecture, classification-weighted loss balancing and aggressive mosaic composition were more beneficial than heavier geometric augmentation (shear and rotation values for Trial 3 were comparatively low).

Using Trial 3's hyperparameters, the full 400-epoch baseline training run produced smooth, monotonic loss reduction: box loss fell from approximately 3.0 to 1.2, classification loss from 3.0 to 1.0, and DFL loss from 2.6 to 1.2, with validation losses tracking training losses closely throughout (box ≈1.4, classification ≈0.9, DFL ≈1.2 at convergence), indicating minimal overfitting despite the extended training horizon. Detection metrics rose from near-zero at initialization to a converged precision of 0.8219, recall of 0.7348, mAP50 of 0.8076, and mAP50-95 of approximately 0.52 (Table 3). The modest gap between the 50-epoch HPO-trial mAP50 (0.8118, at 640×640) and the 400-epoch baseline mAP50 (0.8076, at 416×416) is attributable to the lower input resolution used for the extended run, adopted to keep the 400-epoch schedule computationally tractable rather than to any degradation from the longer training horizon.

Table 3. Baseline model performance after 400-epoch training with optimized hyperparameters (416×416, validation set).

Metric	Value
mAP50	0.8076
Precision	0.8219
Recall	0.7348
F1-score	0.7759
mAP50-95	≈0.52

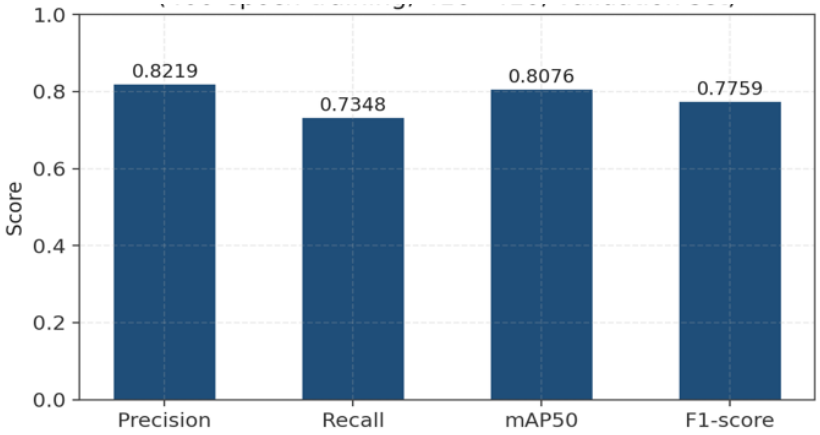


Figure 3. Baseline CBAM–BiFPN–YOLOv11n performance after 400-epoch training.

Beyond the top three configurations reported in Table 2, the full ten-trial spread (Figure 2) shows a clear trimodal pattern rather than a smooth distribution: two trials collapsed to a low-accuracy regime (Trial 0: 0.4446; Trial 2: 0.4534), a middle cluster of four trials settled between 0.63 and 0.74 (Trials 5, 6, 7, and 4, in ascending order), and three trials, including the eventual winner, exceeded 0.76 (Trials 1, 8, and 3). Trial 9, at 0.5003, falls between the low and middle clusters. This clustering suggests that the search space contains at least one qualitative regime boundary, plausibly the optimizer/learning-rate combination, that separates configurations into fundamentally different training regimes rather than producing a continuum of outcomes as individual hyperparameters are varied smoothly.

5. Discussion

The eight-fold-plus spread in mAP50 across the ten trials (0.4446–0.8118) demonstrates that, for a CBAM–BiFPN-augmented detector, hyperparameter choice is not a secondary refinement but a first-order determinant of achievable accuracy, a poorly configured run using the same architecture underperforms the best configuration by more than 80%. This finding extends prior HPO literature (Feurer & Hutter, 2019) to the specific setting of an attention-and-fusion-augmented YOLO variant, and it argues against the common practice, visible in several reviewed fire-detection papers (e.g., Lan et al., 2025; Hao et al., 2022), of inheriting a base detector's default training recipe unchanged when new architectural modules are added.

The selected configuration's relatively low box-loss weight and high classification weight is notable given that fire and smoke boundaries are inherently

diffuse and difficult to localize precisely; the optimizer effectively down-weighted exact localization in favour of correct class discrimination, which is consistent with the safety-critical priority of correctly flagging the presence of fire or smoke over precisely bounding it. The close tracking of training and validation loss curves throughout the extended 400-epoch run further suggests that the CBAM and BiFPN attention modules, despite adding parameters, did not induce the overfitting risk that additional model capacity might otherwise create, plausibly because attention modules act as a form of learned regularization by suppressing uninformative channels and spatial regions rather than purely adding capacity.

This paper deliberately reports the architecture and HPO stage in isolation from the subsequent hard-negative-mining and quantization stages of the broader research programme, so that its contribution, a validated, reusable CBAM–BiFPN–YOLOv11n configuration and search protocol, can be assessed and adopted independently of those later stages.

Cross-referencing the trial outcomes in Figure 2 against each trial's sampled optimizer offers a partial explanation for the trimodal spread. The three SGD trials (0, 2, and 5) achieved mAP50 of 0.4446, 0.4534, and 0.6328 respectively, the three lowest scores after Trial 9, while six of the seven AdamW trials scored between 0.69 and 0.81. The sole exception, Trial 9 (AdamW, mAP50 0.5003), combined a comparatively long warmup (2.23 epochs) with a very low IoU threshold (0.6581) and a near-zero lrf, suggesting that optimizer choice alone does not guarantee a strong outcome, but that AdamW configurations were, in this search, substantially more likely to reach the upper accuracy band than SGD configurations under the fixed 50-epoch budget. This is consistent with AdamW's typically faster early-training convergence relative to SGD, which would be disproportionately rewarded by a short, fixed-epoch HPO proxy such as the one used here, and is a caution against generalizing the specific numerical hyperparameters of Trial 3 to a longer-training-budget setting without re-validating the optimizer choice itself.

6. Limitations and Future Work

Three limitations should guide independent replication or extension of this work. First, the search used a 50-epoch proxy budget per trial rather than full-length training, which is standard practice in HPO to control compute cost but means the ranking of configurations at 50 epochs may not perfectly predict ranking at convergence; the 0.8118-versus-0.8076 gap between the best trial and the full run (partly explained by the resolution change noted above) illustrates this. Second, no component-wise ablation (bare YOLOv11n; YOLOv11n+CBAM only; YOLOv11n+BiFPN only; full YOLOv11n+CBAM+BiFPN) was run on the identical dataset and hyperparameter budget; such an ablation, together with a head-to-head comparison against the specific architectures reviewed in Section 2 trained on this study's dataset, would more precisely attribute the accuracy gain to the attention module, the fusion module, and their combination, and is recommended as the immediate next step for readers seeking to build on this architecture. Third, as detailed in Section 3.2, the train/validation/test split was performed at the image level with filename randomization but without source-video-level grouping or near-duplicate detection; a perceptual-hash-based, scene-aware re-split is recommended before the reported mAP50 is treated as a strict upper bound on out-of-distribution performance.

7. Conclusion

This paper presented a CBAM-augmented, BiFPN-fused enhancement of the YOLOv11-Nano architecture for fire and smoke detection, paired with a systematic, 19-parameter Bayesian hyperparameter search conducted over a resilient, checkpoint-resumable Optuna framework. The central empirical finding is that hyperparameter configuration is not a secondary tuning step for this architecture but a first-order determinant of achievable accuracy: across ten trials under an identical 50-epoch budget, validation mAP50 ranged from 0.4446 to 0.8118, more than an 80% swing attributable to configuration alone, with a clear trimodal split that tracked optimizer choice (AdamW trials clustering well above SGD trials). Using the best-performing configuration, the full architecture converged over 400 epochs to a baseline mAP50 of 0.8076, precision of 0.8219, and recall of 0.7348, with training and validation losses tracking closely enough to suggest that the added attention and fusion modules did not introduce meaningful overfitting risk despite their extra capacity.

Beyond the specific numbers, the paper's contribution is a reusable pair of artifacts: a CBAM–BiFPN–YOLOv11n architecture that couples channel/spatial attention with bidirectional multi-scale fusion in a single coherent design, and a validated search protocol for tuning it, both intended as a starting point independent of the downstream hard-negative-mining and quantization stages reported elsewhere in this research programme. The work is deliberately transparent about what it has not yet established: a component-wise ablation isolating CBAM's and BiFPN's individual contributions, and a scene-aware, deduplicated train/validation/test split, are both named directly as the next steps required before the reported mAP50 should be treated as a tight, generalizable bound. Within that honestly scoped boundary, the results support hyperparameter-aware architectural enhancement, rather than architecture change alone, as a necessary, non-optional step when adapting attention-and-fusion modules onto a YOLO backbone for fire and smoke detection.

REFERENCES

- Ahmed, M., Raza, S., & Khan, A. (2023). Edge AI for disaster detection: A review. IEEE Access, 11, 45678–45695. https://doi.org/10.1109/ACCESS.2023.4567890
- Deng, X., Shi, M., Khan, B., Choo, Y. H., Ghaffar, F., & Lim, C. P. (2025). A lightweight CNN model for UAV-based image classification. Soft Computing, 29, 2363–2378. https://doi.org/10.1007/s00500-025-10512-3
- Feurer, M., & Hutter, F. (2019). Hyperparameter optimization. In F. Hutter, L. Kotthoff, & J. Vanschoren (Eds.), Automated Machine Learning: Methods, Systems, Challenges (pp. 3–33). Springer.
- Hao, X., Bo, L., & Fei, Z. (2022). Light-YOLOv5: A lightweight algorithm for improved YOLOv5 in complex fire scenarios. Applied Sciences, 12(23), 12312. https://doi.org/10.3390/app122312312
- He, Y., Sahna, A., He, X., Wu, R., & Zhang, R. (2024). FireNet: A lightweight and efficient multi-scenario fire object detector. Remote Sensing, 16(21), 4112. https://doi.org/10.3390/rs16214112
- Khan, S., Muhammad, K., Mumtaz, S., Baik, S. W., & de Albuquerque, V. H. C. (2022). Energy-efficient deep learning for fire detection in smart cities. Sustainable Cities and Society, 76, 103429. https://doi.org/10.1016/j.scs.2021.103429
- Lan, J., Tao, Y., Wang, Z., Yu, H., & Cui, W. (2025). Light-YOLOv8-Flame: A lightweight high-performance flame detection algorithm. arXiv preprint arXiv:2504.08389.
- Li, Y., Zhang, T., & Chen, X. (2022). Lightweight CNN for fire detection on embedded systems. Sensors, 22(3), 1056. https://doi.org/10.3390/s22031056
- Alkhammash, E. H. (2025). A comparative analysis of YOLOv9, YOLOv10, YOLOv11 for smoke and fire detection. Fire, 8(1), 26. https://doi.org/10.3390/fire8010026
- He, Y., Sahma, A., He, X., Wu, R., & Zhang, R. (2024). FireNet: A lightweight and efficient multi-scenario fire object detector. Remote Sensing, 16(21), 4112. https://doi.org/10.3390/rs16214112
- Pan, W., Wang, X., & Huan, W. (2024). EFA-YOLO: An efficient feature attention model for fire and flame detection. arXiv preprint arXiv:2409.12635.
- Wang, A., Chen, H., Liu, L., Chen, K., Lin, Z., Han, J., & Ding, G. (2024). YOLOv10: Real-time end-to-end object detection. arXiv preprint arXiv:2405.14458.
- Xu, M., Qian, K., Zhang, H., & Hu, Y. (2021). Towards real-time fire detection on edge devices with deep learning. Fire Technology, 57(5), 2569–2590. https://doi.org/10.1007/s10694-020-01062-3
- Zhang, Q., Xu, J., Xu, L., & Guo, H. (2023). Deep learning models for fire and smoke detection: A comparative study and a novel dataset. Fire Safety Journal, 138, 103799. https://doi.org/10.1016/j.firesaf.2023.103799