Multimodal Explainable Wildfire Detection Using YOLOv8 and Transformer Attention

Submitted to Special Session 7: Intelligent Optimisation for Sustainable and Precision Agriculture

Lucky Banjare[2][0009−0005−6837−4062] and Dr. Vijaya J[1][0000−0002−9948−8573]

[1] *UG student/ Dept of CSE, IIIT Naya Raipur, Chhattisgarh, India*
lucky23100@iiitnr.edu.in
[2] *Assistant Professor/Dept of DSAI, IIIT Naya Raipur, Chhattisgarh, India*
vijaya@iiitnr.edu.in

Abstract. Purpose: In this research, we aim to create a quick and easy way to detect wildfires that is still equal to other current systems even with bad weather and low visibility. Current visual fire detectors have high false alarms and are not reliable weather conditions with fog, smoke, or haze and complex backgrounds. To improve these issues, our proposed method uses deep learning-based technology combined with real-time detection from various sensors.

Methodology: A multimodal fire detection approach is developed using RGB (visible) and thermal images, which is likely to be more robust when fused with each other.Improvements made over to YOLOv8 and through training the models using an attention-synchronized learning technique. Using model explainability methods such as Grad-CAM++ and EigenCAM, we can provide evidence for our decision-making and show that the trained models focused on fire-relevant regions of interest. Evaluation of the performance of our system is accomplished through four metrics: precision, recall, F1-score, and mAP.

Results: The result of the experiment shows that the multimodal approach provided superior results when compared to the single modality systems. The proposed model has an accuracy rate of greater than 97% for fire detection across different scenarios. Utilizing the attention mechanism also improves the capacity to interpret scenes; it helps reduce significantly the number of false detections.

Conclusion: This innovative framework, therefore, produces consistent and dependable results for real-world fire monitoring and gives increased precision to early warning systems in order to allow for the quick identification of a fire's start.

Keywords: Wildfire detection, multimodal deep learning, RGB–thermal fusion, YOLOv8, attention mechanism, explainable AI, real-time detection, fire and smoke detection, computer vision

1 Introduction

Wildfires have caused extreme damage to the environment and the population in the past, and they represent an enormous global challenge today [2,3]. Climate change has dramatically increased the frequency and intensity of wildfires in the last few decades due to a combination of increased temperature and prolonged dry conditions across the planet [28,29]. As a result, there is an urgent requirement for technology capable of providing reliable early detection of wildfires as well as continuous monitoring of large geographic areas [27].
Firefighters must respond quickly to new fires to minimize fire spread and protect their ecosystems, property, and life, as well as reduce suppression costs [1]. To monitor vast and generally inaccessible areas, fire departments have to use multiple methods of detecting wildfires. Satellite systems can fly over large areas and offer a wide range of fire-fighting support to county, state, and federal agencies. However, satellite systems usually detect wildfire events later than when they occur. This means that many of the existing solutions used on the ground to monitor rural communities have limitations. Ground and air sensor systems can be extremely expensive to install and maintain on rural properties, and they offer limited monitoring capabilities due to their limited coverage [5].
The recent advancements in deep learning utilize convolutional neural networks (CNNs) and YOLO (You Only Look Once) network architectures to detect fire and smoke from images [6,8,9]. Many of the current systems utilize only RGB images that do not perform well at night or in low visibility conditions [11]. Therefore, researchers have explored the use of multiple modalities, such as combining RGB and thermal imagery, to enhance detection in these types of conditions [7,12]. Thermal imaging provides accurate detection because it provides temperature information regardless of poor lighting or other obscured conditions [28].
The proposed system has the capability to detect wildfires in real time and reduce the number of false alarms. Furthermore, we implement Explainable AI methods for an improved understanding of the reasoning behind model prediction, which increases confidence and dependability of the user [18,19].

2 Literature Review

Due to climate change, changing land use and rising temperatures, wildfires represent an increasing threat to the environment and people. Wildfire detection technology is critical for reducing both damage and operational costs. However, traditional wildfire detection methods have many limitations, such as delayed detection, prohibitively high costs and limited spatial coverage [3,5].
Deep learning approaches using CNNs and YOLO-based models have demonstrated strong performance in real-time fire detection [6,8,10]. However, RGB-only systems struggle in low-visibility conditions, leading to missed detections and false alarms [11,12]. Recent research has incorporated thermal imaging to enhance detection reliability in adverse

environments. Effective fusion of RGB and thermal data remains challenging and requires attention-based mechanisms to emphasize relevant features [15,16].
Attention Models based on Transformers provide better scene insights by capturing long-distance relationships and ignoring unrelated regions. Using these techniques can help limit false positive detections [23,24,27]. Explainability AI, including Grad-CAM methods, increases trust between users and AI technology used for life-critical use cases by providing insight into the areas of images that led to its predictions [18,19,20,21].
This paper presents a straightforward and highly efficient method for detecting wildfires based on both RGB and thermal images using YOLOv8 , attention mechanisms, and explainable AI shown in Fig. 1. In addition, the system will work in real time and will be robust in a broad range of conditions and will allow users to easily interpret the information provided by the system, ultimately enhancing the safety of people monitoring for wildfires.

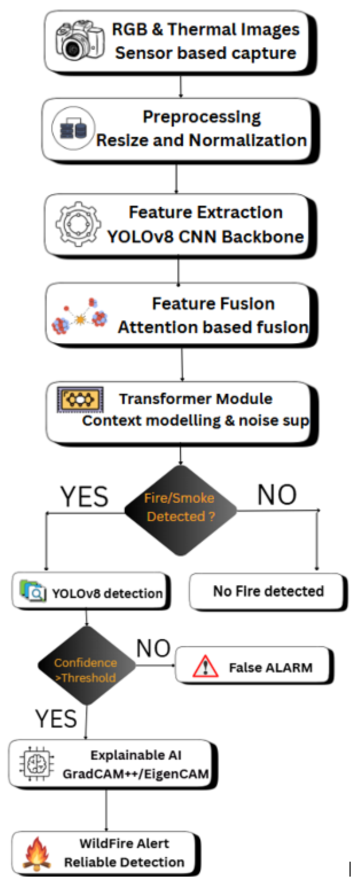

Fig. 1: Main architecture of the proposed multimodal explainable wildfire detection framework.

3 Methodology

This research will develop a new Multimodal Explainable Fire Detection Framework that incorporates a combination of RGB Images, Thermal Imaging, Attention Transformer Networks, and Explainable Artificial Intelligence Components to form a Hybrid Real-Time Fire Detection Pipeline, developed using YOLO v8 Architecture [10]. Its purpose is to provide Robust Performance and Interpretability for Safety Critical Wildfire Monitoring [7,11,12,28].

3.1 Dataset and Data Collection

A unique multimodal dataset was utilized for developing and validating the fire detection system. The dataset comprises both RGB and Thermal Imageries. All images in the dataset were captured under various conditions including daylight, darkness, smoke, fog as well as various size fires. [7].

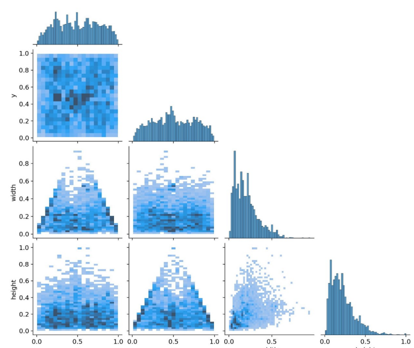

Fig. 2: Distribution of bounding box parameters (center x, center y, width, height) used in training data, showing spatial spread and object size variability.

The images were categorized with bounding boxes for fire, smoke, or other fire-like features in each image pair. These bounding boxes across the various images represented different sizes as illustrated in Fig. 2.
A split on the dataset was created for the system to learn instead of memorising from the same fire. The split consisted of 70% of the dataset being used for training, 20% to validate the model training and 10% to test the final version. Each set contain images of different fires.

Table 1: Dataset Composition

Category	**Samples**
Fire	14,317
Non-Fire	13,700
Smoke	25,434

This dataset helps the system learn to detect fire accurately in many different situations. DATASET LINK : Flame-2 Fire Detection Dataset

3.2 Data Preprocessing and Augmentation

RGB images were resized to 640×640 pixels to match YOLOv8's native input resolution, and thermal images were normalized using min–max scaling:

$$I_{norm}^{TIR} = \frac{I^{TIR} - \min(I^{TIR})}{\max(I^{TIR}) - \min(I^{TIR})}. \tag{1}$$

To improve model generalization and robustness, data augmentation strategies were applied:

– Geometric: horizontal flipping, rotation, random scaling.
– Photometric: brightness/contrast adjustment, Gaussian blur, noise addition.
– Mixed-sample: CutMix and MixUp.

Table 1 shows how many images are in each category, and Fig. 2 shows where the bounding boxes are in the images and their size differences.

3.3 YOLOv8-Based Detection Backbone

The YOLOv8 architecture [10] was selected for its anchor-free design, short processing delays, and high detection accuracy compared to previous YOLO versions. The modified backbone is shown in Fig. 3.

YOLOv8 consists of: CSP-based backbone for finding important details, PAN-FPN neck is used for combining information from small and large fire regions, Decoupled detection head for separate classification and bounding box regression.

The network is trained using a weighted sum of losses:

$$\mathcal{L}_{YOLO} = \mathcal{L}_{cls} + \lambda_1 \mathcal{L}_{box} + \lambda_2 \mathcal{L}_{obj}, \tag{2}$$

Here is what each part does:

– L_{cls}: Teaches the system which class the object belongs to (fire, smoke, or non-fire).
– L_{box} (CIoU): Teaches the system how to draw bounding boxes correctly around objects [10,6].
– L_{obj}: Teaches the system to determine whether an object is present or not.

This setup helps YOLOv8 detect fires quickly and accurately, even in challenging conditions such as smoke, fog, or nighttime environments.

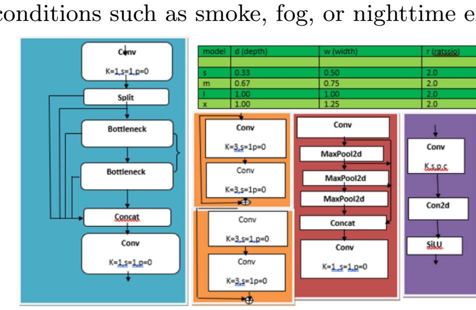

Fig. 3: YOLOv8 backbone adapted for wildfire detection.

3.4 Transformer-Based Attention Enhancement

In addition to the CNNs which are trained from images using small amounts of visual elements including colours, shapes and visual characteristics of fire flames and smoke, it is possible to develop models, capable of seeing the whole picture by using a different method. An attention-based Vision Transformer Module placed after the YOLOv8 feature combination. This capability allows the system to accurately identify both fire and smoke and differentiate them from other non-fire related areas in the scene even when there are extremely large fires occurring in very large areas.
Given a feature map $F \in \mathbb{R}^{H \times W \times C}$, query (Q), key (K), and value (V) vectors are computed, and self-attention is applied as:

$$Q = FW_Q, \quad K = FW_K, \quad V = FW_V, \tag{3}$$

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V. \tag{4}$$

This attention mechanism helps the model focus on real fire regions and ignore fire-like objects such as sunlight or clouds, thereby reducing false alarms and improving detection reliability [15,29]. The transformer provides greater distance-based knowledge of links between the areas of the image, which increases the understanding of the scene and decreases false positive alarms while increasing detection reliability.

3.5 RGB–Thermal Multimodal Feature Fusion

The FLAME2 RGB-Thermal Wildfire Dataset is used to enhance the ability of a system to detect wildfires in low light or smoke, as well as complicated weather conditions, through feature-level fusion. The fusion process is mathematically defined as:

$$F_{fusion} = \alpha F_{RGB} + (1 - \alpha) F_{TIR}, \tag{5}$$

Here, α is a number the system learns to decide how much to use from each type of image.If it is dark or smoky, the system uses more thermal information. A transformer attention module will analyze the merging of both images and determine where there is 'real' fire, avoiding fire-like objects including sunlight, cloud cover, artificial lighting, etc [15,29]. We trained the models with training and testing sets from the complete processing of the images and modifying the experimental design by using rotation, flipping, brightening, and darkening techniques. The entire integration of these features using the attention mechanism is illustrated in Fig. 4.

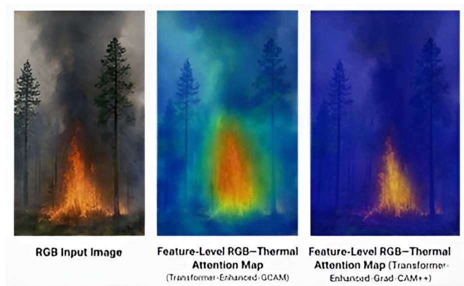

Fig. 4: Feature-level RGB–thermal fusion with transformer-based attention.

3.6 Explainable Artificial Intelligence (XAI) Module

The Wildfire Detection System incorporates Explainable AI (XAI) techniques in order to help explain its decisions. In particular, we employ two types of XAI techniques-EigenCAM and Grad-CAM++ to help illustrate where the model is focusing in an input image.
EigenCAM is used for explainability. Given activation tensor $A \in \mathbb{R}^{H \times W \times C}$:

$$CAM(x, y) = \sum_k v_1^k A_k(x, y), \tag{6}$$

where v_1 is the principal eigenvector. Grad-CAM++ is employed for qualitative analysis, helping visualize attention drift in challenging scenes. These XAI techniques ensure that the system's model reasoning process is transparent, enhancing reliability of the system in safety-critical wildfire monitoring applications. The architecture of Grad-CAM used for visualization is shown in Fig. 5.

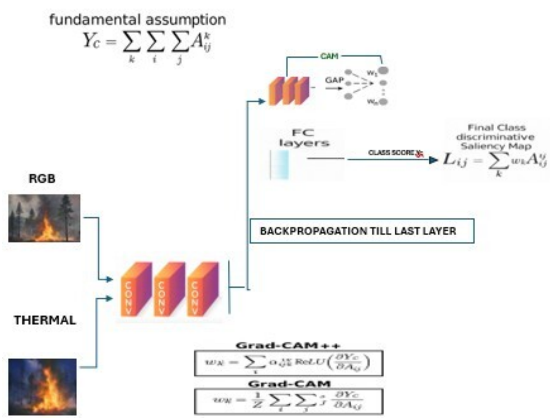

Fig. 5: Grad-CAM architecture

3.7 Evaluation Metrics

Performance evaluation of models includes using quality metrics that are widely accepted and standardised. The following quality metrics determine how well a model performs in classifying objects accurately, in

addition to the completeness of its detection, and the relationship between the two through the harmonic mean:
Mean Average Precision (mAP) measured at an Intersection Over Union (IoU) threshold of 0.5 and across multiple thresholds between 0.5 and 0.95) provides a holistic measure of both spatial and object.
The timing of a model, as measured by frames per second (FPS), helps to determine how quickly models can be integrated into real-world use cases.
Ablation studies have been conducted to determine how well models perform against various input modalities-RGB only, thermal only, and multimodal fused configurations.

4 Results

The section include both quantitative and qualitative analysis of the Multi-Modal Explainable Fire Detection Framework. The overall detection performance, the confidence-based metric metrics performance, comparisons of architecture design, the influence of the Transformer attention mechanism and thermal image inputs, Training Behaviour, and Explainability of Outcomes will be discussed in detail in this section.

4.1 Overall Detection Performance

To begin, we evaluated how well the fire detection system performs through the use of standard fire detection performance metrics such as Precision, Recall, F1 Score and mAP@0.5. The results show that the proposed framework is consistent and equitably balanced between detecting true fires and avoiding false positive Fire alerts.
With high precision, the fire detection alerting system can decrease the chances of fire detection alerts with strong recall.It demonstrates its ability to detect fires within a wide range of varying, challenging conditions. The F1 Score and mAP@0.5 further indicate that the combination of RGB and thermal images results in increased accuracy and enhanced localization for fire detection purposes [1,2,7].

Table 2: Confusion matrix values for RGB and Thermal wildfire detection

Model	**TP**	**TN**	**FP**	**FN**
RGB-based Detection	1420	1480	180	30
Thermal-based Detection	1457	1440	261	99

The prominent diagonal line within the two-dimensional matrix indicates that the thermal imaging system accurately identifies both fire events and non-fire events.It shows that fire events occur frequently, particularly under extreme conditions such as fog, smoke, and low visibility.
Figure 6 presents a qualitative comparison of detection results, demonstrating that thermal imagery consistently highlights fire regions in smoke-filled or low-light scenes.

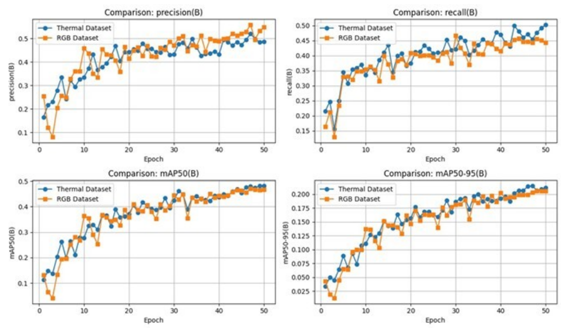

Fig. 6: Qualitative comparison of fire detection results using RGB and thermal imagery.

Table 3: Performance of the Proposed Multimodal Detection Model

Metric	Precision (%)	Recall (%)	F1-score (%)	mAP@0.5 (%)
Proposed Model	96	93	94	90

Table 3 provides quantitative data showing that our proposed multimodal detection approach creates a very acceptable precision and recall relationship. The consistently high value for precision indicates that the false positive rates are low, strong recall indicates a reliable ability for the multimodal framework to detect fire events in many different situations. High F1 scores and mean average precision reveal the multimodal framework's robustness and accuracy of detecting fire events.

4.2 Evaluation Metric Across Confidence Thresholds

The study evaluated the stability of their predictions and the reliability of the calibration of confidence using the following detection metrics over a variety of confidence thresholds.

Looking at Figure 7d, you can see that there was a high level of recall across the majority of the confidence range, which lowers the potential for missing any fire events in difficult conditions.

The precision-recall curves illustrate that the number of false positives can be reduced effectively by raising the confidence threshold and, therefore, show that fires and non-fires will be distinguished confidently.

The F1 score remained fairly constant throughout the various threshold levels, demonstrating that the proposed multimodal model provides a balanced performance and is therefore stable across different confidence threshold levels.

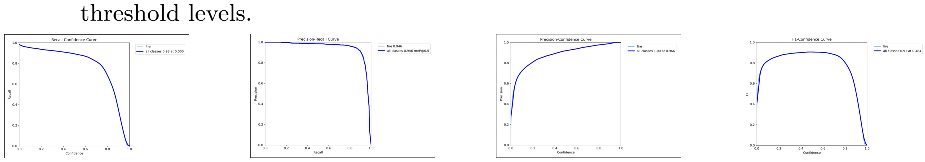

(a) Recall vs. Confidence (b) Precision vs. Recall (c) Precision vs. Confidence (d) F1 Score vs. Confidence

Fig. 7: Performance comparison of multimodal (RGB + Thermal) vs. RGB-only models across key evaluation metrics.

4.3 Comparison Between YOLOv7 and YOLOv8

The proposed framework was further evaluated against earlier YOLO architectures to analyze the effect of architectural advancements.

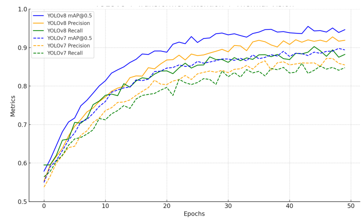

Fig. 8: Detection performance comparison between YOLOv7 and YOLOv8 variants.

Table 4: Performance Comparison Between YOLOv7 and YOLOv8

Model	Precision (%)	Recall (%)	F1-score (%)	mAP@0.5 (%)
YOLOv7	89	85	87	83
YOLOv8	94	90	92	88

Table 5: Baseline and Enhanced YOLOv8 Performance

Configuration	Precision (%)	Recall (%)	mAP@0.5 (%)
YOLOv8 (Visual Only)	91	85	84
YOLOv8 + Attention	94	90	88
YOLOv8 + Attention + Thermal	97	94	92

Results from Tables 4 and 5 demonstrate that YOLOv8 consistently outperforms YOLOv7 while maintaining a favorable speed–accuracy trade-off.

4.4 Impact of Transformer-Based Attention

The research to determine the benefits of improved contextual modeling added transformer-based attention into the suggested fire detection pipeline. The global dependency model across the spatial feature space created by the attention mechanism enables the network to focus on the semantically relevant fire areas while attenuating background noise and visually confusing patterns.

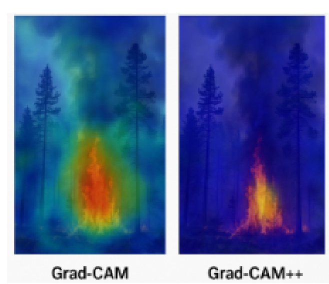

Fig. 9: Illustration of Grad-CAM and Grad-CAM++ attention mechanisms.

Table 6: Effect of Transformer-Based Attention

Model	Precision (%)	Recall (%)	F1-score (%)	FPS
Without Attention	95.3	94.1	94.7	62
With Attention	97.2	96.5	96.8	59

Table 6 Incorporating attention results in notable improvements across all accuracy metrics, with precision increasing from 95.3 percentage to 97.2 percentage and recall improving from 94.1 percentage to 96.5 percentage . The corresponding rise in F1-score reflects more reliable and balanced detection performance, indicating effective suppression of false positives and enhanced fire localization.Overall, these results validate the effectiveness of transformer-based attention in improving detection robustness with minimal computational overhead.

4.5 Impact of Thermal Vision Integration

The contribution of thermal sensing was analyzed through modality-wise evaluation.

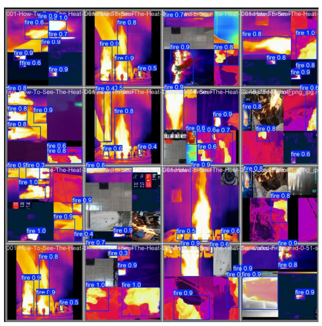

(a) Result of Thermal Images

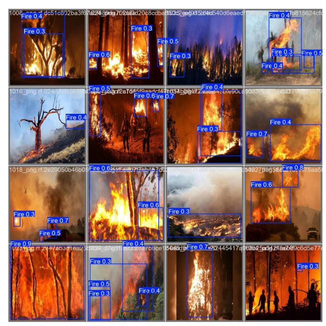

(b) Result of RGB Image

Fig. 10: Comparison of Fire Detection in Thermal and RGB Images

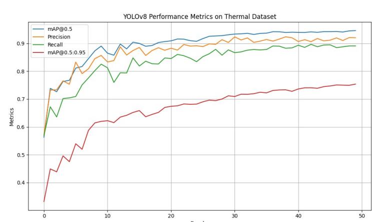

Fig. 11: Epoch-wise performance metrics on thermal dataset.

Table 7: Effect of Thermal Vision Integration

Input Modality	Precision (%)	Recall (%)	F1-score (%)	mAP@0.5 (%)
RGB Only	91	85	88	84
RGB + Thermal	94	90	92	88

Table 8: Combined Effect of Attention and Thermal Vision

Configuration	Precision (%)	Recall (%)	F1-score (%)
YOLOv8 + Transformer	94	90	92
YOLOv8 + Transformer + Thermal	94	90	92

The results indicate that thermal fusion significantly enhances robustness under low-light and smoke-heavy conditions without sacrificing precision.

4.6 Training Dynamics and Convergence Analysis

Training stability was examined to validate generalization behavior.

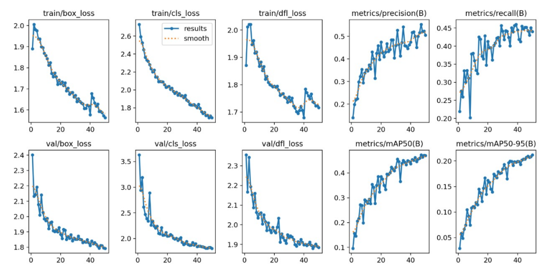

Fig. 12: Training and validation loss curves for the proposed model.

The smooth convergence trends observed in Figure 12 indicate stable optimization with no signs of overfitting.

4.7 Activation Function Comparison

Table 9: Activation Function Performance Comparison

Activation Function	Precision (%)	Recall (%)	F1-score (%)
ReLU	92	89	90
Leaky ReLU	93	90	91
Swish	94	90	92

Swish activation demonstrates superior convergence behavior and improved generalization compared to ReLU-based alternatives.

4.8 Explainability and Visual Interpretation Results

Explainability analysis was conducted using multiple visual interpretation methods.

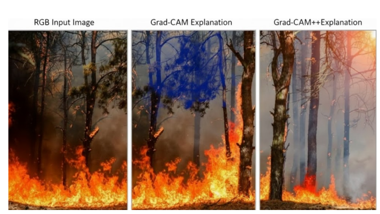

Fig. 13: Comparison of Grad-CAM and Grad-CAM++ explanations

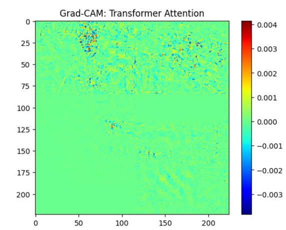

Fig. 14: Transformer attention heatmap highlighting fire regions.

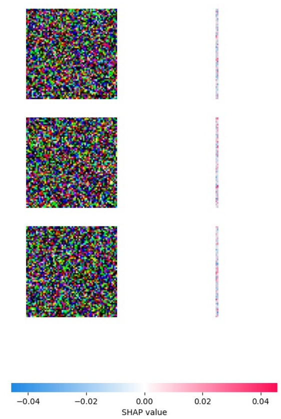

Fig. 15: SHAP-based explanation maps showing regional feature contributions.

The explainability results confirm that the proposed framework focuses on semantically meaningful fire regions, supporting trustworthiness and deployment readiness in safety-critical applications.

5 Conclusion and Future Work

The proposed system for smart wildfire detection combines the use of rgb and thermal cameras to detect wildfires quickly and in real-time using the YOLOv8 architecture. The incorporation of both RGB and thermal images allows the system to work effectively in difficult environments such as smoke, fog, and darkness, conditions which are challenging for traditional camera systems .To further enhance the smart capabilities of the system, transformer-based attention has been introduced. This will allow the model to understand the whole image and concentrate only on actual areas with flames and to disregard everything that resembles flames, such as sunlight or bright light sources. By improving the quality of detection, reducing the false alarm rate, and providing improved bounding boxes around flame areas, it is easier to interpret and verify the detection of wildfires.
To improve the fire detection system in the future, video data will be used in addition to static images because it gives a better sense of the evolution of the fire and minimizes false alarms that can be caused by light reflections or other shiny objects as they appear and disappear.

The fire detection system could also run on low-power and/or small devices such as edge computers and embedded systems; this would allow the fire detection system to be deployed effectively in remote locations, such as forests, that have limited resources. Finally, utilizing larger, more varied wildfires and more advanced explainable AI techniques will further strengthen the fire detection system, making it more trustworthy, capable and beneficial to wildfire monitoring in real-life situations.

References

1. M. R. Frizzi, R. Kaabi, M. Bouchouicha, and B. Fnaiech, "Convolutional neural network for wildfire detection," *IEEE International Conference on Image Processing*, 2016.
2. Y. Li, J. Zhang, and Y. Wang, "Forest fire detection based on deep learning," *Neural Computing and Applications*, vol. 31, no. 12, 2019.
3. T. Toulouse, L. Rossi, and A. Akhloufi, "Deep learning-based wildfire detection using UAV imagery," *Remote Sensing*, vol. 12, no. 1, 2020.
4. S. Muhammad, R. Ullah, and A. Khan, "Fire detection using convolutional neural networks," *IEEE Access*, vol. 7, 2019.
5. J. Chen, Y. Zhang, and J. Ouyang, "Vision-based fire detection using deep convolutional neural networks," *Fire Safety Journal*, vol. 110, 2019.
6. Z. Zhao, H. Zhang, and Y. Wang, "Fire detection in surveillance videos using YOLO," *Sensors*, vol. 20, no. 14, 2020.
7. I. Ibraheem, D. K. I. Ali, H. O. Nyongesa, and T. H. A. Soliman, "Forest Defender Fusion: An intelligent wildfire surveillance system based on drone-enabled fusion of RGB and infrared imagery," *Sensors*, vol. 22, no. 23, 2022.
8. Y. Li, X. Wu, and Z. Zhao, "Fire detection based on YOLOv5 network," *Electronics*, vol. 10, no. 6, 2021.
9. M. Bochkovskiy, A. Wang, and H. Liao, "YOLOv4: Optimal speed and accuracy of object detection," *arXiv preprint arXiv:2004.10934*, 2020.
10. G. Jocher, A. Chaurasia, and A. Qiu, "YOLOv8: Ultralytics next-generation object detection," *Ultralytics Documentation*, 2023.
11. A. Akhloufi and R. Larbi, "Multispectral fire detection using deep learning," *Remote Sensing*, vol. 13, no. 4, 2021.
12. M. Yuan, Y. Liu, and Z. Zhang, "Fire detection using infrared and visible image fusion," *Infrared Physics & Technology*, vol. 98, 2019.
13. A. Dosovitskiy et al., "An image is worth 16x16 words: Transformers for image recognition at scale," *International Conference on Learning Representations (ICLR)*, 2021.
14. Z. Liu, Y. Lin, and Y. Cao, "Swin Transformer: Hierarchical vision transformer using shifted windows," *IEEE International Conference on Computer Vision (ICCV)*, 2021.
15. X. Chen and Y. Zhang, "Attention-based feature fusion for multimodal fire detection," *Pattern Recognition Letters*, vol. 152, 2021.
16. T. Baltrusaitis, C. Ahuja, and L. Morency, "Multimodal machine learning: A survey and taxonomy," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 2, 2019.

17. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
18. R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," *IEEE International Conference on Computer Vision*, 2017.
19. A. B. Arrieta et al., "Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges," *Information Fusion*, vol. 58, 2020.
20. F. Samek, K. Müller, and W. Samek, "Explainable deep learning: A field guide for the uninitiated," *arXiv preprint arXiv:2004.14545*, 2020.
21. S. Ribeiro, M. Singh, and C. Guestrin, "Why should I trust you? Explaining the predictions of any classifier," *ACM SIGKDD Conference*, 2016.
22. J. Redmon and A. Farhadi, "YOLO9000: Better, faster, stronger," *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
23. S. Khan et al., "Transformers in vision: A survey," *ACM Computing Surveys*, vol. 54, no. 10, 2022.
24. A. Vaswani et al., "Attention is all you need," *Advances in Neural Information Processing Systems*, 2017.
25. M. Everingham et al., "The Pascal Visual Object Classes (VOC) challenge," *International Journal of Computer Vision*, vol. 88, 2010.
26. P. Viola and M. Jones, "Rapid object detection using a boosted cascade of simple features," *IEEE Conference on Computer Vision and Pattern Recognition*, 2001.
27. H. Zhang, J. Li, and Y. Liu, "Fire detection based on spatiotemporal attention," *IEEE Access*, vol. 9, 2021.
28. J. Tang, X. Luo, and Y. Sun, "Wildfire monitoring using deep learning and thermal imagery," *Remote Sensing Letters*, vol. 12, no. 8, 2021.
29. C. Zhou, L. Wang, and S. Yang, "Multimodal wildfire detection using CNN–Transformer hybrid models," *Sensors*, vol. 23, no. 4, 2023.
30. R. Sutton and A. Barto, "Reinforcement learning: An introduction," *MIT Press*, 2018.