Hybrid Reinforcement Learning-Based Hyper-Parameter Optimization with Yolov8 Indoor Fire Recognition

Patrick D. Cerna, Harvey C. Quijada, Michael Joseph E. Ortaliz, Kenneth Charles H. Saluna
College of Computer and Information Sciences, Mapua Malayan College Mindanao, Philippines

***Abstract*—This study presents an indoor vision-based fire detection system that integrates a YOLOv8n object detection model with a Reinforcement Learning-based Optimization Algorithm (ROA) for hyperparameter tuning. The research investigates three key aspects: 1) the effectiveness of ROA in improving model performance, 2) the optimal smart camera resolution and placement for indoor fire detection, and 3) the implementation of a real-time dual-channel user notification system. The BantaySunog model iteratively adjusted a few hyperparameters using the Reinforcement Learning-based Optimization Algorithm (ROA) [Talaat & Gamel, 2023]. An episodic framework was used for training, with 15 episodes of 20 epochs each, for a maximum of 300 epochs. Each episode's top weights were carried over to the following one. To balance exploration and exploitation, ROA employed an epsilon-greedy policy with an epsilon value that decreased from 0.9 to 0.2. Experimental results show that while ROA reduced training time and yielded a more conservative prediction strategy, it did not consistently outperform the baseline YOLOv8n in terms of detection metrics such as recall and mAP50-95. Camera deployment tests identified that positioning cameras away from direct light sources significantly improved detection success, with both elevation and resolution contributing to overall system performance. Finally, a dual-channel alert mechanism combining Firebase Cloud Messaging (FCM) and Telerivet SMS API enabled the timely delivery of fire alerts, aligning with real-world standards. The findings contribute to the development of reliable and accessible fire detection systems, especially for densely populated residential areas with limited infrastructure.**

***Keywords*—Fire detection; YOLOv8; reinforcement learning; hyperparameter optimization; convolutional neural network; indoor surveillance; real-time alert system**

I. Introduction

Fires cause significant harm globally due to their destructive potential in residential, commercial, and natural environments [1]. Early detection is essential; for instance, in Seoul, property damage and casualties increase significantly if fires are not reported within five minutes [2]. In the Philippines, over 3,000 incidents occurred in the first two months of 2024 alone, with a majority in residential zones and caused by lighted cigarettes, cooking, and faulty electrical wiring [3]. Davao City has seen a 300% increase in fire damage this year [4-6].

Traditional fire detection systems—based on heat, smoke, or flame sensors—remain dominant due to building code requirements [7], but they suffer delayed activation and high false-alarm rates indoors [8-10]. In contrast, vision-based fire detectors (VFDs) use real-time video analysis and CNNs such as YOLO for faster detection [11, 12], although most current work focuses on outdoor or large-scale fires [7, 13].

This study proposes BantaySunog, an indoor fire detection and alert system combining a YOLOv8n model with hyperparameter tuning via Reinforcement Learning-based Optimization Algorithm (ROA) [14]. It also integrates Firebase Cloud Messaging and the Telerivet SMS API for real-time, dual-channel notifications. The research is guided by three objectives: 1) evaluate the effect of ROA on fire detection performance, 2) identify optimal camera resolution and positioning, and 3) assess the system's ability to deliver real-time alerts under various conditions.

II. Related Works

This section reviews prior studies that informed the design and evaluation of BantaySunog. It focuses on six domains: 1) fire incident trends, 2) traditional vs. vision-based fire detection, 3) reinforcement learning-based optimization of CNNs, 4) image dataset augmentation, 5) camera deployment factors, and 6) multi-channel user notification systems.

Fire outbreaks remain a major threat worldwide due to human error, electrical faults, and inadequate infrastructure [1]. In South Korea, architectural fires account for over 80% of casualties and 88% of fire-related property damage, emphasizing the vulnerability of indoor environments [2]. In the Philippines, smoking and open flames were leading causes of fire in early 2024 [3], with Davao City alone reporting a 300% increase in fire damage [4-6]. Traditional detectors rely on sensors for heat, smoke, or UV radiation [15, 16], but they are susceptible to delayed activation and false positives, particularly in ventilated indoor environments [8-10, 17]. In contrast, VFDs use computer vision and CNNs to detect fire visually in real-time [11, 12], allowing earlier intervention. However, these systems demand high computation and visibility, and can still produce false detections if affected by lighting conditions or obstructions [13, 18, 19].

YOLO-based models have become increasingly favored due to their balance of speed and accuracy [20, 21]. Several studies highlight their adaptability to indoor and smart city scenarios [7, 11, 19]. ROA is a recent reinforcement learning-based method designed to optimize CNN hyperparameters through Q-learning [14]. It uses an agent that iteratively refines network-trained hyperparameters (NTHs) such as learning rate and loss weights to improve training outcomes. ROA operates by updating a Q-table that guides hyperparameter selection based on prior training rewards. While promising for image classification tasks, its effect on object detection remains underexplored. Data augmentation is essential in training robust models under limited data conditions [22, 23]. Techniques such as flipping, rotation, and color distortion increase dataset variety and reduce overfitting [24-26]. Larger datasets have been created through systematic augmentation [11, 27], often using specialized tools such as Albumentations [28].

Camera resolution, orientation, and elevation play a vital role in real-world deployment. Higher resolution enhances detail detection but increases computational load [29-32]. Camera orientation significantly affects detection accuracy; glare or reflections from windows can lead to false positives [33-37]. Studies suggest that facing cameras away from bright sources improves reliability [38, 39].

Elevation also influences the field of view and detection success. Ceiling-mounted cameras offer wider coverage [40, 41], while lower placements may capture finer details for specific tasks [42]. For timely alerts, fire detection systems must notify users through redundant channels. Firebase Cloud Messaging (FCM) supports rich, image-based mobile notifications [43], while Telerivet's REST API ensures SMS delivery even without internet [44]. Prior research supports combining both for higher reliability in emergencies [45-47]. Such dual-channel systems align with best practices in crisis response and are particularly useful in areas with unstable connectivity.

III. Materials and Methods

This study employed a quantitative research design structured around three objectives: 1) to evaluate the impact of ROA on training a YOLOv8n model, 2) to determine optimal smart camera configurations for fire detection, and 3) to implement a dual-channel alert system for real-time notifications. The conceptual framework in Fig. 1 illustrates the overall workflow of the study.

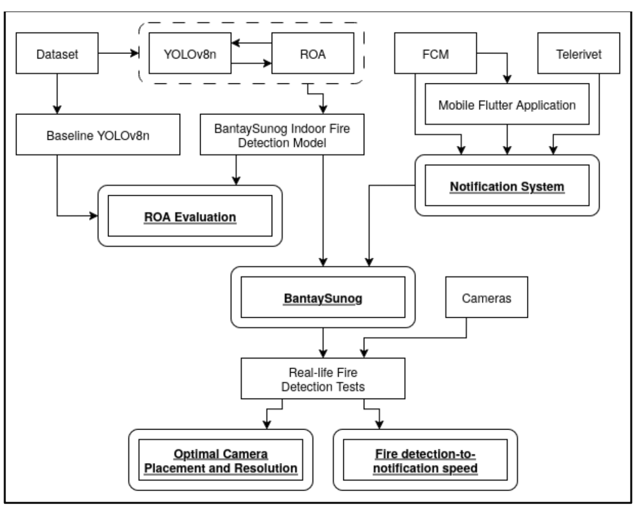

Fig. 1. Conceptual framework.

A. Dataset Augmentation

The "Indoor Fire Detection Dataset" in roboflow [48] was used, comprising 920 labeled images split into training, validation, and test sets. To improve generalization and simulate visual variability, the dataset was expanded to 3,680 images using the Albumentations library [28]. Augmentation steps included horizontal flipping (50% chance), random rotation (±15°, 30% chance), and fixed-angle rotations of 90°, 180°, and 270°. A comparison of the original and augmented versions is shown in Fig. 2.

Fig. 2. Sample indoor fire image before and after augmentation.

B. Model Training

Model training was based on YOLOv8n [20] using Ultralytics's YOLO module in Python, executed in Google Colab Pro with NVIDIA T4 GPUs.

The BantaySunog model employed the Reinforcement Learning-based Optimization Algorithm (ROA) [Talaat & Gamel, 2023] to iteratively tune selected hyperparameters. Training followed an episodic structure: 15 episodes of 20 epochs each, totaling 300 maximum epochs. The best weights from each episode were carried over to the next. ROA used an epsilon-greedy policy with the epsilon value decaying from 0.9 to 0.2 to balance exploration and exploitation. Fig. 3 visualizes the episodic training process of the BantaySunog model.

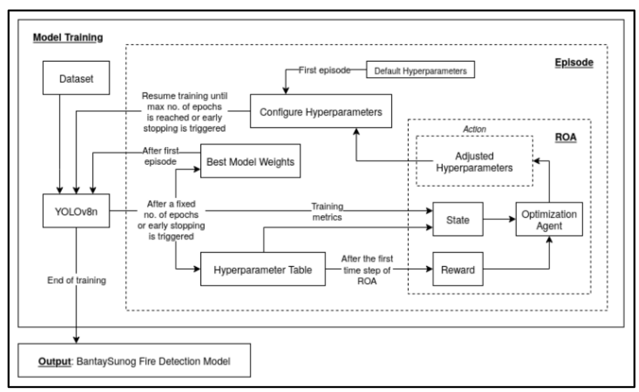

Fig. 3. The BantaySunog model training process.

The hyperparameter search space included four network-trained hyperparameters (NTHs) and three augmentation parameters, as shown in Table I, along with the candidate value ranges for each.

To provide a baseline for comparison with the BantaySunog model, YOLOv8n was trained continuously for up to 300 epochs using default hyperparameters. Early stopping was enabled with a patience of 75 epochs.

TABLE I. HYPERPARAMETER SEARCH SPACE AND VALUE RANGES

Hyperparameter	Default values	Range of values
Final learning rate	0.01	[0.001, 0.005, 0.01, 0.05, 0.1]
Box loss weight	7.5	[7, 7.25, 7.5, 7.75, 8]
Classification loss weight	0.5	[0.38, 0.44, 0.5, 0.56, 0.63]
Distribution focal loss weight	1.5	[1, 1.25, 1.5, 1.75, 2]
Degrees of rotation	0	[0, 45, 90, 135, 180]
Shear factor	0	[0, 3.75, 7.5, 11.25, 15]
Perspective transformation	0	[0, 0.00025, 0.0005, 0.00075, 0.001]

C. ROA Evaluation

This study employed a systematic evaluation comparing the BantaySunog model's training and testing results against those of the baseline YOLOv8n model trained with default values to determine the effectiveness of ROA as a hyperparameter optimization strategy. The following key metrics were considered in the training and testing results comparisons: 1) precision, which measures the proportion of correctly identified fire instances among all detections made, 2) recall, which captures the model's ability to detect actual fire instances present in the scene, 3) mAP50, which represents the average precision across all classes using a single Intersection over Union (IoU) threshold of 0.50, and 4) mAP50-95, a more comprehensive metric that evaluates the model's fine-grained localization accuracy and consistency.

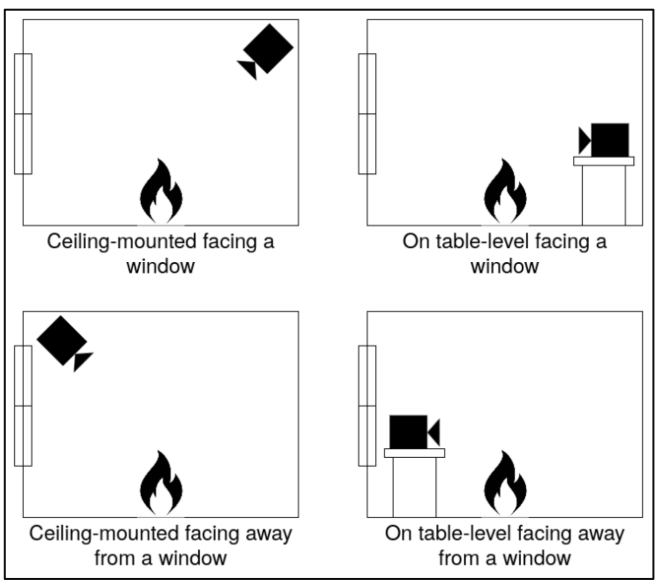

Fig. 4. Camera placements for testing.

D. Camera Resolution and Placement

Three TP-Link Tapo cameras were tested: C200 (1080p), C211 (2K 3MP), and C220 (2K QHD). Each camera was deployed in four positions: 1) ceiling-mounted facing a window, 2) table-level facing a window, 3) ceiling-mounted facing away from a window, and 4) table-level facing away. Each position is illustrated in Fig. 4. This experiment evaluated how resolution, glare exposure, and elevation influence fire detection accuracy [29, 38, 40].

E. Notification System

To ensure prompt alerts under various connectivity conditions, a dual-channel system was developed. Firebase Cloud Messaging (FCM) [43] was used for push notifications via a Flutter mobile app. For offline scenarios, Telerivet's REST API [44] was integrated into the model to trigger SMS alerts through a dedicated Android gateway.

F. Expert Consultation on Real-world Standards

The detection-to-notification delay was benchmarked based on guidance from Fire Officer 1 Jesardnel S. Vargas, FDAS Specialist at BFP Davao. He specified a maximum 30-second detection window from the fire's incipient stage. During camera tests, a stopwatch mechanism was used to measure detection and alert times relative to this benchmark.

IV. RESULTS AND DISCUSSION

On the original training set, the baseline model completed training in 117 epochs (0.378 hours), whereas the ROA-optimized BantaySunog model converged slightly faster in 106 epochs (0.362 hours). This trend persisted on the augmented training set, where the baseline required 176 epochs (2.265 hours) compared to BantaySunog's 149 epochs (2.045 hours). These results indicate that ROA reduced both the number of epochs and the amount of time needed for training.

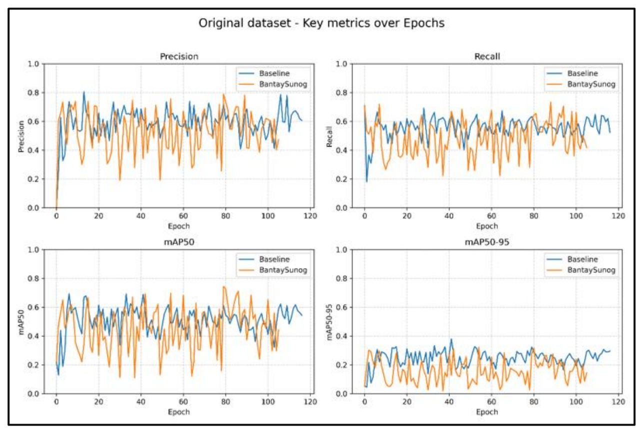

Fig. 5. Training metrics on the original dataset.

However, on the original training set, the baseline model outperformed BantaySunog in training stability and performance, as seen in Fig. 5. Precision and mAP50 curves were higher and more stable for the baseline, while BantaySunog showed greater volatility—especially in mAP50–95, signaling weaker localization quality. Recall was similarly less consistent, with the baseline maintaining centered, stable performance across epochs while BantaySunog showed frequent dips and underperformance.

The performance gap widened on the augmented training set. The baseline model consistently led across all metrics—precision, recall, mAP50, and mAP50–95—while BantaySunog's curves remained flatter and lower in amplitude, as seen in Fig. 6. This suggested minimal learning progress per epoch and an inability to scale with the increased dataset complexity. Although BantaySunog occasionally peaked in precision during early epochs, this was not sustained and came at the cost of reduced recall and poor localization.

While ROA improved convergence speed, it failed to enhance—and in some cases undermined—training quality. The training metrics indicate that ROA guided the model toward more conservative but less effective learning behavior, especially on complex datasets. These results raise concerns about ROA's generalizability and its adaptability to broader training conditions beyond fast convergence.

On the original testing set, the baseline model outperformed the BantaySunog model in nearly all key metrics. Most notably, it achieved higher recall (0.9227 vs. 0.7879) and mAP50–95 (0.5548 vs. 0.3646), indicating better detection coverage and localization accuracy. Although BantaySunog reported higher precision (0.9217 vs. 0.8354), this came at the expense of missing more true positives, consistent with a more conservative detection strategy. This behavior likely stems from ROA's tuning bias toward reducing false positives. Fig. 7 compares the testing result metrics from both models on each dataset, while Fig. 8 and Fig. 9 visually illustrate this conservativeness, showing sparser bounding box predictions compared to the baseline's more liberal detection.

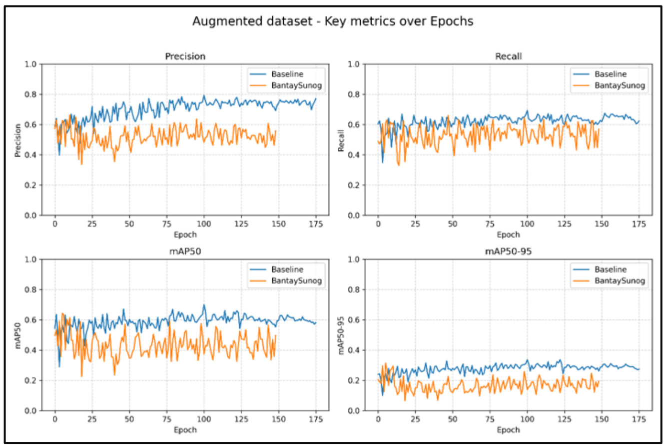

Fig. 6. Training metrics on the augmented dataset

On the augmented testing set, performance between the two models became more comparable, though the baseline still led in key areas. It slightly surpassed BantaySunog in precision (0.8896 vs. 0.8841) and mAP50–95 (0.5564 vs. 0.5153), while BantaySunog exhibited minor gains in recall (0.8664 vs. 0.8546) and mAP50 (0.9179 vs. 0.8898). These results suggest that ROA may provide some benefit in detecting more fire instances and improving bounding box coverage under data-rich conditions, but not enough to yield consistently superior overall performance.

ROA appears to promote cautious predictions but falls short in enhancing generalization. The occasional metric gains under augmentation are outweighed by consistent underperformance in stricter quality measures such as mAP50–95. This reinforces concerns about ROA's limited scalability and its tendency to prioritize false-positive minimization over comprehensive detection quality.

The ROA-optimized BantaySunog model consistently trained faster than the baseline YOLOv8n, converging in fewer epochs and with reduced wall time. However, this efficiency came at the cost of performance. Training curves were flatter, especially on the augmented dataset, indicating limited learning progress. Testing results echoed this trend: BantaySunog had lower recall and mAP50–95 on the original dataset and only marginal improvements under augmentation. Its higher precision reflected a conservative prediction style, but this also meant missing more true positives and weaker localization. Overall, while ROA reduced training duration, it did not enhance—and sometimes degraded—detection quality, suggesting that further refinement is needed before ROA can reliably optimize CNN-based object detectors for safety-critical applications.

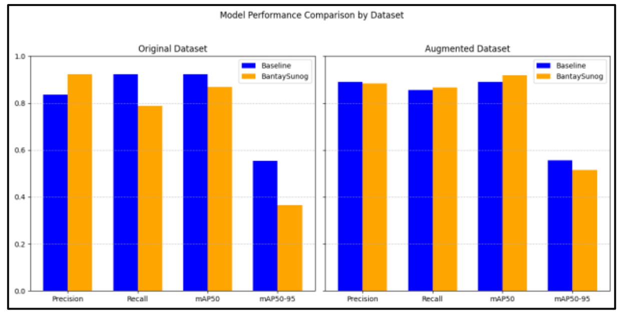

Fig. 7. Comparison of testing results.

Fig. 8. Original testing set predictions from each model.

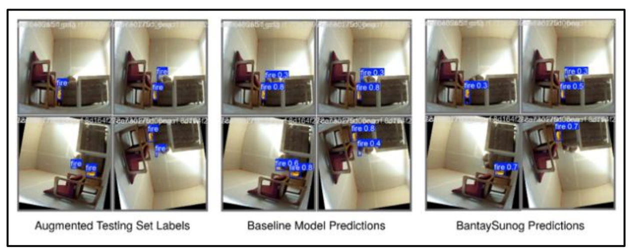

Fig. 9. Augmented testing set predictions from each model.

A. *Optimal Camera Resolution and Placement*

To determine optimal camera deployment for indoor fire detection, three smart cameras—TP-Link Tapo C200 (1080p), C211 (2K 3MP), and C220 (2K QHD)—were tested across four configurations: ceiling-mounted and table-level, each facing toward or away from a window. These combinations were chosen to evaluate the effects of resolution, elevation, and light source orientation on detection reliability.

Results revealed that camera orientation relative to light sources had the greatest influence on detection success. As shown in Table II, all three cameras consistently achieved better results when positioned away from windows, regardless of resolution or mounting height. Cameras facing windows frequently failed to detect fire instances due to glare, reflections, or contrast loss—particularly in brightly lit environments.

TABLE II. FIRE DETECTION RESULTS FROM EACH CAMERA AND POSITION

Cameras	Facing window		Facing away from window	
	On ceiling	On table-level	On ceiling	On table-level
C200	Fail	Fail	Success	Success
C211	Fail	Success	Success	Success
C220	Fail	Success	Success	Success

Fig. 10. Fire detection frames from ceiling-mounted cameras.

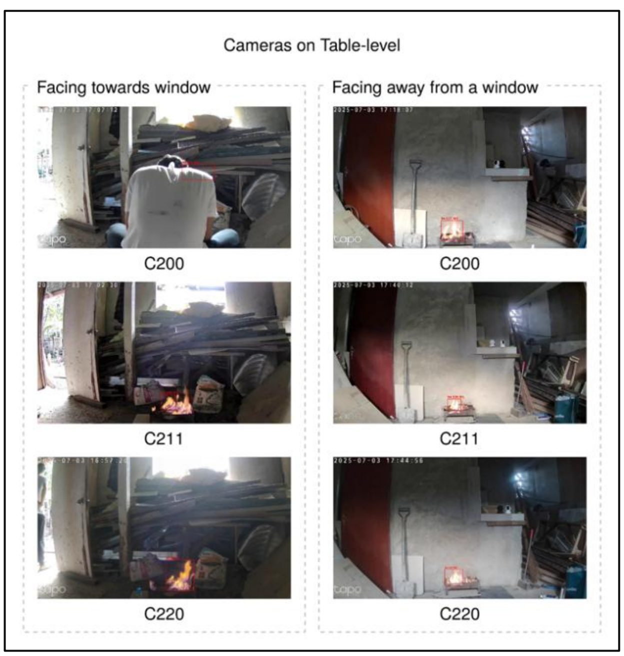

Fig. 11. Fire detection frames from table-level cameras.

While higher resolution (2K) offered marginal improvements in detecting small or distant flames, the gains were not substantial enough to justify the added computational overhead if placement was suboptimal. In particular, the C211 (2K 3MP) showed slightly better overall performance across conditions, but even this model struggled in unfavorable orientations. Visual results in Fig. 10 and Fig. 11 support these findings, with clearer and more accurate detections appearing in camera views facing away from light.

These results suggest that proper orientation and strategic positioning—avoiding direct light interference—are more critical to reliable fire detection than increasing resolution or changing elevation alone. Such insights are vital for effective system deployment in real-world environments, especially where infrastructure constraints limit camera options.

B. Implementation of the Notification System

To ensure timely user alerts in both connected and offline environments, BantaySunog integrates a dual-channel notification system using Firebase Cloud Messaging (FCM) and the Telerivet SMS API. Upon detecting a fire, the system automatically sends both an image-rich push notification and a backup SMS containing fire detection details.

Fig. 12 shows the FCM and SMS notifications on the end-user smartphone. This confirms the successful reception of alerts from BantaySunog through FCM and Telerivet in the end-user smartphone.

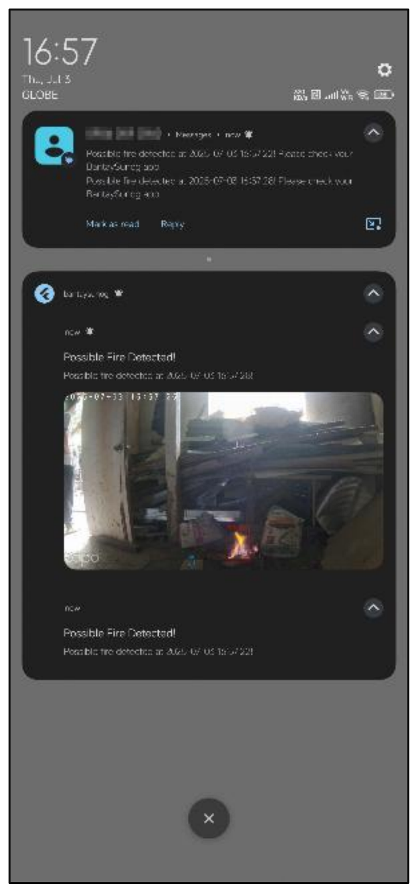

Fig. 12. FCM and SMS notifications.

TABLE III. FIRE DETECTION-TO-NOTIFICATION TIMES FROM SUCCESSFUL FIRE DETECTIONS

Successful camera configurations	Time (seconds)
C220, Table-level, Facing window	18.749
C211, Table-level, Facing window	8.495
C200, Table-level, Facing away from window	9.658
C211, Table-level, Facing away from window	9.446
C220, Table-level, Facing away from window	7.264
C220, Ceiling-mounted, Facing away from window	15.192
C220, Ceiling-mounted, Facing away from window	7.396
C220, Ceiling-mounted, Facing away from window	9.324

Table III presents the fire detection-to-notifications times from the successful fire detections during the tests to determine the optimal camera configuration. The average delay between fire detection and full alert receipt was 10.69 seconds. This includes model processing, message generation, and cross-platform delivery. All alerts were received in under 30 seconds, meeting the benchmark advised by FO1 Vargas of BFP Davao.

These results affirm that BantaySunog's notification system meets the real-time responsiveness required for fire emergencies and offers high reliability through multi-channel delivery.

V. CONCLUSIONS AND RECOMMENDATIONS

This study introduced BantaySunog, an indoor vision-based fire detection system that integrates a YOLOv8n object detection model with a reinforcement learning-based optimization algorithm (ROA), paired with a dual-channel alert system. The research evaluated three primary objectives: improving training performance through ROA, identifying optimal camera configurations for deployment, and ensuring real-time alert delivery under realistic conditions. Findings show that while ROA reduced training time and promoted a conservative prediction strategy, it failed to consistently improve detection performance. The BantaySunog model often underperformed in recall and mAP50–95, especially on the original dataset, suggesting that ROA's policy may require further tuning to handle more complex or diverse training scenarios. This indicates that, in its current form, ROA is not a reliable drop-in optimization tool for CNN-based fire detection in safety-critical applications.

Camera deployment tests revealed that orientation—specifically avoiding direct exposure to windows or light sources—had a greater impact on detection reliability than resolution or elevation. Proper placement significantly reduced false negatives, underscoring the importance of environment-specific configuration in real-world deployments. The dual-channel alert system using FCM and Telerivet successfully delivered notifications within an average of 10.69 seconds, well below the 30-second threshold advised by the Bureau of Fire Protection. This confirms the system's readiness for deployment in infrastructure-limited environments where redundancy and speed are critical. Future work should explore adaptive versions of ROA or alternative reinforcement learning strategies tailored for object detection. Additional evaluation in real fire conditions, including smoke or occlusion scenarios, would strengthen generalizability. Deployment across multiple households with varied layouts could further validate the system's robustness in uncontrolled environments.

REFERENCES

[1] AG Global Strategies, "Fire, accident & emergency." [Online]. Available: https://www.ag-globalstrategies.com/fire-accident-and-emergency

[2] J. Won, "Seoul fire response improvement plan," Policy Report, 2017.

[3] I. Laqui, "BFP: Smoking, open flames top causes of fire incidents in Q1," The Philippine Star, Mar. 1, 2024. [Online]. Available: https://www.philstar.com/headlines/2024/03/01/2337302

[4] E. Hinutan, "Fire damages in Davao City increased by 300%," NewsLine Philippines, Jan. 11, 2024. [Online]. Available: https://newsline.ph/top-stories/2024/01/11/fire-damages-in-davao-city-increased-by-300/

[5] I. C. Espinosa, "Davao City fire razes 350 houses," MindaNews, Dec. 27, 2023. [Online]. Available: https://mindanews.com/top-stories/2023/12/davao-city-fire-razes-350-houses/

[6] A. Colina IV, "2 fire incidents destroy 18 houses and a commercial building," MindaNews, Jan. 1, 2024. [Online]. Available: https://mindanews.com/top-stories/2024/01/2-fire-incidents-destroy-18-houses-and-a-commercial-building/

[7] Y. Ahn, H. Choi, and B. S. Kim, "Development of early fire detection model for buildings using computer vision-based CCTV," J. Build. Eng., vol. 65, p. 105647, 2023.

[8] S. Chen, J. Ren, Y. Yan, M. Sun, F. Hu, and H. Zhao, "Multi-sourced sensing and support vector machine classification for effective detection of fire hazard in early stage," Comput. Electr. Eng., vol. 101, p. 108046, 2022.

[9] T. G. Cleary, A. A. Chernovsky, W. L. Grosshandler, and M. D. Anderson, "Particulate entry lag in spot-type smoke detectors," Fire Safety Science, vol. 6, pp. 779–790, 2000.

[10] G. Holmstedt, S. E. Magnusson, and P. Thomas, "Detector environment and detector response: A survey," Tech. Rep. 3039, Lund Univ., Sweden, 1987.

[11] F. M. Talaat and S. A. Gamel, "RL based hyper-parameters optimization algorithm (ROA) for convolutional neural network," J. Ambient Intell. Humaniz. Comput., vol. 14, no. 10, pp. 13349–13359, 2023.

[12] M. Ahrens, "Smoke alarms in US home fires," Quincy, MA: National Fire Protection Association, Fire Analysis and Research Division, 2009.

[13] C. T. Truong, T. H. Nguyen, V. Q. Vu, V. H. Do, and D. T. Nguyen, "Enhancing fire detection technology: A UV-based system utilizing fourier spectrum analysis for reliable and accurate fire detection," Appl. Sci., vol. 13, no. 13, p. 7845, 2023.

[14] J. F. Middleton, "Developments in flame detectors," Fire Safety J., vol. 6, pp. 175–182, 1983.

[15] L. Craig and R. Awati, "Convolutional neural network (CNN)," TechTarget. [Online]. Available: https://www.techtarget.com/searchenterpriseai/definition/convolutional-neural-network

[16] K. Avazov, M. Mukhiddinov, F. Makhmudov, and Y. I. Cho, "Fire detection method in smart city environments using a deep-learning-based approach," Electronics, vol. 11, no. 1, p. 73, 2021.

[17] J. Pincott, P. W. Tien, S. Wei, and J. K. Calautit, "Development and evaluation of a vision-based transfer learning approach for indoor fire and smoke detection," Build. Serv. Eng. Res. Technol., vol. 43, no. 3, pp. 319–332, 2022.

[18] Q. Zhang, B. Li, and Z. Mao, "Real-time fire detection using YOLOv3," IEEE Access, vol. 7, pp. 115540–115550, 2019.

[19] Y. Wang, Q. Zhang, and B. Zhao, "Enhancing fire detection in smart buildings: An integration of YOLOv8 and reinforcement learning," Sustainability, vol. 15, no. 3, p. 1794, 2023.

[20] G. Jocher, A. Chaurasia, and J. Qiu, "Ultralytics YOLO (Version 8.0.0)" [Computer software]. Available: https://github.com/ultralytics/ultralytics

[21] F. M. Talaat and H. ZainEldin, "An improved fire detection approach based on YOLO-v8 for smart cities," Neural Comput. Appl., vol. 35, no. 28, pp. 20939–20954, 2023.

[22] W. J. Jang, "Enhancing fire image classification through transfer learning and data augmentation in convolutional neural networks," Asia-Pac. J. Converg. Res. Interchange, vol. 10, no. 1, pp. 13–27, 2024.

[23] J. Talukdar, A. Biswas, and S. Gupta, "Data augmentation on synthetic images for transfer learning using deep CNNs," in Proc. 5th Int. Conf. Signal Process. Integr. Netw. (SPIN), pp. 215–219, Feb. 2018.

[24] A. Salunke, "Boosting object detection accuracy: A comparative study of image augmentation techniques," 2023. [Online]. Available: https://doi.org/10.13140/RG.2.2.21787.23844

[25] B. Yilmaz and U. Kutbay, "YOLOv8-based drone detection: Performance analysis and optimization," Computers, vol. 13, no. 9, p. 234, 2024.

[26] B. Zoph, E. D. Cubuk, G. Ghiasi, T. Y. Lin, J. Shlens, and Q. V. Le, "Learning data augmentation strategies for object detection," in Proc. Eur. Conf. Comput. Vis., pp. 566–583, Aug. 2020.

[27] N. Yunusov, B. M. S. Islam, A. Abdusalomov, and W. Kim, "Robust forest fire detection method for surveillance systems based on you only look once version 8 and transfer learning approaches," Processes, vol. 12, no. 5, p. 1039, 2024.

[28] A. Buslaev et al., "Albumentations: Fast and flexible image augmentations," Information, vol. 11, no. 2, p. 125, 2020.

[29] Y. Hao et al., "Understanding the impact of image quality and distance of objects to object detection performance," in Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS), pp. 11436–11442, Oct. 2023.

[30] V. Růžička and F. Franchetti, "Fast and accurate object detection in high resolution 4K and 8K video using GPUs," in Proc. IEEE HPEC, pp. 1–7, Sept. 2018.

[31] E. Elshebli and F. Erdõs, "Optimizing video resolution for machine learning-based traffic monitoring systems: A performance analysis," in Proc. Int. Conf. Eng. Emerg. Technol. (ICEET), pp. 1–5, Dec. 2024.

[32] C. N. Clark et al., "Investigating the resolution-performance trade-off of object detection models in support of the Sustainable Development Goals," IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens., vol. 16, pp. 5695–5713, 2023.

[33] M. Z. Alam, Z. Kaleem, and S. Kelouwani, "How to deal with glare for improved perception of Autonomous Vehicles," arXiv preprint arXiv:2404.10992, 2024.

[34] S. Carvalho, J. Humphries, N. Dunne, and S. Leahy, "Impact of light flickering on object detection accuracy using convolutional neural networks," in Proc. Telecoms Conf. (ConfTELE), pp. 1–6, Feb. 2021.

[35] T. H. Chen, P. H. Wu, and Y. C. Chiou, "An early fire-detection method based on image processing," in Proc. Int. Conf. Image Process. (ICIP), vol. 3, pp. 1707–1710, Oct. 2004.

[36] J. A. Milke and T. J. McAvoy, "Analysis of fire and non-fire signatures for discriminating fire detection," Fire Safety Science, vol. 5, pp. 819–828, 1997.

[37] P. Sawant, A. Jadhav, N. Waghmode, P. Bansode, and K. Rajgude, "Pyro Alert: Fire Detection using Computer Vision and Alert System," in Proc. 5th Int. Conf. Commun. Inf. Process. (ICCIP), 2023.

[38] M. Sadok and R. R. Zakrzewski, "U.S. Patent No. 7,245,315," U.S. Patent and Trademark Office, Washington, DC, 2007.

[39] J. Pincott, P. W. Tien, S. Wei, and J. K. Calautit, "Indoor fire detection utilizing computer vision-based strategies," J. Build. Eng., vol. 61, p. 105154, 2022.

[40] C. H. Hsieh, Q. F. Zhang, and Z. Y. Chen, "An indoor surveillance system using fisheye camera and YOLOv4 object detection," in Proc. IEEE Int. Conf. Consum. Electron.–Taiwan, pp. 469–470, July 2022.

[41] X. Li, F. Lu, G. Tian, B. Huang, and S. Zhao, "Ceiling camera based environment objects recognition of intelligent space," in Proc. IEEE Int. Conf. Autom. Logist., pp. 1512–1516, Sept. 2008.

[42] R. Bodor, A. Drenner, P. Schrater, and N. Papanikolopoulos, "Optimal camera placement for automated surveillance tasks," J. Intell. Robot. Syst., vol. 50, no. 3, pp. 257–295, 2007.

[43] Firebase, "Firebase Cloud Messaging," 2025. [Online]. Available: https://firebase.google.com/docs/cloud-messaging

[44] Telerivet, "Telerivet Python Client Library," 2024. [Online]. Available: https://telerivet.com/docs/api/client/python

[45] C. Reuter and M.-A. Kaufhold, "Fifteen years of social media in emergencies: A retrospective review and future directions for crisis informatics," J. Conting. Crisis Manag., vol. 26, no. 1, pp. 41–57, 2018.

[46] J. Fajardo and C. Oppus, "A mobile disaster management system using the Android technology," WSEAS Trans. Commun., vol. 9, no. 6, pp. 343–353, 2010.

[47] S. Latif et al., "Leveraging volunteer and crowdsourced data for disaster resilience by using artificial intelligence and machine learning: A survey," Future Internet, vol. 10, no. 7, p. 1, 2019.

[48] Sih, "Indoor Fire Detection Dataset," Roboflow Universe, Sept. 2023. [Online]. Available: https://universe.roboflow.com/sih-hisav/indoor-fire-detection.