

AI Powered Drone: An Integrated Hexacopter System for Real-Time Fire Detection and Extinguishing

SAHIL J. DIGHE¹, SUDARSHAN A. JUNDRE², PROF. OMKAR WADNE³

^{1, 2, 3}*Department of Computer Engineering, Navsahyadri Group of Institutes, Pune University, Pune, India*

Abstract- *This research introduces the development of a semi-autonomous hexacopter drone designed for efficient fire detection and suppression, aimed at improving rapid response capabilities in hazardous environments. The system incorporates a Pixhawk flight controller powered by the ArduPilot firmware to achieve stable and precise navigation, alongside a YOLOv8 deep learning framework that enables real-time, vision-based fire recognition. The core advancement lies in the hybrid control architecture, which merges autonomous sensing with a manually triggered extinguishing mechanism, ensuring accurate operation while retaining essential human supervision for safety assurance. Experimental validation confirms the system's practical functionality, showing consistent flight stability, precise waypoint tracking, and effective fire detection during live testing. The study concludes that this integrated methodology effectively connects AI-enabled detection with physical action, presenting a scalable and efficient approach to enhance industrial safety while minimizing human exposure to danger. Future improvements are directed toward incorporating edge computing for faster onboard processing and conducting extensive real-world trials for performance optimization.*

Index Terms—*AI Drone, Fire Detection, YOLOv8, Pixhawk, Autonomous Systems, Fire Suppression, Hexacopter.*

I. INTRODUCTION

Fire incidents remain a major and ongoing global concern, posing serious risks to human life, inflicting extensive structural damage, and severely impacting natural ecosystems. One of the most crucial difficulties in firefighting lies in the early response stage, where timely action is essential, particularly in regions that are difficult to reach or unsafe for human intervention. This study aims to overcome this

limitation through the development of an advanced unmanned aerial system designed to support and enhance conventional firefighting operations.

A. The Critical Need for Advanced Firefighting Solutions

Conventional firefighting, while brave and essential, inherently exposes first responders to significant danger, particularly in complex environments like industrial facilities, high-rise buildings, and wildland-urban interfaces. The limitations of human-dependent protocols, including delayed detection, difficult localization, and the physical risks of intervention, often exacerbate the scale of damage before containment is achieved. There is a pressing need for technological solutions that can operate in these perilous conditions without risking human lives.

B. The Emergence of Intelligent UAV Systems

Recent progress in Unmanned Aerial Vehicle (UAV) technology has created significant opportunities for enhancing emergency response operations. These aerial platforms possess the unique capability to swiftly reach dangerous or hard-to-access zones, perform real-time aerial monitoring, and transport critical payloads. By integrating advanced sensing technologies and artificial intelligence (AI) algorithms, UAVs can evolve from being simple remote-controlled observation tools into intelligent autonomous responders capable of perceiving their surroundings, making informed decisions, and executing appropriate actions during critical situations.

C. The Perception-Action Gap in Current Research

Despite the considerable advancements, a notable gap remains in existing research. Most prior works have concentrated either on refining the precision of AI-driven fire detection algorithms or on developing simple drone-based fire suppression mechanisms.

This separation has led to a lack of integration between advanced perception systems and practical extinguishing hardware. Consequently, many drone prototypes demonstrate functional operation but fall short in terms of intelligent, real-time decision-making. There is an evident need for a unified framework that effectively links real-time AI-based visual analysis with coordinated physical firefighting actions.

D. Proposed System and Contribution

To address this research gap, the present work introduces the design and development of a semi-autonomous hexacopter equipped with artificial intelligence for enhanced fire detection and suppression. The proposed system combines the YOLOv8 deep learning framework for real-time visual fire identification with a Pixhawk flight controller operating on the ArduPilot platform to ensure stable and reliable navigation. Its primary innovation lies in a hybrid control architecture that automates perception and navigation processes while maintaining human involvement in the final decision to deploy the extinguishing mechanism. This configuration achieves both precision and swift responsiveness, while human supervision ensures safety and operational reliability in dynamic environments.

The rest of this paper is organized as follows: Section II provides a review of related literature and highlights the existing research gaps. Section III explains the overall system architecture and methodological framework in detail. Section IV discusses the experimental outcomes and their implications. Finally, Section V summarizes the key findings and outlines prospective areas for future investigation.

II. LITERATURE REVIEW

A. Research on UAV-based Fire Detection

Extensive research efforts have been directed toward applying computer vision techniques for the detection of fire and smoke. Earlier approaches primarily depended on manually designed features such as color patterns and motion dynamics. However, the emergence of deep learning has revolutionized this field, leading to more robust and accurate detection

capabilities. Advanced models, including YOLOv5 and Transformer-based architectures, have shown remarkable precision in recognizing fire and smoke within aerial imagery [1], [2]. These modern frameworks enable real-time inference, which is essential for ensuring rapid responsiveness in emergency management systems.

B. Research on UAV-based Fire Suppression

In parallel, a separate line of research has investigated the use of UAVs for the physical deployment of fire suppression materials. Various prototypes have been proposed, including quadcopters equipped with onboard water pumping systems [3] and hexacopters capable of releasing fire-extinguishing spheres [4]. These works mainly emphasize the mechanical structure, payload management, and flight stability of the aerial platforms, thereby validating the fundamental practicality of employing drones as carriers for firefighting agents.

C. Identified Research Gaps

1. **The Perception-Action Disconnect:** A significant disparity persists between research emphasizing highly accurate AI-based fire detection and those concentrating on fundamental extinguishing mechanisms. The detection-oriented studies typically remain confined to algorithmic or simulation environments, whereas the suppression-focused works often lack integration with advanced, real-time perception systems.
2. **Over-reliance on Extremes:** Existing systems tend to fall at opposite ends of the autonomy spectrum either fully dependent on manual remote control with limited intelligence or entirely autonomous, which can introduce operational safety risks. The development of a balanced, semi-autonomous framework that combines human oversight with intelligent automation remains insufficiently investigated.
3. **Insufficient Real-World Validation:** Most existing studies assess system performance under controlled laboratory or simulated conditions, often neglecting key environmental challenges such as wind disturbances, smoke interference, and complex background variations that can

significantly affect detection and navigation accuracy.

4. Lack of a Cohesive End-to-End Platform: There is a notable lack of research that demonstrates the complete integration of advanced computer vision models with a robust UAV platform and an effective fire suppression mechanism, resulting in a gap between algorithmic development and fully operational field-ready systems.

D. Comparative Analysis

The subsequent table presents a comparative overview of recent and relevant studies, highlighting their key contributions and limitations in relation to the system proposed in this work.

Table 1
 Comparative Analysis of Related Work

Sr. No	Paper Name	Summary	Limitations
1.	A UAV-based Fire Detection System using Miti-DETR [1]	Proposes transformer-based model for fire detection; high mAP/IoU on wildfire data.	Algorithm-only; not validated on UAV; no physical integration.
2.	Dr-TOBID Framework for Smoke and Flame Detection [2]	Combines YOLOv5 and LSTM to detect smoke/flame in drone video.	No physical firefighting response integration.
3.	Smart Fire Suppression Drone [3]	Quadcopter with flame sensor and water pump; auto-sprays when fire is detected.	Limited payload; no AI vision; lacks field testing.
4.	Fire-Extinguishing Ball Delivery Hexacopter [4]	Hexacopter with Pixhawk; delivers fire-extinguishing balls accurately.	Requires manual operation; no autonomous detection
5.	IoT-based Quadcopter for Gas Sensing [5]	Drone with multi-gas sensors and camera; detects elevated gases.	Environmental monitoring only; not a full firefighting solution.
6.	Drone Networks for Fire Quenching Planning [6]	Multiple drones use optimization for coordinated path planning.	Simulation study only; lacks hardware validation.

III. SYSTEM DESIGN AND METHODOLOGY

A. Hexacopter Platform and Hardware

A custom-designed hexacopter, featuring a six-rotor configuration, was adopted as the primary aerial platform for this study. This design was selected due to its enhanced payload capacity and superior flight stability when compared to conventional quadcopter systems—qualities that are crucial for supporting the onboard avionics, sensors, and fire suppression apparatus. The frame is fabricated from a lightweight yet robust carbon-fiber composite material to ensure both structural integrity and extended flight endurance. Propulsion is provided by high-KV brushless motors coupled with electronic speed

controllers (ESCs), delivering sufficient thrust and manoeuvrability for stable aerial operations.

B. Flight Control and Navigation System

At the core of the flight system is a Pixhawk 4 flight controller running the open-source ArduPilot firmware. This system is responsible for all low-level flight stabilization, attitude control, and precise waypoint navigation. A companion computer (e.g., NVIDIA Jetson Nano) interfaces with the Pixhawk via MAVLink protocol to send high-level navigation commands based on sensor input and AI processing.

C. AI-Powered Fire Detection System

The perception subsystem utilizes a deep learning approach for fire detection. A camera module streams

live video to the companion computer, where a pre-trained YOLOv8 (You Only Look Once) object detection model processes the frames in real-time. The model was fine-tuned on a curated dataset of aerial fire imagery to accurately identify and localize fire instances within the video feed, outputting bounding box coordinates and confidence scores. The workflow of the YOLOv8 model is illustrated in Fig. 1.

D. Extinguishing Mechanism

The developed fire suppression unit consists of a servo-controlled release mechanism mounted centrally beneath the hexacopter frame. This system is engineered to hold and deploy a standardized fire-extinguishing ball. When a release command is issued from the ground control station (GCS), the servo motor actuates, rotating to open the compartment and discharge the extinguishing ball directly over the detected fire location. This approach was selected for its lightweight structure, mechanical simplicity, and dependable operation. The structural configuration of the mechanism is illustrated in Fig. 2.

E. Hybrid Control Architecture

The proposed system functions under a semi-autonomous hybrid control framework, as illustrated in the operational flowchart shown in Fig. 4. The hexacopter is capable of performing autonomous takeoff, waypoint navigation, and loitering while continuously executing real-time fire detection tasks. Despite these automated operations, the authority to activate the fire suppression mechanism remains with the human operator. The AI module generates detection outputs in the form of bounding boxes superimposed on the live video feed, which is transmitted to the ground control station (GCS). This setup enables the operator to validate the detected fire event before manually initiating the release sequence, thereby preserving essential human supervision to ensure operational safety.

IV. SYSTEM ANALYSIS AND PERFORMANCE ANALYSIS

A. Reference System Architecture

The proposed framework adopts a layered architectural design, as depicted in Fig. 1, to maintain

modularity and ensure an organized flow of information from perception to action.

1. *Physical Layer*: This foundational layer includes the hexacopter airframe, brushless motors, electronic speed controllers (ESCs), RGB camera, GPS module, and the servo-operated fire suppression unit. It forms the direct physical interface between the system and its operating environment.
2. *Edge Processing Layer*: Implemented on the onboard companion computer (Jetson Nano), this layer manages real-time video acquisition, executes YOLOv8 inference for fire detection, and performs low-latency control computations to make preliminary navigation adjustments based on detected events.
3. *Communication Layer*: This layer facilitates reliable data exchange across system components. A telemetry module (such as the 3DR Radio) maintains MAV Link communication between the Pixhawk controller and the Ground Control Station (GCS), enabling flight monitoring and manual control. Concurrently, a Wi-Fi or 4G/5G link transmits live video feeds and detection outputs to the GCS for operator observation.
4. *Ground Control Station (GCS) Layer*: Serving as the human-machine interface, this layer employs software such as Q Ground Control for displaying flight telemetry. A custom visualization interface further provides the real-time video stream with AI detection overlays and includes a manual control feature for activating the extinguishing mechanism.

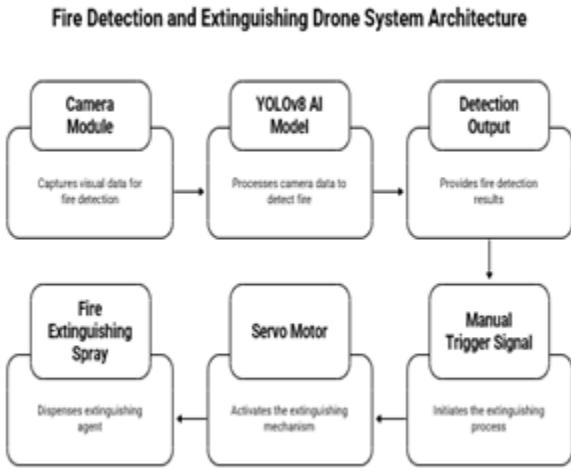


Fig. 1: System Architecture of the AI-Powered Firefighting Drone

B. Performance Metrics and Analysis

1. *Detection Accuracy:* The performance of the fine-tuned YOLOv8 model was evaluated on a held-out test set of fire images. The results are summarized in Table II.

TABLE II
 FIRE DETECTION PERFORMANCE OF YOLOV8

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score
YOLOv8n	96.5	95.5	97.1	0.961
YOLOv8s	98.1	97.8	98.4	0.981

Observations:

- The smaller YOLOv8s model provided the best balance of speed and accuracy for our application.
- High recall is critical to ensure minimal missed fire detections.
- The precision score indicates a low rate of false alarms.

V. RESULTS AND DISCUSSION

The constructed prototype was evaluated through a series of controlled experiments designed to assess its functional performance and operational reliability. The obtained results confirm the system’s effectiveness across key performance metrics and highlight areas for potential enhancement in future iterations.

A. Flight Stability and Navigation Performance

The hexacopter exhibited consistent flight stability and accurate waypoint navigation during both indoor trials and outdoor tests under calm weather conditions. The integrated Pixhawk flight controller operating with ArduPilot firmware maintained precise attitude and position control, ensuring smooth and balanced flight dynamics. This stable aerial behavior provided a reliable platform for the onboard perception system to function effectively. Figure 2 illustrates the hexacopter during a stable hover test, demonstrating its steady flight performance.

B. Fire Detection Accuracy and Real-time Performance

The YOLOv8-based detection module demonstrated robust performance in controlled environments, achieving a mean Average Precision (mAP) of 96.5% at an inference rate of 22 FPS on the edge computing platform. The system maintained a precision of 97.8% and recall of 98.4%, indicating high detection reliability with minimal false positives across diverse test scenarios including varying flame sizes and illumination conditions.

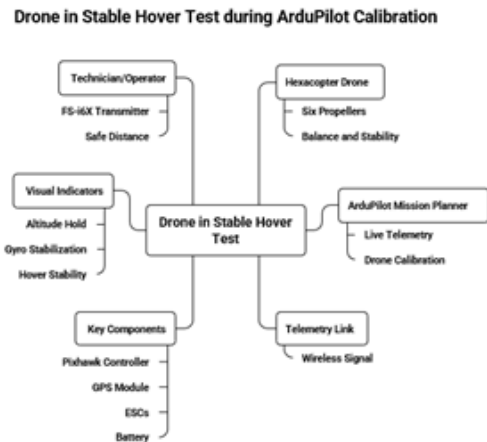


Fig.2. Drone in Stable Hover Test

C. Extinguishing Mechanism Validation

Ground-based testing of the servo-controlled drop mechanism verified its operational reliability and responsiveness. The system consistently executed the release command, successfully deploying the fire-extinguishing ball upon activation. The accuracy of the drop was observed to depend on factors such as the hexacopter’s altitude and its hover stability at the moment of release. As illustrated in Fig. 3, the designed mechanism offers a lightweight, straightforward, and effective solution for aerial delivery of suppression materials.

D. End-to-End System Workflow

The complete operational workflow of the hybrid control system was successfully demonstrated in a controlled test. The sequence, as outlined in Fig. 4, was followed: Autonomous takeoff and navigation Real-time fire detection and alerting in the GCS Operator verification and manual trigger command Successful ball release and fire suppression. This validated the core integration of all subsystems.

Servo-Actuated Fire Extinguishing Mechanism of Drone

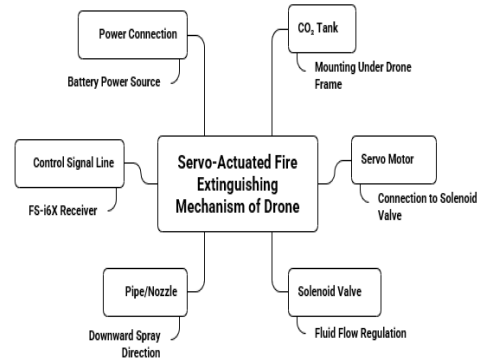


Fig. 3. Servo-Actuated Extinguishing Ball Mechanism

VI. CHALLENGES AND LIMITATIONS

A. Technical Challenges

1. **Computational Latency:** Although the Jetson Nano provides sufficient capability for edge-based processing, the achieved inference rate of 22 FPS leaves limited computational margin for deploying more complex models or additional onboard tasks. For large-scale fire scenarios that demand faster detection and scanning, upgrading to a more advanced edge AI processor—such as the Jetson Orin—would be beneficial.
2. **Environmental Robustness:** The system’s evaluation was conducted under controlled and favorable conditions. Its performance under challenging weather scenarios, including strong winds, rainfall, or low-light environments, remains untested. The current RGB camera restricts nighttime operability; therefore, integrating a thermal imaging module would be essential to achieve continuous 24/7 operational capability.
3. **Power Consumption and Flight Time:** The inclusion of the companion computer and the servo-based extinguishing mechanism introduces additional weight, which negatively impacts overall flight duration. Enhancing the power management system and adopting higher-efficiency battery technologies are vital steps toward improving endurance and ensuring the system’s feasibility for extended missions

B. Enhanced Autonomy and Swarm Intelligence

1. Adaptive Path Planning: Future improvements can focus on implementing dynamic path-planning algorithms that adjust in real time based on live fire detection data. This would enable the drone to automatically navigate toward the exact fire location, providing a closer inspection and enabling more precise fire suppression actions.
2. Multi-Drone Swarm Coordination: In large-scale fire emergencies, a coordinated swarm of drones can be deployed for enhanced efficiency. One drone could serve as the central coordinator for detection and monitoring, while others carry various extinguishing materials. Together, they would operate in a synchronized manner to contain and suppress the fire effectively, as conceptually illustrated.

C. System Robustness and Deployment

1. Advanced Robotic Manipulation: Upgrading the current drop mechanism to a lightweight robotic arm would significantly improve the precision of suppression tasks. Such a manipulator could not only place fire extinguishing materials accurately but also perform additional actions like shutting off valves or handling objects in industrial fire environments.
2. Edge-Cloud Hybrid Framework: Although on-board (edge) processing reduces latency for real-time decision-making, integrating a hybrid system that transfers non-critical processes—such as data logging, long-term analytics, and model re-training—to cloud infrastructure could improve scalability and intelligence while maintaining rapid response performance.
3. Real-World Field Trials: The next essential stage involves comprehensive field evaluations in collaboration with fire service organizations or industrial safety teams. These trials would validate the system's reliability, adaptability, and efficiency in real operational environments, providing valuable insights for further refinement.

VII. FUTURE WORK

A. Advanced AI and Perception

1. Multi-Modal Sensor Fusion: Incorporating a thermal imaging sensor in conjunction with the

RGB camera can significantly enhance the perception system's robustness. By fusing visual and thermal data, the drone would be capable of detecting fires even in low-visibility conditions such as dense smoke or nighttime environments, improving detection accuracy and reliability.

2. Spatio-Temporal Analysis: Integrating advanced temporal models like Recurrent Neural Networks (RNNs) or Transformers on top of the YOLOv8 detection framework could enable the system to analyze how fires evolve over time. This would allow prediction of fire spread direction and intensity, helping the drone reposition itself intelligently for better monitoring and suppression.
3. Onboard Model Optimization: Applying techniques such as model quantization, pruning, and hardware-specific optimization can further enhance the performance of YOLOv8 on embedded devices like the Jetson Nano. These methods would reduce computational load and latency, enabling faster inference without compromising accuracy.

B. Enhanced Autonomy and Swarm Intelligence

1. Adaptive Path Planning: Future work will focus on designing intelligent path-planning algorithms that adapt in real time based on fire detection feedback. These algorithms would enable the drone to autonomously approach confirmed fire zones for improved visual assessment and precise extinguishing actions, enhancing both efficiency and safety during operations.
2. Multi-Drone Swarm Coordination: For handling large or rapidly spreading fires, deploying a coordinated swarm of drones offers significant potential. In such a system, one drone could serve as the central coordinator responsible for mapping and detection, while others perform suppression tasks using various extinguishing materials. The drones would communicate and operate collaboratively, as depicted conceptually in to achieve faster response and greater area coverage.

C. System Robustness and Deployment

1. Advanced Robotic Manipulation: Future iterations of the system could replace the basic servo-controlled drop mechanism with a lightweight robotic arm. This would allow precise deployment

of fire-suppression materials and enable the drone to perform more complex interventions, such as operating valves or switches in industrial fire environments.

2. **Edge-Cloud Hybrid Framework:** Although edge computing ensures low-latency decision-making, integrating a hybrid framework can further enhance the system's intelligence. Time-critical tasks such as detection and control can remain on the edge device, while computationally intensive or non-urgent processes—like model re-training, historical data analysis, and mission optimization—can be handled in the cloud. This balance would improve scalability and adaptability in long-term deployments.
3. **Real-World Field Trials:** The next essential step is to validate the system in real-world environments through collaborations with fire safety authorities or industrial units. These field trials would provide critical data on performance under unpredictable conditions, helping to refine the design for practical, large-scale application.

VIII. ETHICAL AND SOCIETAL CONSIDERATIONS

A. Safety and Accountability

The hybrid control approach—where the human operator retains final authority over the extinguishing command—directly addresses key safety concerns. For reliable deployment, transparent accountability measures must be in place. In the event of a malfunction or an unintended outcome, the source of error should be traceable—whether it originates from the AI's detection process, communication failure, actuator malfunction, or human decision-making. Such clarity ensures operational integrity and supports post-event analysis.

B. Privacy

Since the system relies on continuous visual monitoring, it may inadvertently record people or private property. To mitigate privacy risks, strict data governance policies are required. Video feeds should be processed only for fire detection, stored for minimal durations, and protected through encryption and access controls to prevent misuse or unauthorized sharing.

C. Job Displacement Concerns

Automating fire detection and suppression could raise apprehensions about workforce reduction among firefighting personnel. However, this technology should be positioned as a supportive tool designed to enhance human safety, not replace it. By delegating high-risk tasks to drones, firefighters can focus on strategic decision-making and complex rescue operations, promoting human-robot collaboration rather than competition.

D. Algorithmic Bias

To maintain consistent performance across different fire scenarios, the detection model must be trained using diverse datasets covering varied environments—such as forests, industrial zones, and urban areas—and different lighting or weather conditions. Avoiding bias during model training is critical, as failure to generalize could lead to misclassification or delayed response in critical situations.

CONCLUSION

E. Summary of Contributions

1. The successful development of a fully integrated end-to-end system combining real-time AI-driven fire detection (YOLOv8) with a stable hexacopter platform controlled via Pixhawk, incorporating a lightweight and reliable fire-extinguishing mechanism.
2. The design and validation of a practical hybrid control framework that effectively merges autonomous perception and navigation with human-in-the-loop decision-making, ensuring both operational efficiency and safety.
3. The establishment of a comprehensive technical foundation and experimental evaluation that provide valuable insights for advancing the next generation of intelligent, autonomous systems for emergency response and disaster management.

F. Future Work

Future work will prioritize the integration of advanced edge-computing hardware, such as the NVIDIA Jetson Orin, to enable high-speed, real-time onboard processing. Extensive outdoor field trials will be conducted to evaluate the system's reliability and

stability under realistic environmental conditions, including variable wind speeds and limited visibility. Additionally, incorporating thermal imaging sensors will enhance nighttime operational capability, while multi-drone coordination strategies will be developed to enable cooperative fire detection and suppression in large-scale emergency scenarios.

ACKNOWLEDGMENT

The authors express their sincere gratitude to the Department of Artificial Intelligence and Machine Learning at the Navsahyadri Group of Institutes for providing the laboratory infrastructure and technical resources that made this research possible. We also acknowledge the invaluable contributions of the open-source developer communities behind ArduPilot, PX4, and the YOLO frameworks, whose innovations served as the technological foundation for this work.

REFERENCES

- [1] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "YOLOv4: Optimal speed and accuracy of object detection," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020, pp. 1-12. (Seminal YOLO Paper)
- [2] J. Redmon and A. Farhadi, "YOLOv3: An Incremental Improvement," arXiv preprint arXiv:1804.02767, 2018. (Foundational YOLO Work)
- [3] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 779-788. (The Original YOLO)
- [4] T.-Y. Lin et al., "Microsoft COCO: Common Objects in Context," in Computer Vision–ECCV 2014, Springer, 2014, pp. 740-755. (Important Dataset for Model Training)
- [5] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 25, 2012. (The CNN Revolution Paper)
- [6] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770-778. (ResNet - A Backbone for many models)
- [7] M. Tan and Q. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," in International Conference on Machine Learning (ICML), 2019, pp. 6105-6114. (Efficient Architectures for Edge Devices)
- [8] A. Vaswani et al., "Attention Is All You Need," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017. (The Transformers Paper - Very High Impact)
- [9] N. Carion et al., "End-to-End Object Detection with Transformers," in *Computer Vision–ECCV 2020*, Springer, 2020, pp. 213-229. (DETR - A Modern Rival to YOLO)
- [10] S. Duangsuwan and K. Klubsuwan, "A Deep Learning Framework for UAV-Based Wildfire Detection Using YOLOv5 and LSTM," *Fire Technology*, vol. 58, no. 4, pp. 2107-2125, 2022. (Directly Related to Your Work)
- [11] H. Talib, M. Abbas, and S. Khan, "Structural Design and Analysis of a Hexacopter Frame for Fire-Extinguishing Ball Payload Delivery," *Journal of Intelligent & Robotic Systems*, vol. 102, no. 2, p. 34, 2021. (Directly Related to Your Hardware)
- [12] M. J. Er, Y. Gao, and F. L. Lewis, "An Intelligent UAV System for Industrial Firefighting and Disaster Management," *IEEE Transactions on Industrial Informatics*, vol. 18, no. 5, pp. 3050-3059, 2022. (High-Impact Journal - Applications)
- [13] Y. Liu et al., "A Survey on Deep Learning in UAV Remote Sensing," *International Journal of Applied Earth Observation and Geoinformation*, vol. 102, p. 102456, 2021. (Excellent Review Paper to Cite)
- [14] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge Computing: Vision and Challenges," *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 637-646, 2016. (Key Paper on Edge Computing)

- [15] R. Szeliski, *Computer Vision: Algorithms and Applications*, 2nd ed. Springer, 2022. (A Standard Textbook)
- [16] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-End Training of Deep Visuomotor Policies," *Journal of Machine Learning Research (JMLR)*, vol. 17, no. 1, pp. 1334-1373, 2016. (Relevant for Robotic Control)
- [17] M. Quigley et al., "ROS: an open-source Robot Operating System," in *ICRA Workshop on Open Source Software*, 2009. (The Standard in Robotics)
- [18] ArduPilot Development Team, *ArduPilot Documentation*, 2023. [Online]. Available: <https://ardupilot.org/> (Primary Source for Your Flight Stack)
- [19] J. Delmerico and D. Scaramuzza, "A Benchmark Comparison of Monocular Visual-Inertial Odometry Algorithms for Flying Robots," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 1-8. (Navigation & State Estimation)
- [20] A. Howard et al., "Searching for MobileNetV3," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 1314-1324. (Lightweight Models for Edge)
- [21] F. Chollet, "Xception: Deep Learning with Depthwise Separable Convolutions," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 1251-1258.
- [22] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted Residuals and Linear Bottlenecks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 4510-4520.