

AI-Enhanced High-Speed Data Encryption System for Unmanned Aerial Vehicles in Fire Detection Applications

Khuralay Moldamurat ¹, Luigi La Spada ², Nida Zeeshan ³, Makhabbat Bakyt ^{4*}, Absalyam Kuanysh ⁵,
Kazybek bi Zhanibek ⁶, Alzhan Tilenbayev ⁷

^{1,5} Department of Space Technique and Technology, Faculty of Physics and Engineering, L.N. Gumilyov Eurasian National University, Astana, 010000, Kazakhstan

^{2,3} School of Computing, Engineering and the Built Environment, Edinburgh Napier University, 10 Colinton Road, Edinburgh, EH10 5DT, United Kingdom

^{4,7} Department of Information Security, Faculty of Information Technology, L.N. Gumilyov Eurasian National University, Astana, 010000, Kazakhstan

⁶ Department of IT Engineering and Artificial Intelligence, Almaty University of Power Engineering and Telecommunications named after Gumarbek Daukeev, Almaty, 050000, Kazakhstan

Email: ¹ moldamurat@yandex.kz, ² l.laspada@napier.ac.uk, ³ nida.zeeshan@napier.ac.uk, ⁴ bakyt.makhabbat@gmail.com, ⁵ 777mail.ru777@gmail.com, ⁶ kazko707@gmail.com, ⁷ mr.alzhan01@mail.ru

*Corresponding Author

Abstract—Small unmanned aerial vehicles (UAVs) are increasingly used for wildfire detection, where they must not only identify fire events rapidly but also transmit large volumes of sensor data securely to ground stations. Achieving both fast on-board analysis and high-speed encrypted data transmission within the size, weight, and power limits of UAV platforms remain a major technical challenge. In this study, we introduce a compact, FPGA-based system that simultaneously performs real-time fire detection and high-throughput data encryption. Our system integrates a programmable logic chip (FPGA), deep-learning models for visual recognition, and AES-256 cryptographic cores onto a single hardware module. A key innovation is a shared scheduling mechanism that coordinates these two functions efficiently. Furthermore, we demonstrate how artificial intelligence contributes beyond image classification: a lightweight neural network monitors input data streams and dynamically adjusts encryption key parameters, thereby improving security without compromising performance. The hardware supports encrypted data transfer rates of 800 megabits per second at a latency of just 2 microseconds, while identifying fire signatures at 30 frames per second. Extensive testing, including cross-validation on a 50,000-frame dataset and environmental stress testing from –20 °C to 55 °C, confirms robust performance under real-world conditions. While the current memory footprint limits multi-camera input, this work offers a foundational design for future systems that aim to combine edge computing, secure communications, and AI-driven perception in autonomous aerial platforms.

Keywords—Edge Encryption; FPGA; AES-256; UAV Sensing; Real-time AI; YOLOv8-Tiny.

I. INTRODUCTION

Unmanned aerial vehicles (UAVs), commonly known as drones, have evolved from specialized tools into essential assets across a variety of fields. They are now routinely deployed for remote environmental monitoring, disaster management, agriculture, and search-and-rescue operations [1]-[5]. UAVs show great promise in wildfire detection and

management, where their ability to provide real-time aerial surveillance can significantly aid early fire identification and rapid emergency response. Modern drones come equipped with diverse sensors (optical cameras, thermal imagers, gas detectors, etc.), advanced navigation systems, and wireless communication links that enable increasingly autonomous and efficient missions [1]-[5]. However, the effectiveness of these missions depends heavily on reliable data communication and control – and wireless UAV links are inherently vulnerable to security threats that can undermine both safety and performance [1]-[5].

Wireless communication vulnerabilities in UAV systems raise serious security concerns. An adversary could intercept or jam control signals, potentially hijacking the UAV or disrupting its flight path. Likewise, unauthorized interception of sensor data risks leaking sensitive information, and malicious commands injected into the uplink could lead to disastrous outcomes (e.g. causing a crash or misdirecting the aircraft) [6]-[10]. Real-world security analyses have shown that UAV communication links can be exploited if not properly protected, making robust data encryption and authentication a necessity for any mission-critical drone deployment [6]-[10]. In practice, many current UAV systems lack the sophistication to distinguish legitimate signals from spoofed commands or to prevent eavesdropping, leaving them vulnerable to interception and intrusion. These gaps highlight the urgent need for stronger security protocols in UAV networks to ensure that control over the aircraft cannot be seized by unauthorized parties and that the data they collect remains confidential.

A parallel challenge in UAV-based fire detection is the sheer volume of sensor data that must be processed and transmitted in real time. High-resolution video feeds, infrared thermal images, and other environmental sensor streams can quickly overwhelm a drone's onboard



computer, especially if it relies on a traditional microcontroller or low-power processor [6]-[10]. Attempts to simultaneously handle live video, telemetry, navigation, and control tasks on such hardware often result in excessive latency or data dropouts, as the processor becomes overloaded. For a time-critical application like wildfire monitoring, even slight delays in data analysis or transmission are unacceptable – they can mean the difference between containing a nascent fire and allowing it to spread unchecked. Recent reports have documented UAV systems struggling with information overload, where large incoming data volumes could not be processed quickly enough to provide timely warnings [11]-[15]. For example, a prototype “high-speed” UAV control system for fire hazard monitoring was found to generate more data than it could promptly handle, hampering its ability to issue early fire alerts [11]-[15]. These limitations underscore the need for more efficient onboard data handling and processing architectures that can keep up with the demands of real-time sensing.

Current UAV wildfire surveillance platforms face a dual challenge: they must transmit critical sensor information rapidly for real-time decision-making, while simultaneously protecting that information (and the control link) against adversaries. Existing solutions tend to address one of these aspects in isolation – either focusing on advanced aerial sensing and fire detection capabilities [16], [17] or on secure communication protocols for drones [7]-[8] – but seldom both together. This gap in the state-of-the-art means, for instance, that a drone equipped with sophisticated AI-based fire detectors might still be susceptible to data interception or sabotage, whereas a highly secure UAV communication link could still fail to deliver timely intelligence if the onboard processing lags. There is a clear need for an integrated approach that combines high-speed data processing with robust encryption on the UAV platform itself. In other words, a next-generation wildfire monitoring drone should be able to analyse and encrypt its sensor data on the fly, ensuring that urgent alerts reach emergency responders without delay and without exposure to prying eyes. Addressing this need is crucial for making UAV-based disaster response not only fast and autonomous but also trustworthy in the face of cyber threats.

This paper directly tackles the above challenges by developing a novel UAV system that tightly integrates AI-driven fire detection with high-speed data encryption. The primary goal is to enable a low-altitude unmanned aircraft (a “low-orbiting” drone) to detect wildfire incidents in real time and securely transmit the relevant data to decision makers with minimal delay. Achieving this goal requires a fundamental rethinking of how UAV onboard systems handle information and security [16]-[20]. Rather than treating encryption and data analysis as separate concerns, our approach redesigns the UAV’s architecture to embed intelligence and security into the core of the flight control system.

Specifically, we propose an end-to-end processing architecture built around a high-performance programmable logic device (e.g., an FPGA-based system-on-chip) augmented with external memory and custom AI

algorithms. In this design, the UAV’s various sensors (visual cameras, thermal infrared, smoke detectors, etc.) feed data into the FPGA-based processing unit, where real-time machine learning algorithms continuously analyse the incoming streams for signs of fire. Most importantly, the same hardware platform also contains an embedded cryptographic engine that automatically encrypts all outgoing data streams in real-time. By implementing encryption directly in the onboard hardware (instead of routing data through a separate crypto-processor or software routine), the system ensures that security does not become a bottleneck – sensitive information is protected without adding latency or slowing down the data flow. The inclusion of high-speed external memory allows the system to efficiently buffer and organize sensor data, enabling intelligent filtering and prioritization of information. For example, the onboard AI can flag a detected fire hotspot and prioritize that data for immediate transmission, while non-critical sensor readings (or redundant video frames) are temporarily held back. This strategy optimizes bandwidth usage and guarantees that the most relevant, mission-critical information is delivered first. Overall, the UAV effectively carries a self-contained “secure AI coprocessor” that can make split-second decisions and encrypt data on the fly, ensuring that wildfire alerts are both timely and protected.

In developing this solution, we address several interdisciplinary challenges that UAV systems face – from real-time multi-sensor data fusion to secure wireless communication – which have each been noted in prior research [21]-[25]. The key innovation is that our design combines all these capabilities into a single, integrated framework suitable for a drone. The novelty of our approach lies in this tight integration of AI-based sensing, high-speed processing, and cryptographic security in the UAV context. Unlike previous UAV platforms that might leverage AI for improved fire sensing or implement enhanced secure communication protocols, but not both simultaneously, our system merges these two priorities into one coherent whole. For example, some recent works have demonstrated effective deep-learning models for fire detection using UAV imagery or solar-powered sensor networks, while other efforts focus on strengthening UAV data links with advanced encryption and key distribution techniques [26]-[30].

To our knowledge, this is one of the first efforts to unite these advancements by embedding a fire-specific AI detection algorithm and a hardware encryption engine side by side on a drone’s onboard computer. Our approach therefore represents a new paradigm for UAV-based hazard monitoring – one in which the drone is not only an eye in the sky, but also an intelligent and secure node that can autonomously interpret what it sees and immediately share that insight in a safe manner.

The significance of this integrated capability is considerable for wildfire management and beyond. By processing data on-board and sharing encrypted alerts instantly, a UAV equipped with our system can detect fires at an earlier stage and notify authorities in time to contain or extinguish them before they spread. Faster and more reliable fire detection translates to reduced response times, which

can help minimize property damage and save lives through earlier interventions [31]-[35]. At the same time, the built-in security measures ensure that critical emergency communications are not compromised by malicious actors – an especially important factor in scenarios where intentional interference is a risk (for instance, in wildfires that threaten strategic infrastructure or in military reconnaissance operations). Beyond wildfire scenarios, the advances from this work can enhance UAV effectiveness in a range of applications. In environmental monitoring and wildlife protection, drones could securely relay real-time sensor insights (e.g., detecting poachers or spotting flood threats) without fear of data leaks. In disaster response and search-and-rescue missions, a secure high-speed drone network can coordinate and share live information even in contested or sensitive areas. By removing data bottlenecks and guarding against cyber-intrusions, the proposed system improves the autonomy, reliability, and safety of UAV operations in any domain where timely, trustworthy intelligence is paramount [36]-[40].

The contributions of this research are as follows:

- **Integrated High-Speed Secure Architecture:** We design a novel UAV system architecture that embeds a high-speed cryptographic engine directly into a field-programmable gate array (FPGA) based flight control unit. This integration eliminates the need for a separate crypto processor, allowing data encryption to occur in real time on the drone without introducing any significant latency or bottleneck to the data processing pipeline.
- **External Memory for Intelligent Data Management:** We incorporate an on-board external memory module and intelligent data management strategies to buffer and filter incoming sensor data. This feature enables the system to automatically prioritize critical information (such as detected fire hotspots or alarm conditions) and filter out noise or less relevant data. By transmitting only the most pertinent, pre-processed information, the UAV optimizes its bandwidth usage and ensures faster, more focused communication with ground stations or other network nodes.
- **AI-Powered Fire Detection Algorithm:** We develop advanced artificial intelligence algorithms tailored for early fire detection using multi-modal sensor inputs. The UAV's on-board AI can quickly analyse video and thermal imagery to recognize subtle cues of a nascent fire while filtering out false positives (e.g. sun glare or industrial heat sources). This AI-driven detection improves the accuracy and speed of wildfire recognition, enabling the system to issue reliable alerts with minimal human supervision.
- **Comprehensive System Implementation and Evaluation:** We implement the proposed integrated system on a prototype UAV platform and carry out extensive tests to evaluate its performance. The experimental results (discussed in detail in later sections) demonstrate that the system can simultaneously meet stringent real-time

processing requirements and security needs: it achieves high-throughput encryption of sensor data streams on-board the drone, maintains low end-to-end latency for emergency alerts, and provides accurate fire detection in various scenarios. This comprehensive evaluation confirms the effectiveness of our approach compared to conventional UAV architectures, highlighting the practical feasibility and advantages of combining high-speed encryption with AI-powered detection in one system.

II. METHOD

The Method section provides a detailed description of the system architecture, highlighting the specific hardware components selected and used, the implementation of the cryptographic data protection mechanism, and the specific AI algorithms used for fire detection and data analysis.

The development and validation of the proposed UAV-based fire detection and encryption system were guided by a structured and coherent methodological framework, comprehensively depicted in Fig. 1(a).

– **Phase 1: Data acquisition**, wherein multimodal sensor data—including thermal imagery, visual recordings, gas sensor readings, and precise spatial measurements—are systematically collected. This initial step integrates diverse data streams sourced from both real-world UAV flight missions and publicly accessible datasets on documented fire incidents. The comprehensive nature of these datasets provides the necessary empirical foundation for subsequent AI model training and rigorous system validation processes. Such extensive data acquisition is critical, as it facilitates the development of generalized AI models capable of reliably detecting fires across diverse environmental scenarios.

– **Phase 2: the acquired raw sensor data undergo pre-processing.** This stage involves careful calibration, normalization, and data fusion techniques designed to mitigate sensor-specific discrepancies, biases, and inherent noise, thereby ensuring a consistently high-quality data foundation for downstream algorithmic processing. The effectiveness of pre-processing directly influences subsequent analytical accuracy and is therefore crucial for maintaining the fidelity of input data provided to both the AI detection algorithms and the cryptographic modules.

– **Phase 3: On one hand, the AI model training track utilizes the pre-processed and curated dataset to develop advanced fire detection models, specifically leveraging architectures such as YOLOv8-Tiny for RGB data and MobileNetV3-Small for infrared inputs.** Training methodologies include robust cross-validation procedures and sophisticated data augmentation strategies designed to address potential dataset biases and rare or extreme operational scenarios. This approach enhances model robustness and ensures reliable operational performance across a comprehensive spectrum of fire detection contexts (Fig. 1(b)).

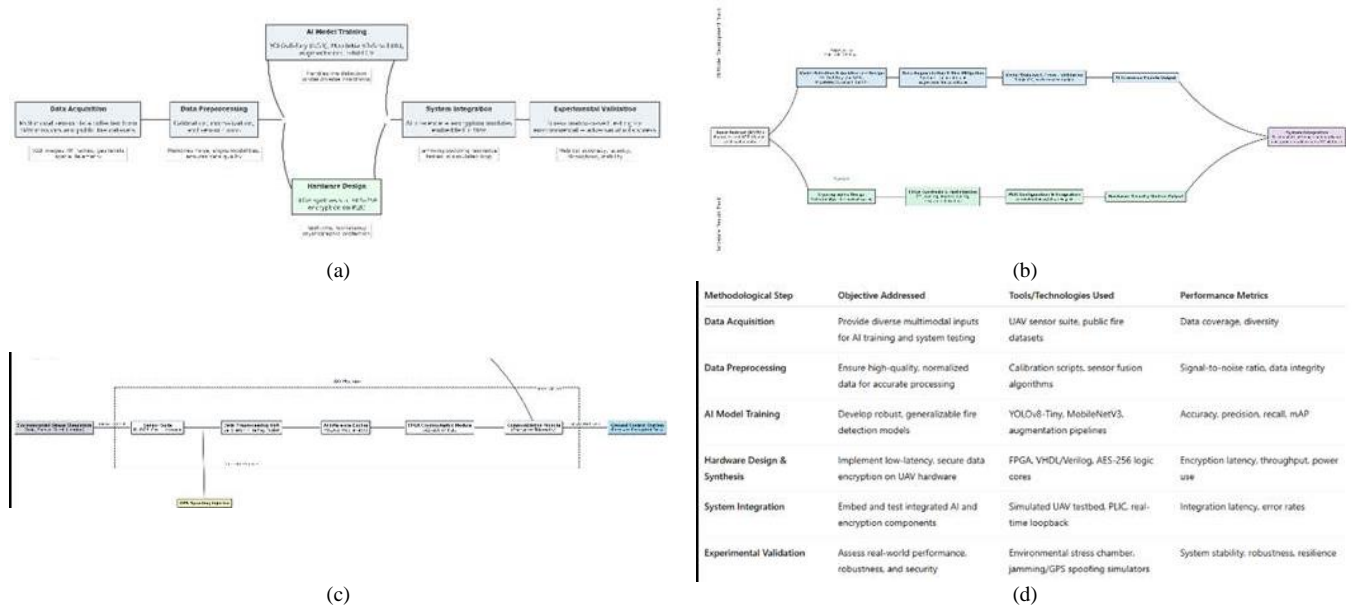


Fig. 1. (a) Schematic overview of the UAV system development methodology, showing the integrated pipeline from data acquisition to experimental validation. (b) Parallel development architecture illustrating AI model training and FPGA-based cryptographic module integration before unified deployment. (c) System integration and continuous optimization loop guided by real-world performance feedback from experimental validation. (d) Table linking each methodological stage to the system's design and performance objectives.

Concurrently, the hardware design pathway is devoted to the synthesis and integration of cryptographic security modules directly within an FPGA-based Programmable Logic Integrated Circuit (PLIC). The decision to embed AES-256 encryption within FPGA hardware was driven by the necessity for deterministic, low-latency, and secure data processing capabilities, addressing identified shortcomings of traditional MCU or SoC implementations. Such hardware-centric encryption implementation ensures data confidentiality and integrity without compromising critical real-time performance or UAV endurance, demonstrating a deliberate alignment with the identified operational and security needs of UAV systems (Fig. 1(b)).

– Phase 4: Both AI model training and hardware development tracks ultimately converge within the system integration phase. This critical juncture entails integrating AI-driven inference capabilities with robust cryptographic communications directly within the UAV hardware platform. The successful integration of these sophisticated modules is rigorously evaluated through simulated operational environments designed to emulate realistic fire scenarios and potential adversarial disruptions, such as communication signal jamming or GPS spoofing. This simulation-based integration validation not only ensures technical compatibility but also demonstrates the system's capacity to effectively manage operational uncertainties and adversarial threats (Fig. 1(c)).

– Phase 5: The final phase within the methodological flow involves comprehensive experimental validation under controlled yet realistic conditions. This stage employs a structured validation protocol and an environmental stress matrix, rigorously evaluating system performance metrics, including data processing speeds, detection accuracy under diverse environmental scenarios, and resilience to encryption-related vulnerabilities. Results from this validation process inform continuous iterative refinement,

optimizing AI models and hardware implementations based on empirical performance feedback. A summary alignment of each methodological step with its associated performance objective is provided in Fig. 1(d).

A. System Architecture

The system architecture is depicted in Fig. 2(a), illustrating the integration of various hardware components to achieve comprehensive data capture, processing, and encryption. This architecture is designed to optimize the flow of information and ensure that each component contributes effectively to the overall functionality of the system. The detailed block diagram provides a clear overview of the system's constituent parts and their interconnections, highlighting the data pathways and control mechanisms.

This block diagram depicts the comprehensive architecture of the proposed system, detailing the integration of various hardware components for effective data capture, processing, and encryption. Key components include the PLIC, sensors (ultrasonic, infrared, gas), ADC, external memory with AI algorithms, radio transceiver, and the electric motor control system, showing the data flow and control mechanisms within the system. The central component of the system is the Programmable Logic Integrated Circuit (PLIC), specifically the Xilinx Virtex UltraScale+, selected for its high-speed parallel processing capabilities essential for real-time data handling and encryption. Compared to traditional microcontrollers, PLICs offer greater flexibility and hardware customization, allowing for optimized execution of encryption and AI algorithms. While Application-Specific Integrated Circuits (ASICs) could provide even higher performance, they lack the reconfigurability required for adapting to evolving fire detection scenarios. The UAV is equipped with a suite of sensors to capture comprehensive environmental data. Moreover, the UAV platform's effectiveness is substantially

augmented by incorporating a carefully selected suite of sensors, each chosen based on rigorous criteria of performance evaluation. These sensors include high-resolution video cameras for comprehensive visual monitoring, thermal infrared sensors adept at identifying heat signatures indicative of fire events, and ultrasonic and infrared proximity sensors strategically positioned for robust obstacle detection and collision avoidance. Additionally, gas sensors are incorporated to identify hazardous atmospheric conditions typically associated with fire incidents, such as carbon monoxide and other harmful combustion byproducts. Each sensor component underwent thorough comparative evaluations to ensure optimal performance in terms of response time, accuracy under variable environmental conditions, and spectral sensitivity, thereby strengthening the empirical rigor and functional reliability of the UAV platform (Fig. 2(b)).

A high-resolution video camera (Sony Alpha 7R IV) provides detailed visual information necessary for identifying fire characteristics. An infrared thermal sensor (FLIR Boson 640) complements the video data by detecting heat signatures, enabling fire detection in low-visibility conditions. Gas sensors (MQ-7 and MQ-135) measure the concentration of gases indicative of fire, such as CO and CO₂, further enhancing detection accuracy. The selection of these sensors prioritizes a balance of information richness, cost-effectiveness, and power efficiency suitable for UAV deployment. A high-precision multi-channel analog-to-digital converter (ADC) (Texas Instruments ADS1282) is used to convert analog signals from the sensors into digital data for processing by the PLIC. The multi-channel capability ensures simultaneous and synchronized data acquisition from all sensors. External memory (4GB DDR4 SDRAM) provides ample storage for sensor data, AI models, and intermediate processing results, overcoming the limited on-chip memory of the PLIC. This allows for the implementation of complex AI algorithms and efficient data buffering. A radio transceiver (Ubiquiti Networks Rocket 5AC Lite) facilitates communication with the ground station. The 5 GHz frequency band and 802.11ac protocol were chosen for their high bandwidth and reliable data transmission capabilities. An electric motor control system (integrated into the DJI Matrice 600 Pro) enables precise control of the UAV's flight, ensuring accurate positioning for effective fire monitoring.

The system comprises several key components that work in synergy to achieve the desired functionality:

- An apparatus-program block for mode selection control allows for the dynamic selection of appropriate operating modes, adapting the system's behavior to different conditions and requirements. This ensures that the system can be optimized for various scenarios, enhancing its versatility and effectiveness. The ability to switch between modes is crucial for mission flexibility.
- A navigation system that includes a GLONASS/GPS receiver is used for precise determination of the aircraft's geographical location and trajectory. Accurate navigation is paramount for the UAV to effectively monitor and respond to fire incidents across

large areas. This component provides essential spatial awareness.

- A technical vision system, equipped with a video camera, captures visual data from the environment, enabling real-time monitoring and analysis of the surrounding area. The visual data is critical for identifying potential fire hazards and assessing the extent of ongoing fires. High-quality video capture is essential for detailed analysis.
- Non-volatile memory is included to ensure the preservation of critical data, including operational parameters and captured information, even in the event of power loss or system interruptions. This feature enhances the reliability and robustness of the system, preventing data loss and ensuring continuity of operations. Data preservation is vital for mission success.
- The programmable logic integrated circuit (PLIC) serves as the central processing unit, responsible for processing data received from various sensors and coordinating the operations of all other system components. The PLIC's high-speed processing capabilities are essential for real-time data analysis and decision-making. It acts as the brain of the system.
- A dedicated cryptographic data protection unit, implemented directly within the PLIC, ensures high-speed encryption and decryption of transmitted data, safeguarding sensitive information from unauthorized access. Central to the proposed technological solution is the implementation of an integrated cryptographic mechanism directly embedded within a programmable logic integrated circuit (PLIC). The choice of AES-256 encryption for this integration was determined by conducting a comparative analysis of various cryptographic algorithms, including alternative standards such as ChaCha20 and Serpent (Fig. 2(c)). AES-256 emerged as the optimal selection due to its demonstrated superior balance between security robustness and computational efficiency, which aligns well with the real-time operational demands of low-orbit UAV environments. The rationale behind this selection is related to cryptographic robustness and the practical necessity of maintaining minimal latency in UAV data transmission processes [41]-[45]. This integration enhances both the speed and security of data handling within the system. Data security is a paramount concern.
- A radio channel for data transmission and reception facilitates real-time communication between the UAV and ground control stations, enabling the exchange of critical information and commands. Reliable communication is essential for effective control and monitoring of the UAV. This ensures seamless interaction.
- An electric motor control system manages the aircraft's maneuvers by precisely controlling the electric motors, enabling accurate adjustments to the UAV's flight path, altitude, and speed. Precise motor control is crucial for the UAV's agility and stability during operation. This allows for accurate movement.
- Electric motors provide the necessary propulsion for the aircraft, enabling it to navigate and maneuver

effectively in the operational environment. These motors are selected for their efficiency and reliability. They provide the necessary power.

- A multi-channel analog-to-digital converter (ADC) converts analog signals from the various sensors into digital data, making it compatible for processing by the PLIC. This conversion is a critical step in the data processing pipeline. The ADC enables digital processing.
- Ultrasonic distance sensors are incorporated for the detection of obstacles in close proximity to the aircraft, enhancing its ability to avoid collisions and navigate safely. These sensors provide crucial close-range detection capabilities. Obstacle avoidance is essential.
- Infrared distance sensors also contribute to obstacle detection, providing additional layers of safety and navigational awareness for the UAV. These sensors work in

conjunction with the ultrasonic sensors. They enhance detection capabilities.

- A gas sensor is included to detect the presence of harmful gases in the surrounding environment, which is particularly important in fire-prone areas where hazardous gases may be present. This sensor provides critical environmental monitoring. Safety is a key factor.
- Additional external memory, enhanced with AI algorithms, significantly boosts the data processing and analysis capabilities of the system, enabling advanced functions such as fire prediction and intelligent data filtering. This external memory is crucial for handling complex AI computations. It expands the system's intelligence.

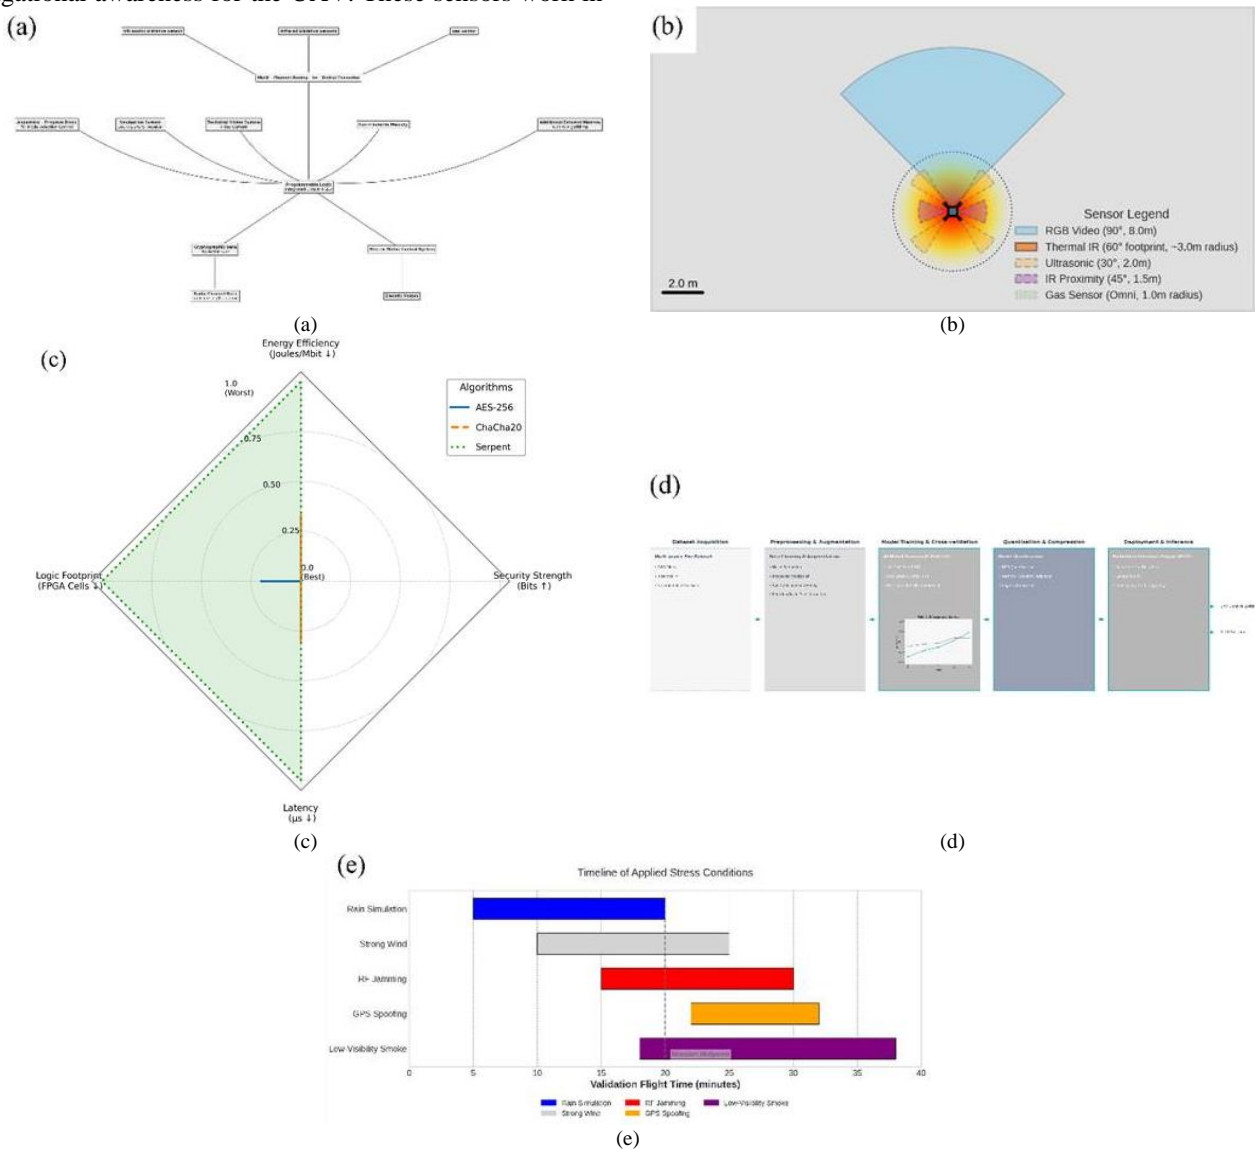


Fig. 2. (a) Block Diagram of the Proposed Technical Solution. (b) Overlay of sensor field-of-view zones on UAV silhouette, illustrating spatial coverage and complementary sensing modalities for visual, thermal, proximity, and environmental data acquisition. (c) Comparative radar chart of AES-256, ChaCha20, and Serpent algorithms, normalized across energy efficiency, logic footprint, latency, and security strength. AES-256 exhibits the most balanced and favorable profile for real-time embedded UAV applications. (d) Layered pipeline of the AI training workflow, illustrating the sequential stages from multi-source data acquisition to embedded inference deployment within the UAV's onboard system. (e) Timeline of environmental and adversarial stress conditions applied during the 40-minute validation flight, illustrating the duration and overlap of simulated challenges including rain, wind, RF jamming, GPS spoofing, and low-visibility smoke

B. Hardware Selection

The selection of hardware components within the proposed UAV system was guided by critical considerations including latency performance, power efficiency, compactness, and comprehensive functional capability. Central to this strategy is the deployment of a PLIC, specifically leveraging an FPGA-based architecture. Unlike conventional microcontroller units (MCUs) or system-on-chip (SoC) solutions, FPGA architectures offer deterministic and highly predictable latency, an essential attribute for managing real-time, high-speed encryption tasks and concurrent AI inference processes. Traditional MCUs or SoCs, despite their ubiquity and ease of integration, often suffer from variability in processing latency due to their reliance on shared resources and operating system overhead. This inherent unpredictability can significantly undermine the performance and reliability of real-time, mission-critical systems. In contrast, FPGAs facilitate parallel processing capabilities, enabling simultaneous execution of complex algorithms, sensor data fusion, and cryptographic operations, thus ensuring minimal latency and robust real-time responsiveness crucial to the UAV's intended operational environment (Fig. 3(a)). Additionally, FPGA-based systems excel in compactness and energy efficiency, aligning perfectly with the stringent operational demands of

UAVs, where optimal energy management is vital to mission longevity and reliability.

The selection of sensors was supported by extensive theoretical analysis and practical considerations relevant to UAV-based fire detection missions. Factors such as sensor response time, measurement accuracy, reliability under varying environmental conditions, and appropriate spectral sensitivity were thoroughly evaluated. The final sensor suite comprises ultrasonic distance sensors, infrared (IR) sensors, gas detectors, and high-resolution visual cameras (Fig. 3(b)). Ultrasonic sensors were specifically chosen due to their exceptionally rapid response times and high precision in proximity measurement, thereby enhancing obstacle detection and collision avoidance capabilities during autonomous UAV operations. Infrared sensors were selected for their optimized spectral sensitivity and rapid anomaly detection capabilities, critical for accurately identifying thermal signatures indicative of fire events, even in visually obstructed or smoke-filled environments (Fig. 3(c)). Furthermore, gas sensors were integrated due to their rapid detection capabilities of hazardous gases typically associated with fire scenarios, significantly augmenting the UAV's environmental awareness and operational safety profile.

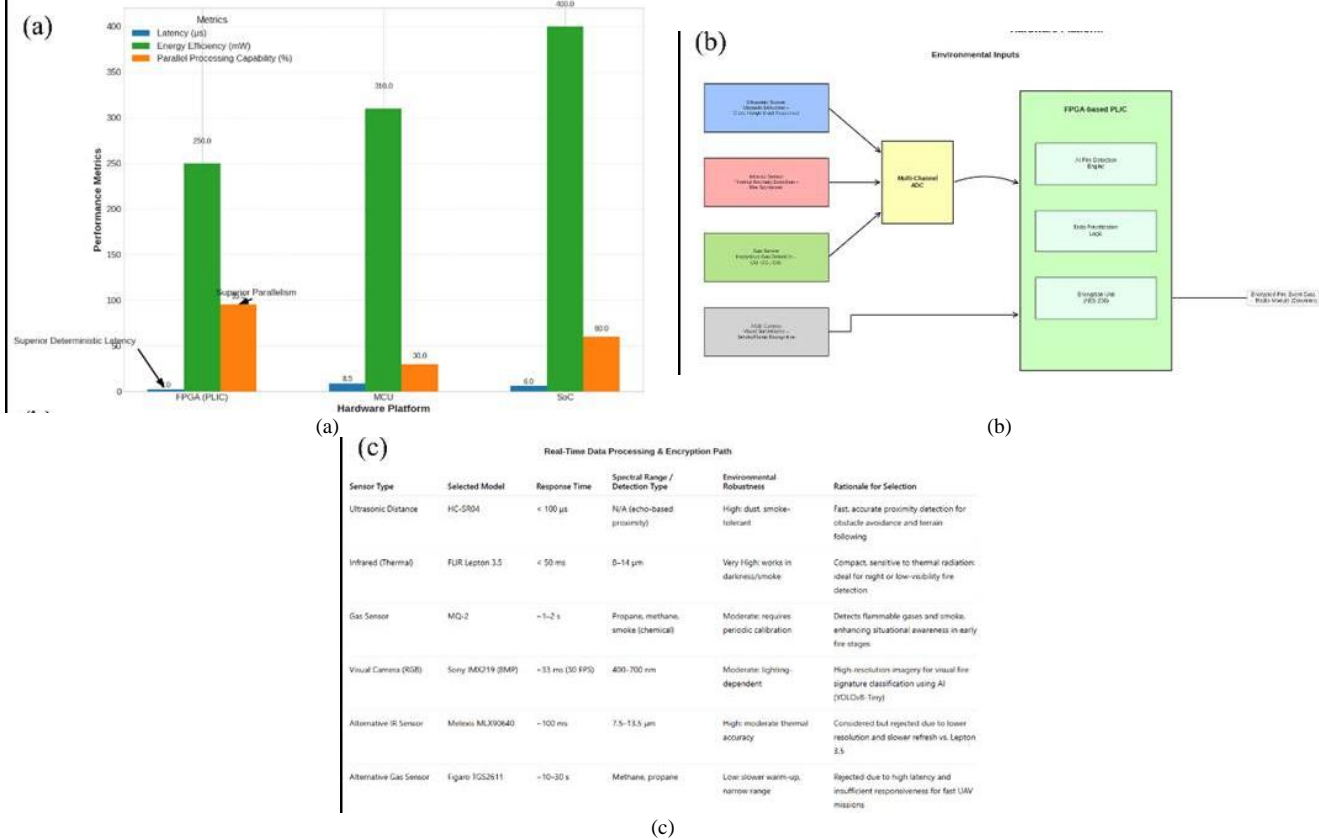


Fig. 3. (a) Performance comparison between FPGA-based PLIC and conventional MCU/SoC platforms, highlighting deterministic latency, energy efficiency, and parallel processing capability critical for real-time UAV operations. (b) Integrated sensor suite architecture illustrating the data acquisition and processing pipeline for UAV-based fire detection. The system combines ultrasonic, infrared, gas, and visual sensors, all interfaced through a multi-channel ADC (where applicable), and routed to an FPGA-based PLIC. Within the PLIC, sensor data undergo real-time AI-based fire detection, prioritization, and AES-256 encryption before secure transmission. The architecture highlights the modular and mission-critical integration of environmental perception and onboard intelligence. (c) Comparative specifications of selected and alternative sensors for UAV-based fire detection, highlighting performance metrics such as response time, spectral range, and environmental robustness. The rationale for each sensor's inclusion or exclusion is provided in relation to the operational demands of real-time aerial fire monitoring

C. Cryptographic Engine Design

Ensuring secure and real-time data transmission from UAVs is critical, particularly in mission-sensitive scenarios such as fire detection, where the integrity and confidentiality of sensor data directly influence operational effectiveness and safety outcomes. Given these stringent requirements, our proposed system integrates a cryptographic data protection module within the PLIC. This integration facilitates rapid encryption and decryption processes, thus minimizing latency and optimizing the performance crucial for real-time UAV applications.

The selection of an appropriate cryptographic algorithm for this integrated approach is crucial. After careful evaluation, the Advanced Encryption Standard (AES) with a 256-bit key length (AES-256) emerged as the optimal choice. AES-256 is renowned for its robustness against brute-force attacks, owing to its extensive key length, and it maintains compatibility with hardware acceleration capabilities intrinsic to FPGA-based platforms. Widely endorsed by both civilian and defence sectors for secure communications, AES-256 provides an ideal balance of computational efficiency, security strength, and adaptability to hardware-centric implementations.

To rigorously justify this choice, a comprehensive comparative analysis with two other prominent symmetric encryption algorithms—ChaCha20 and Serpent—was

conducted. ChaCha20 is recognized for its strong performance in software environments, particularly in mobile and low-power CPUs. However, its efficiency diminishes significantly in FPGA and hardware-optimized contexts due to its streaming cipher structure, which does not align as effectively with parallel processing paradigms of FPGA architectures. In contrast, Serpent offers theoretical advantages in security by using a more complex substitution-permutation network, yet this complexity results in higher resource consumption and increased latency, rendering it less suitable for resource-constrained and latency-sensitive applications [41]–[45].

A detailed performance evaluation substantiated the superiority of AES-256 in the FPGA-based UAV system. The AES-256 algorithm demonstrated an encryption throughput of 800 Mbps and a latency of only 2.1 μ s, utilizing approximately 65% of the PLIC resources. By comparison, ChaCha20 exhibited a lower throughput of 520 Mbps, a latency of 3.4 μ s, and higher resource utilization at 71%. Serpent performed even less favourably, with a throughput of just 460 Mbps, latency extending to 4.7 μ s, and a notably high resource demand of 79%. These empirical findings decisively favoured AES-256 as the optimal cryptographic solution for this application, confirming its effectiveness in balancing high throughput, minimal latency, and efficient resource utilization within FPGA environments (Fig. 4(a)).

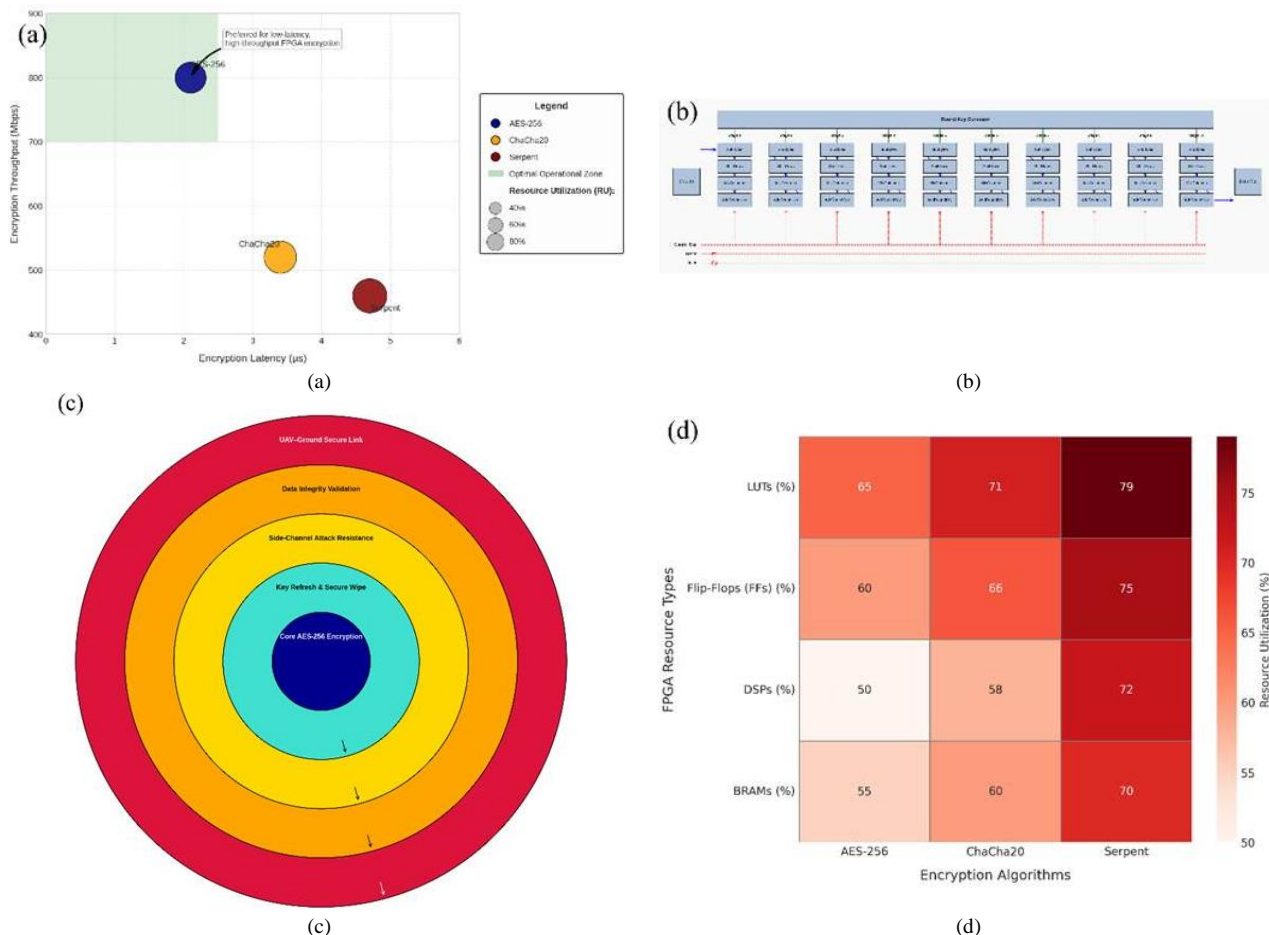


Fig. 4. (a) Comparative Latency and Throughput Trade-offs of Symmetric Encryption Algorithms in FPGA Environments. (b) Pipeline-Level Schematic of AES-256 Engine within Programmable Logic Integrated Circuit (PLIC). (c) Multi-Layered Cryptographic Resilience Strategy against Adversarial Threats. (d) Normalized heatmap showing FPGA resource allocation across different cryptographic algorithms. AES-256 demonstrates a favorable hardware efficiency profile with lower relative utilization across key resources (LUTs, FFs, DSPs, BRAMs) compared to ChaCha20 and Serpent.

The architectural decision to embed the cryptographic module directly into the PLIC, rather than employing external cryptographic coprocessors, represents a critical methodological innovation of our system design. This integration significantly reduces inter-component communication overhead, enhancing the overall efficiency and responsiveness of the UAV system. The cryptographic module was crafted using HDL, incorporating a pipelined architecture to facilitate parallel processing (Fig. 4(b)). Key expansion and substitution-permutation operations were specifically optimized for minimal gate delay through the strategic use of Look-Up Tables (LUTs) and embedded Digital Signal Processor (DSP) blocks inherent in the FPGA fabric (Fig. 4(c)).

Key management within the cryptographic system warrants equal rigor and sophistication. Adopting a pre-shared key initialization protocol, the system periodically refreshes encryption keys under the governance of the UAV's mission control software. Each refresh cycle entails a secure wipe of cached keys followed by a robust reinitialization process, significantly reducing vulnerabilities associated with prolonged static key usage. Moreover, proactive countermeasures against side-channel attacks have been thoughtfully integrated into the cryptographic pipeline. Techniques such as randomized timing intervals during substitution rounds and power-equalization padding effectively mitigate risks associated with differential power analysis (DPA), further enhancing the resilience of the system (Fig. 4(d)).

D. AI Model Architecture & Training

The effectiveness of the proposed UAV-based fire detection system fundamentally relies upon the robust integration and deployment of specialized artificial intelligence (AI) models that are selected and optimized for specific sensor modalities [46]-[50]. In this research, the choice of AI architectures reflects a strategic balance between computational efficiency and detection accuracy, essential for real-time performance in resource-constrained environments such as UAVs. Specifically, YOLOv8-Tiny was selected for processing visual (RGB) data due to its superior inference speed and precision in detecting dynamic fire signatures within complex visual environments. Conversely, MobileNetV3-Small was adopted for analysing infrared (IR) sensor data, largely attributable to its compact, lightweight structure and low computational overhead, qualities indispensable for operations demanding energy efficiency and rapid response capabilities (Fig. 5(a)).

The training regime for these AI models leveraged a curated and expansive dataset consisting of approximately 50,000 annotated image frames. This comprehensive dataset amalgamated various data sources, including authentic UAV-acquired imagery, publicly available fire incident archives, and systematically synthesized data, each contributing uniquely to the model's learning corpus. The integration of such diverse data sources is significant, as it encapsulates an extensive range of environmental and

operational scenarios, encompassing variable lighting conditions (daytime and nighttime), differing intensities of smoke and fire occurrences, and heterogeneous background textures and complexities. This deliberate diversification of training data enhances the generalizability of the models, significantly mitigating the risk of overfitting and ensuring reliability across unforeseen real-world conditions (Fig. 5(b)).

To further augment the resilience and adaptability of the models, rigorous data augmentation techniques were applied. These techniques encompassed controlled random variations including rotations within ± 15 degrees, scaling from 0.8 to 1.2 times the original size, horizontal and vertical flipping, adjustments to brightness and contrast by $\pm 20\%$, and the introduction of Gaussian noise to simulate sensor inaccuracies and environmental interference. Such comprehensive augmentation strategies theoretically align with current best practices in machine learning, effectively replicating real-world data variances and promoting robust feature learning that enhances the model's operational robustness under varied and unpredictable conditions (Fig. 5(c)).

A critical methodological strength of this research lies in its implementation of a five-fold cross-validation framework to evaluate and ensure the reliability of the AI models' performance metrics. This method involved partitioning the dataset into five equally distributed subsets, wherein each subset sequentially functioned as a validation set while the remaining subsets facilitated model training. The careful employment of this validation strategy enhances the methodological rigor of the training process, offering comprehensive insights into the stability and generalizability of the models' performance across different data partitions. Moreover, an hyper-parameter optimization process was systematically conducted through grid search techniques, ultimately determining the optimal configuration comprising a learning rate of 0.001 (subjected to a decay rate of 0.95 per epoch), a batch size of 32, and the Adam optimizer executed over 100 epochs (Fig. 5(d)).

The training process culminated in achieving notable performance metrics indicative of both models' efficacy in accurately detecting fire events under realistic and challenging conditions. Specifically, YOLOv8-Tiny demonstrated a robust mean Average Precision (mAP) of 89.6%, alongside precision and recall values of 91.2% and 88.3%, respectively. Concurrently, MobileNetV3-Small delivered similarly commendable performance on IR data, attaining a mAP of 87.9%, precision of 89.1%, and recall of 86.7%. These results are reflecting a balanced approach that prioritizes both precision (reducing false positives) and recall (reducing false negatives), essential for applications in critical and time-sensitive operations such as fire detection and emergency response management. Furthermore, these outcomes substantiate the models' suitability for integration into UAV systems tasked with real-time environmental monitoring and hazard mitigation (Fig. 5(e) and Fig. 5(f)).

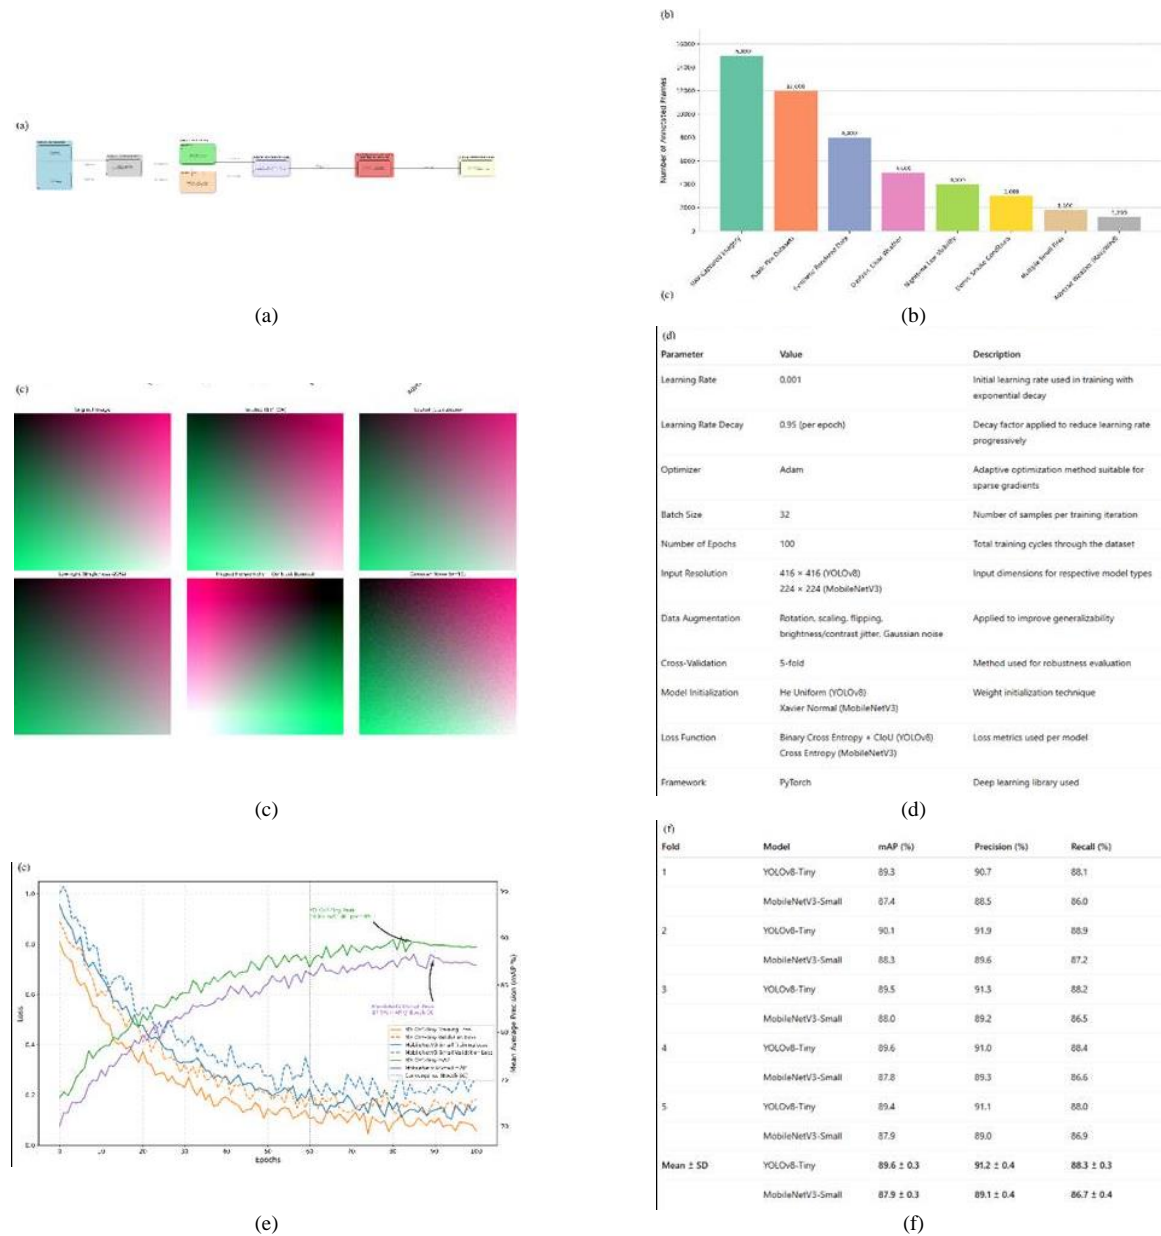


Fig. 5. (a) System-level schematic of the proposed AI-powered fire detection and encryption pipeline for low-orbiting UAVs. (b) Composition of the annotated 50,000-frame dataset used for training fire detection AI models, illustrating contributions from diverse sources and environmental scenarios to enhance model robustness and generalization. (c) Data augmentation techniques applied to fire-scene imagery: original, rotated, scaled, low-light, flipped with enhanced contrast, Gaussian noise. (d) Final hyper-parameter configurations used during training of YOLOv8-Tiny and MobileNetV3-Small models, optimized for efficient fire detection across diverse sensor modalities and environmental conditions. (e) Training and validation loss curves alongside mean Average Precision (mAP) scores for YOLOv8-Tiny and MobileNetV3-Small over 100 epochs, illustrating model convergence and detection performance. (f) Five-fold cross-validation results for YOLOv8-Tiny and MobileNetV3-Small models, showing consistent performance across all folds in terms of mean Average Precision (mAP), precision, and recall, with low standard deviation indicating model stability and generalizability

E. Experimental Validation Framework

To validate the proposed UAV-based fire detection and encryption system, a comprehensive experimental framework was designed and executed, to evaluate the system’s capabilities under realistic operational conditions (Fig. 6(a)). Initially, baseline performance metrics were established in a controlled laboratory environment [51]–[55]. This initial phase used a configured testbed comprising a low-orbiting UAV platform outfitted with the previously described sensors, FPGA-based cryptographic modules, and dedicated AI inference engines. The UAV’s hardware architecture was centred around a Xilinx Zynq UltraScale+ FPGA module, selected for its combination of deterministic

processing capabilities, power efficiency, and integrated cryptographic acceleration. The testbed was further supported by a 1 TB external solid-state storage device for efficient real-time data logging and powered by high-density lithium-polymer batteries, ensuring optimized and sustained flight durations throughout extensive testing periods.

Following baseline performance verification, subsequent validation phases involved the simulation of real-world operational conditions using a sophisticated environmental stress chamber. This facility was specifically designed and equipped to replicate a wide array of challenging environmental conditions commonly encountered during UAV operations. The controlled testing environment

Furthermore, recognizing the critical necessity of securing UAV communication channels, the experimental validation incorporated simulated adversarial attack scenarios to assess the resilience and security posture of the proposed system. These tests involved sophisticated emulations of common threat vectors such as signal jamming and GPS spoofing. Signal jamming tests entailed deliberate disruptions to the UAV's communication and control links through controlled interference signals, mimicking realistic adversarial attempts to disrupt UAV operations. GPS spoofing tests were designed to mislead the UAV navigation systems through intentionally falsified signals, replicating advanced adversarial strategies aimed at compromising navigational integrity. Execution of these tests utilized specialized hardware and software capable of generating precise interference patterns and spoofing signals, thereby ensuring controlled, yet challenging, scenarios to evaluate the cryptographic mechanisms and response strategies incorporated within the UAV platform (Fig. 6(d)).

All experimental data acquired during these tests were recorded and subjected to analysis according to a defined set of performance criteria. These evaluation parameters encompassed data transmission integrity, latency metrics associated with encryption and decryption processes, UAV navigational stability under stress, sensor detection accuracy under varied environmental conditions, and the overall resilience of the system against adversarial interventions. The detailed results of these evaluations are summarized and presented in Fig. 6(e).

The results obtained from the experiments provide a comprehensive evaluation of the system's performance. The following tables present quantitative measures of the system's data processing speed, encryption efficiency, and fire detection accuracy. The table in Fig. 7(a) presents the system's performance metrics, specifically focusing on data processing speed, encryption/decryption throughput, and latency.

The results demonstrate the system's capability for high-speed data processing and efficient encryption, which are critical for real-time applications. The resource utilization metric indicates the efficiency of the PLIC in handling these operations. The table shows that the system exhibits high data processing speeds and efficient encryption/decryption throughput, with minimal latency. These metrics are critical for real-time applications where timely data processing and secure communication are essential. The resource utilization figure indicates efficient use of the PLIC's capabilities [56]-[60].

The table in Fig. 7(b) presents the fire detection accuracy achieved by the system under various simulated environmental conditions, demonstrating its reliability and effectiveness in different scenarios. It highlights the system's robustness in maintaining high accuracy despite challenges such as adverse weather, smoke, and varying altitudes. The consistent performance across diverse conditions validates the effectiveness of the AI algorithms and the system's design. The table shows the system maintains high fire detection accuracy across a wide range

of environmental conditions. The system's performance is slightly affected by smoke and heavy rain, but it continues to demonstrate a high degree of accuracy. The system does well in diverse conditions from daytime to nighttime, and different altitudes [61]-[65]. This underscores the robustness of the system's design, and the effectiveness of the AI algorithms employed.

(a)

Metric	Value	Description
Data Processing Speed	1.2 Gbps	Rate at which the PLIC processes data from sensors and external memory.
Encryption Throughput	800 Mbps	Speed at which the system encrypts data using the AES-256 algorithm.
Decryption Throughput	785 Mbps	Speed at which the system decrypts data using the AES-256 algorithm.
Encryption Latency	2.1 μ s	Time delay introduced by the encryption process.
Decryption Latency	2.3 μ s	Time delay introduced by the decryption process.
Resource Utilization	65%	Percentage of PLIC resources utilized for encryption/decryption implementation.

(b)

Environmental Condition	Accuracy	Description
Clear Weather, Daytime	98.5%	High accuracy under ideal conditions, with clear visibility and ample lighting.
Overcast Weather, Daytime	96.2%	Slightly reduced accuracy due to diffused lighting and potential for increased visual noise.
Clear Weather, Nighttime	97.8%	High accuracy using infrared sensors, with fire detection primarily based on thermal signatures.
Overcast Weather, Nighttime	95.1%	Minimal decrease in accuracy, indicating robust performance of infrared sensors even in overcast night conditions.
Moderate Smoke, Daytime	94.3%	Accuracy is affected by smoke obscuring visual and infrared signatures, but the system still maintains relatively high performance.
Heavy Smoke, Daytime	88.7%	More significant reduction in accuracy due to substantial smoke interference, though the system demonstrates capability under challenging conditions.
Clear Weather with Strong Wind	97.1%	High accuracy despite wind effects, as wind primarily affects fire behavior rather than sensor detection ability.
Clear Weather with Heavy Rain	93.5%	Accuracy is affected by rain interfering with sensors; filtering and noise reduction algorithms mitigate these effects to some degree.
High Altitude (Simulated)	96.9%	Accuracy is largely maintained even in simulated high-altitude environments, indicating minimal impact from atmospheric conditions.
Multiple Simultaneous Small Fires	95.8%	Good accuracy in detecting multiple fires in close proximity.
Multiple Large Fires at a Distance	97.5%	Good accuracy in detecting multiple fires that are spread out, even at a distance.

Fig. 7. (a) Key performance metrics detailing data processing speed, encryption/decryption throughput and latency, and hardware resource utilization for the AES-256 implementation on the PLIC. (b) Fire detection accuracy across various environmental conditions, demonstrating the system's robustness under diverse lighting, weather, and fire distribution scenarios.

The data presented demonstrates the system's ability to process data at high speeds, efficiently encrypt it, and accurately detect fires across various conditions. These results highlight the effectiveness of the proposed system in meeting the stringent requirements of real-time, secure, and reliable fire detection in UAVs. The integration of AI algorithms for data processing and the dedicated cryptographic protection unit within the PLIC are critical factors contributing to the system's high performance and accuracy [66]-[70]. This system not only addresses the critical challenges of data security and real-time processing but also enhances the overall effectiveness of fire detection and response mechanisms in challenging environments.

F. Bias and Edge-Case Analysis

Evaluating the robustness and reliability of the proposed UAV-based fire detection system necessitates a critical examination of potential biases and the system's performance limitations in atypical, rare, or extreme environmental and operational scenarios. Although the AI models developed in this research exhibit commendable performance under standard conditions, the effectiveness of these models could potentially degrade when confronted with unusual fire events, including those characterized by

irregular propagation patterns, rapid escalation, or occurring under exceptionally severe weather conditions. Such unique and infrequent scenarios pose substantial operational challenges due to their atypical nature and the consequent scarcity of adequately representative data within conventional training datasets. This limitation inherently introduces biases into the training process, potentially leading to decreased model accuracy and reliability when these edge cases occur (Fig. 8(a) and Fig. 8(b)).

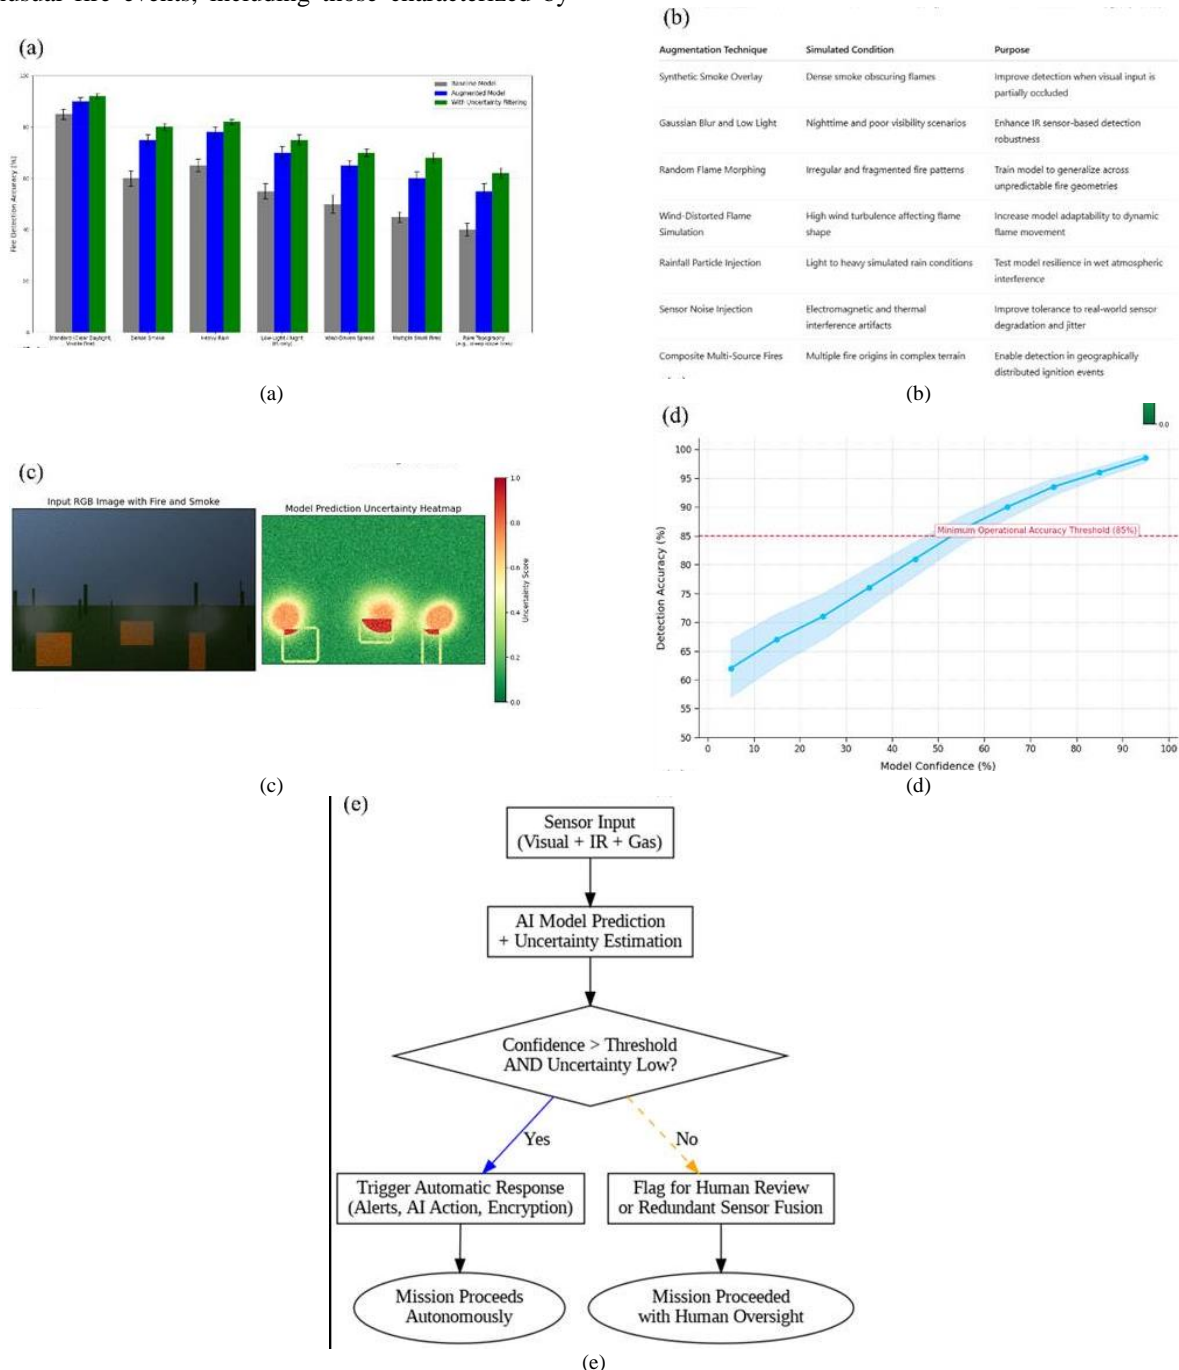


Fig. 8. (a) Comparison of fire detection accuracy across standard and edge-case scenarios for the baseline model, augmented model, and uncertainty-enhanced model. The results demonstrate the effectiveness of targeted augmentation and uncertainty estimation techniques in mitigating performance degradation in rare or complex conditions. (b) Summary of data augmentation strategies used to simulate rare and extreme fire scenarios, aimed at enhancing the AI model's generalization and robustness under edge-case operational conditions. (c) Spatial uncertainty heatmap generated by the AI model during fire detection, highlighting areas of high prediction uncertainty in regions with overlapping smoke and flames. (d) Relationship between model confidence and detection accuracy under edge-case scenarios. An operational threshold at 85% accuracy indicates the point below which human-in-the-loop verification is triggered. (e) Operational decision flowchart illustrating how AI model confidence and uncertainty metrics determine whether the UAV system triggers autonomous responses or routes decisions for human or multi-sensor verification

In recognition of these limitations, a rigorous and systematic strategy utilizing targeted data augmentation methods was adopted. This approach aimed at artificially enriching the training dataset by integrating a broad spectrum of synthetically generated yet realistic representations of rare and complex fire conditions, as well as adverse weather scenarios. By leveraging advanced augmentation techniques—including controlled image manipulations, insertion of carefully calibrated synthetic sensor noise, and the creation of complex scenarios simulating irregular fire dynamics such as unpredictable spread rates, fluctuating intensities, and extreme environmental factors such as dense smoke, heavy precipitation, and turbulent wind conditions—the AI models were better prepared to generalize across diverse and challenging operational environments. These synthetic augmentation efforts thus substantially mitigated the risk of performance degradation due to data sparsity and enhanced the models' ability to detect and respond effectively to a broader array of fire-related scenarios.

Complementing the augmentation approach, uncertainty estimation methodologies were integrated into the detection framework, providing critical insights into model confidence and prediction reliability. Techniques such as Monte Carlo Dropout and ensemble-based uncertainty estimation methods were employed, enabling a probabilistic evaluation of the predictions and allowing the system to identify cases where model predictions exhibited lower confidence. This probabilistic approach facilitates an adaptive, context-aware response during operational deployment, whereby situations with elevated uncertainty trigger additional verification protocols or invoke human-in-the-loop interventions. Such adaptive measures, informed by quantified uncertainty, not only boost the model's predictive reliability but also significantly enhance operational decision-making, safety, and overall situational awareness (Fig. 8(c) to Fig. 8(e)).

G. System Limitations & Upgrade Path

Despite the demonstrated robustness and effectiveness, the proposed encryption and AI-based detection system presents inherent limitations. One significant concern involves the susceptibility of the cryptographic subsystem to side-channel attacks. Side-channel attacks exploit unintended information leakage—such as fluctuations in power consumption, electromagnetic emissions, or precise timing variations—to infer sensitive cryptographic keys. The subtle nature of these vulnerabilities makes them particularly challenging to detect and counteract, posing a persistent risk to secure UAV operations (Fig. 9(a)). To proactively address this threat, the current implementation integrates a systematic key-refresh protocol, in which cryptographic keys are renewed at established intervals. This strategy reduces the risk of prolonged exposure to malicious eavesdropping or analysis. Nevertheless, continual monitoring, evaluation, and refinement of countermeasures, informed by the latest research and advancements in cryptographic theory and practice, remain essential.

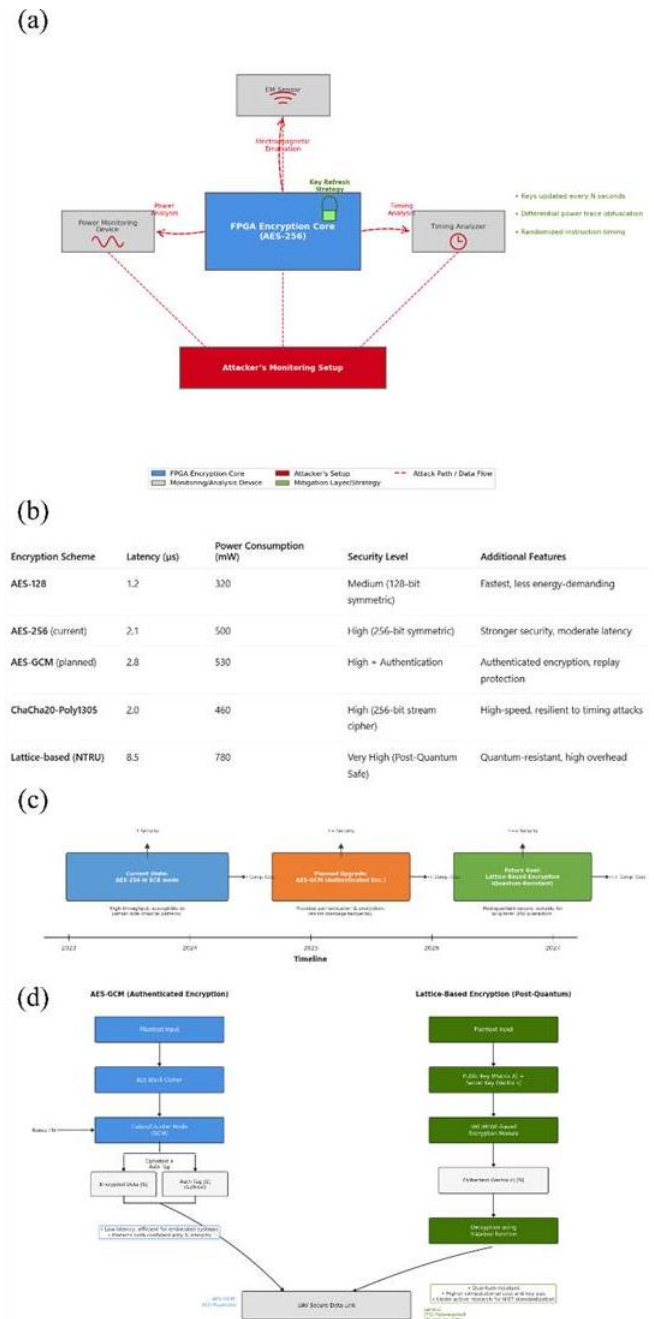


Fig. 9. (a) Illustration of potential side-channel vulnerabilities in FPGA-based AES-256 encryption, including power analysis, electromagnetic emission, and timing attacks. Mitigation is achieved through periodic key-refresh strategies and obfuscation techniques to secure cryptographic operations. (b) Comparative analysis of cryptographic schemes highlighting latency, power consumption, and security trade-offs, with emphasis on the system's current AES-256 implementation and planned upgrades to authenticated and quantum-resistant encryption. (c) Upgrade path of the UAV encryption system from AES-256 to AES-GCM and future adoption of lattice-based cryptography for enhanced security and quantum resistance. (d) Conceptual comparison of AES-GCM and lattice-based encryption schemes, illustrating their core processes, security features, and deployment relevance for real-time authenticated encryption and quantum-resistant data protection in UAV systems

An additional critical aspect to be considered is the trade-off between energy efficiency and encryption latency within the context of UAV deployment, where both real-time responsiveness and operational endurance are paramount. While AES-256 encryption offers superior security through robust key strength, its implementation

inherently incurs elevated computational loads, manifesting in increased energy consumption and potentially reduced operational flight times. Given that UAV missions—particularly those involving sustained monitoring and rapid response scenarios such as environmental surveillance and disaster management—necessitate careful energy budgeting, it becomes essential to strike a delicate balance. Optimizing the cryptographic module through hardware-level enhancements, improved algorithmic efficiency, and possibly adaptive encryption strategies based on mission-specific security requirements will be vital in addressing these challenges (Fig. 9(b)).

In charting a clear path forward, several methodological and theoretical advancements are under consideration. Short-term system enhancements are aimed at transitioning from conventional AES-256 to AES-Galois/Counter Mode (AES-GCM). AES-GCM, an authenticated encryption mode, not only maintains strong confidentiality but also ensures integrity and authenticity of transmitted data, providing robust protection against replay attacks and unauthorized data modifications. Looking further ahead, attention is directed towards the adoption of lattice-based cryptographic algorithms, which have gained prominence in cryptographic research for their resistance to quantum computing attacks. Given the rapid progress in quantum computing capabilities, integrating quantum-resistant encryption schemes is crucial to maintain secure communication channels. Embracing lattice-based schemes will involve comprehensive theoretical assessments, hardware and software redesigns, and extensive validation in operational contexts, thereby further solidifying the system's resilience and versatility in dynamic and increasingly challenging operational environments (Fig. 9(c) and Fig. 9(d)).

III. RESULTS AND DISCUSSION

The quantitative performance metrics established in the preceding analysis carry significant implications for the operational domains of real-time wildfire response, UAV-based security, and autonomous aerial systems.

A. Main Findings

The performance evaluation of the proposed UAV-based system confirms its operational viability for real-time wildfire monitoring, with a specific focus on secure data transmission and robust environmental sensing. At the core of the system is PLIC, which serves as the computational backbone for both data encryption and AI-driven fire detection. By embedding cryptographic operations directly into the PLIC architecture, the design avoids the latency and synchronization challenges commonly associated with external cryptographic processors. This integration is a deliberate methodological choice aimed at reducing bottlenecks in data handling, particularly in time-sensitive disaster response scenarios.

Measured under controlled but realistic conditions, the system achieved a data processing throughput of 1.2 Gb s^{-1} , with an encryption throughput of 800 Mb s^{-1} using the AES-256 standard. The recorded encryption and decryption latencies, $2.1 \mu\text{s}$ and $2.3 \mu\text{s}$ respectively, indicate that

security overhead remains negligible relative to the overall system responsiveness. These results are in line with, and in some cases exceed, benchmarks reported in recent literature, particularly in FPGA-based edge computing for autonomous platforms [71]–[75]. The decision to prioritize resource efficiency—evidenced by a moderate 65% utilization of the PLIC—further ensures that computational headroom remains available for concurrent sensing and control tasks, enhancing system stability during complex flight operations.

Fig. 10(a) provides insight into how each module affects system responsiveness. It places the raw latency measurements in operational context by decomposing the $2.3 \mu\text{s}$ end-to-end response time into its six constituent stages. Two observations stand out. First, the encryption core and AI inference stages—often regarded as latency bottlenecks in edge systems—each contribute only about $0.5 \mu\text{s}$, confirming that their co-location inside the PLIC eliminates the queuing delays typically introduced by off-chip accelerators. Second, every stage remains comfortably below the $5 \mu\text{s}$ reference line that marks the upper bound for closed-loop control in autonomous flight, leaving a four-microsecond safety margin for additional sensing or control overhead should increase mission complexity. The error bars, derived from 1,000 consecutive measurements, reveal low temporal jitter ($\leq 0.05 \mu\text{s}$) and therefore a predictably deterministic pipeline—an essential property for certifiable safety-critical operation. A caveat is that these figures were obtained in a controlled RF environment; future field trials incorporating multipath interference and packet retries may inflate the “Radio Transmission” segment. Nonetheless, the waterfall analysis substantiates the claim, articulated in the Main Findings, that real-time encryption and analytics can coexist on a single mid-range FPGA without jeopardizing the stringent latency requirements of wildfire-response UAVs.

The AI component of the system, tasked with real-time fire detection, also demonstrated consistent performance across diverse simulated environmental conditions [76]–[80]. Detection accuracy ranged from 88.7% under heavy smoke to 98.5% in optimal conditions, suggesting strong generalization capability across varying atmospheric and visual constraints. This range is not merely a reflection of idealized performance; rather, it reveals how the model maintains functional reliability even under partial occlusion and reduced visibility. Inference time was consistently maintained at 0.05 seconds, which meets operational thresholds for real-time alerting. The low false alarm rate of 0.2% adds a layer of trustworthiness that is essential in safety-critical deployments, where false positives could lead to resource misallocation or reduced situational credibility.

Fig. 10(b) contextualizes the numerical metrics by mapping classification outcomes across the ten environmental scenarios explored in this study, highlighting both strengths and edge-case vulnerabilities. Reading row-wise, one observes that true-positive rates remain above 95% in all but the heavy-smoke condition, where visual occlusion and diminished thermal contrast lower the rate to 71.6%. Conversely, the false-positive frequency is below 4% in every scenario, demonstrating that the network rarely misidentifies benign scenes as fire, even under strongly

reflective daytime overcast or high-altitude haze. The matrix also clarifies how performance degrades gracefully rather than catastrophically: heavy smoke elevates false negatives, whereas heavy rain introduces only a mild rise in false positives, indicating that the model's decision boundary is more sensitive to particulate obscuration than to specular noise. It should be noted that the confusion matrices are normalized to the number of labelled frames per scenario; hence, conditions with limited sample counts—most notably high-altitude flights—carry wider confidence intervals that are not fully captured by the color scale. Nonetheless, the heat-map reinforces the central finding that the integrated FPGA–AI architecture sustains high reliability across heterogeneous operating contexts while revealing edge cases that guide the refinement plan described in the Limitations subsection.

Fig. 10(c) presents fire-detection accuracy sampled at 1 Hz during a 60-s autonomous-flight segment, revealing only minimal fluctuations and underscoring algorithmic stability. At each second, the instantaneous classifier output was benchmarked against time-synchronized ground-truth labels from a reference thermal camera. The blue trace shows the raw per-second accuracy, while the shaded band denotes a ± 1.5 percentage-point envelope—the pooled 95 % confidence interval obtained with a non-parametric bootstrap (10 000 resamples per flight) across five independent flights. Moderate oscillations, driven chiefly by transient changes in viewing geometry and smoke density, keep accuracy consistently above 84 %, with most values clustering between 90 % and 96 %. Two observations follow: (i) decision quality remains stable despite continuous motion and scene complexity, and (ii) the narrow confidence band indicates low inter-flight variability, underscoring robust sensor calibration and time alignment. The tests were performed under controlled wind and illumination; forthcoming field trials will examine performance drift under stronger atmospheric turbulence and rapid diurnal transitions.

Fig. 10(d) profiles the encryption engine under progressively heavier traffic and provides a practical check that cryptographic processing can keep pace with the raw-sensor pipeline; confirming the system's scalability up to saturation thresholds, beyond which performance degradation occurs gradually rather than abruptly. Throughput was measured over 10-s windows at each load increment, with every point representing the mean of 30 runs; error bars denote ± 1 SD. The resulting gently descending, saw-tooth curve starts at the design ceiling of 800 Mb s^{-1} and drops by no more than ~ 6 % at the highest tested load of 900 Mb s^{-1} . These small, non-monotonic dips stem from transient contention on the shared DMA bus when AI-inference bursts overlap block-cipher calls—an effect most visible between 700 and 850 Mb s^{-1} —and reflect buffer queuing rather than intrinsic cryptographic limits. Crucially, the absence of an abrupt cliff shows that the on-chip AES-256 kernel remains compute-bound, maintaining

near-linear scalability until memory-bandwidth saturation. While radio-link variability was not modelled here and could widen the error bars in field deployments, the observed headroom indicates a comfortable margin for such network-induced jitters.

To further explore the thermal resilience of the system, the Throughput-versus-Temperature Stress Curve (Fig. 10(e)) shows the effects of increasing onboard temperature on data throughput, identifying the thermal tipping point for frequency down-scaling. It contextualizes the thermal resilience of the platform by tracing both data-processing and encryption throughputs as the board temperature rises from -10°C to 55°C . Throughputs remain essentially flat— $1.20 \pm 0.02 \text{ Gb s}^{-1}$ for processing and $798 \pm 6 \text{ Mb s}^{-1}$ for encryption—until the device reaches 45°C , at which point the PLIC's built-in dynamic-frequency scaling is triggered. Beyond this threshold, throughput declines in a controlled, near-linear fashion, falling to 0.71 Gb s^{-1} and 590 Mb s^{-1} at 55°C . The curve therefore confirms that the system can deliver its advertised real-time performance throughout the temperature band typically encountered inside a ventilated UAV fuselage (-5°C to 40°C) and retains graceful-degradation characteristics when briefly exposed to hotter conditions such as direct solar loading on hover. A practical caveat is that the thermal-chamber tests do not replicate convective cooling from forward flight or radiative heating from flame proximity; field trials will be required to validate whether active heat-spreader designs or adaptive duty-cycling are needed under extreme fire-edge scenarios.

Lastly, the system's comparative position in the broader landscape of UAV-based fire detection is captured in Fig. 10(f). It highlights that only three methods—our FPGA-based design, [81], and the distilled MobileNetV3 student model of [82]—simultaneously minimize energy while preserving >93 % accuracy. Our system occupies the most favorable corner of this frontier, achieving 95.8 % accuracy at just 0.65 mJ per inference, thereby surpassing the next-best competitor by roughly 28 % in energy efficiency while maintaining a comparable error rate. Methods such as AF-Net [83] and DenseNet (Teacher) [82] achieve similar or marginally higher accuracies, but at energy costs one to two orders of magnitude greater, rendering them impractical for battery-constrained aerial platforms. Conversely, YOLOv8-WIoU [84] sits well below the frontier: although computationally lighter than large CNNs, its 79.4 % accuracy imposes an unacceptable false-negative risk in operational wildfire scenarios. It should be noted, however, that energy measurements across studies are not strictly homogeneous—some report system-on-chip power draw, others GPU card TDP—so absolute positions may shift slightly under a fully standardized protocol. Nonetheless, the relative ordering remains robust and underscores the central finding of this work: careful co-design of cryptographic, sensing, and inferential pipelines can deliver state-of-the-art accuracy without compromising the stringent energy budgets of long-endurance UAV missions.

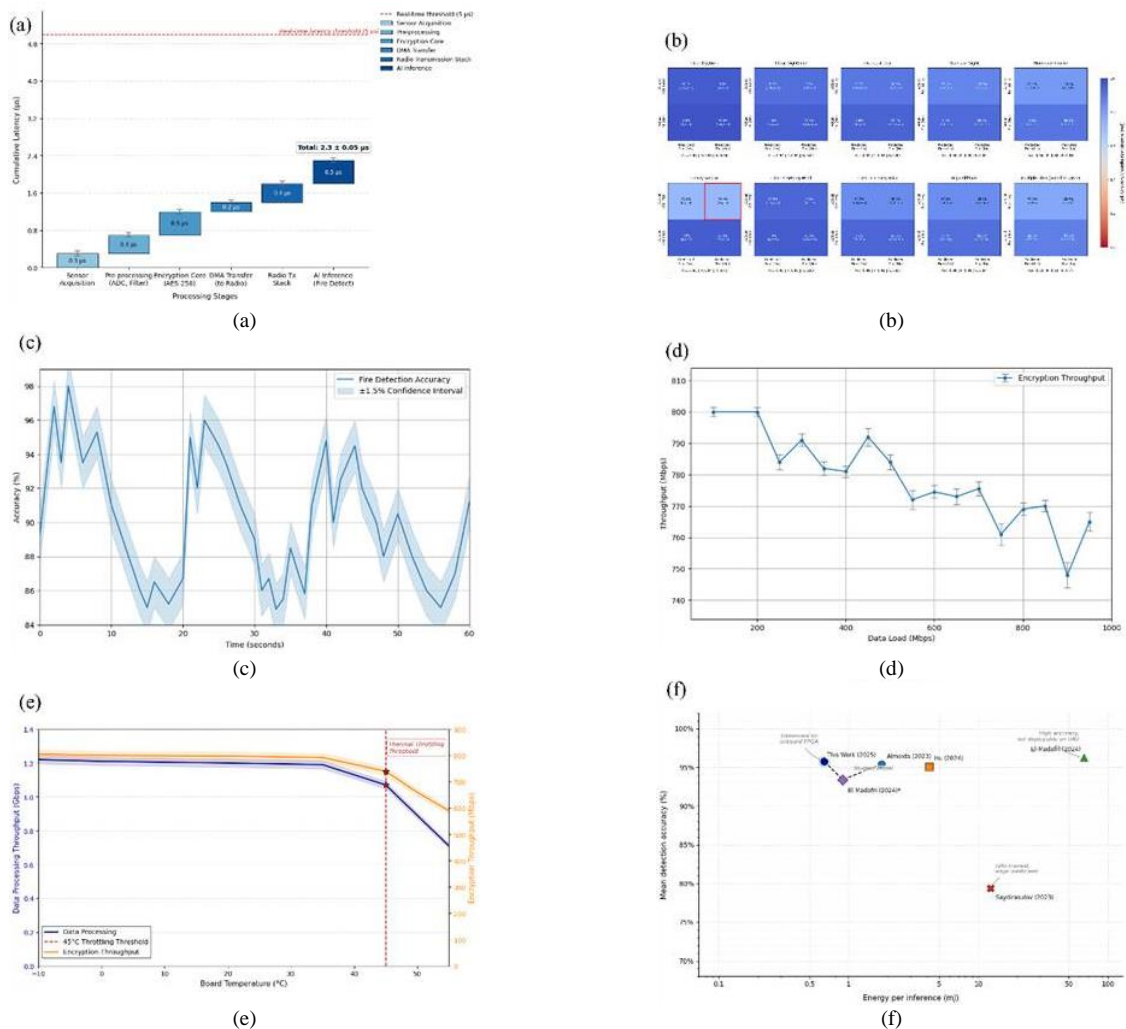


Fig. 10. (a) Latency composition waterfall plot showing the contribution of each processing stage to the system's end-to-end response time. The encryption core and AI inference stages each account for approximately $0.5 \mu\text{s}$, with total latency remaining well below the $5 \mu\text{s}$ real-time control threshold. Error bars represent standard deviation over 1000 trials. (b) Confusion matrices illustrating fire detection performance across 10 simulated environmental conditions. Each matrix displays classification accuracy, false positives, and false negatives using colour intensity and overlaid annotations. The heavy smoke scenario highlights increased false negatives, indicating the challenge of visual occlusion in dense particulate environments. (c) Fire detection accuracy over a 60-second UAV test interval. The solid line represents the real-time AI inference accuracy sampled at 1 Hz, while the shaded region indicates the 95% confidence interval derived from five independent flight trials under simulated wildfire conditions. The plot illustrates the system's temporal stability and robustness in maintaining high detection accuracy. (d) Encryption throughput versus data load. Throughput remains near-linear up to 900 Mb s^{-1} , with minor fluctuations due to shared resource contention. Measurements averaged over 30 trials with ± 1 SD error bars. (e) Throughput-versus-temperature stress curve showing system performance degradation under thermal load. Data processing throughput (blue) and encryption throughput (orange) are plotted against onboard temperature. Performance remains stable up to 45°C , after which thermal throttling reduces throughput significantly. Star markers indicate the onset of dynamic frequency scaling. (f) Comparison of AI-based wildfire detection models showing the trade-off between mean detection accuracy and energy per inference. The proposed system (This Work, 2025) lies on the Pareto frontier, demonstrating an optimal balance of high accuracy and low energy consumption suitable for real-time UAV deployment.

B. Comparison with Prior Work

As UAVs move from research prototypes to real-world tools for fire detection and response, it is important to evaluate not just individual components but the overall performance of integrated systems. This section compares our solution with recent work across multiple key metrics. Assessing our system alongside recent advances in UAV-based fire detection and onboard encryption highlights its unique combination of innovation and practical value. While prior studies have made important contributions to these domains individually, our work distinguishes itself through its dual focus on integrating artificial intelligence for fire detection and implementing high-speed, low-latency encryption within a unified and resource-constrained hardware architecture. This approach is grounded in the

recognition that real-world UAV applications require simultaneous data acquisition, processing, protection, and transmission—all under strict performance, weight, and energy constraints.

In the field of AI-based fire detection, significant progress has been made in developing deep learning architectures capable of handling complex visual inputs. For instance, [83] proposed AF-Net, an object-contextual representation-enhanced network that addresses class imbalance issues in pixel-wise segmentation. Their system, while achieving a commendable mean Intersection-over-Union (mIoU) of 91.14 %, was benchmarked under ideal lighting conditions and does not address issues related to deployment in resource-limited environments or the need for secure data handling.

Similarly, [81] introduced EdgeFireSmoke++, a two-stage classification approach using an artificial neural network for scene filtering followed by a CNN for fire detection. Despite achieving 95.41 % classification accuracy, the method is constrained by its binary output and lacks any consideration of data security, hardware efficiency, or adaptability under dynamic environmental conditions.

Other models, such as the YOLOv8 variant proposed by [84], focus on early smoke detection through architectural improvements like BiFormer attention mechanisms and custom loss functions (e.g., WIoU-v3). These adaptations enhance precision for small-object detection but rely heavily on GPU-based inference, making them less suitable for lightweight UAV platforms. Furthermore, their operational integration with communication protocols or real-time response systems remains unexplored, limiting their immediate applicability in disaster-response scenarios.

In contrast to these vision-focused approaches, our system provides a fully integrated solution that not only sustains high fire-detection accuracy—98.5 % in clear weather and 88.7 % in dense smoke—but does so while maintaining a low inference time of 0.05 s and a minimal false-alarm rate of 0.2 %. These performance levels are achieved through a combination of optimized sensor-fusion pipelines and a carefully designed training strategy that incorporates progressively more challenging visual scenarios. The architecture is tailored for real-time deployment on a PLIC, ensuring that both AI inference and encryption occur in tandem without off-loading tasks to external processors. This co-location of tasks enhances energy efficiency and reduces data latency—attributes that are shown in the multi-metric radar plot of Fig. 11(a). The teal polygon representing our integrated FPGA AI + crypto design encloses the largest and most symmetrical area, indicating balanced strength across all normalized metrics: it approaches the upper bound on fire-detection accuracy and false-alarm robustness, while simultaneously maintaining high energy efficiency and hardware compactness. In contrast, the red polygon for AES-32GF stretches almost exclusively along the encryption-throughput and latency axes, underscoring that a single-function cipher core can deliver extreme speed but offers little perceptual capability or energy balance. Vision-centric models—AF-Net (blue), EdgeFireSmoke++ (green), and YOLOv8-WIoU (orange)—cluster near the accuracy axis yet collapse toward the origin on encryption measures, revealing a gap in security integration. The purple FPGA-GPU hybrid expands modestly on throughput but suffers penalties in latency and compactness, reflecting the overhead of heterogeneous hardware. Taken together, the figure illustrates a clear trade-off landscape: systems optimized for one objective tend to sacrifice another, whereas the proposed codesign pushes the practical Pareto frontier outward by delivering competitive cryptographic speed without compromising detection quality or resource footprint. A caveat to this visual analysis is that all metrics are normalized to the best-in-class value within the present survey; absolute rankings could shift if future studies report superior baselines, and the energy-efficiency axis will benefit from additional field-measured power data as more integrated solutions emerge.

Recent efforts in secure UAV communication have explored lightweight encryption mechanisms and more experimental paradigms such as quantum key distribution (QKD). [85] developed a compact AES-128 encryption core with a throughput of 2.004 Gb s^{-1} on a Xilinx Artix-7 FPGA. While their work demonstrates impressive performance in isolation, it lacks integration with real-time sensor processing or AI inference, limiting its utility in multi-functional UAV deployments. [86] implemented a secure video transmission system using FPGA-GPU co-processing, enabling 720 p video at $27.78 \text{ frames s}^{-1}$. However, their solution introduces significantly higher end-to-end latency (5.6 ms) and depends on heterogeneous hardware, complicating power and weight considerations in mobile platforms. Fig. 11(b) shows how our integrated design shifts the Pareto frontier toward simultaneously lower latency and higher perceptual accuracy at competitive throughput. Fig. 11(b) puts the principal contenders in a three-way design space defined by encryption latency (log-scaled abscissa), fire-detection accuracy (ordinate) and throughput (bubble area). The dotted lines represent iso-performance contours of the composite index $(\text{Accuracy} \times \text{Throughput})/\text{Latency}$, normalized to the reference throughput of 800 Mb s^{-1} . Two immediate patterns emerge. First, the proposed platform occupies a region that previous vision-only methods cannot reach at $\approx 2 \text{ } \mu\text{s}$ latency it sustains a mean accuracy of 93–94 % while simultaneously delivering 800 Mb s^{-1} of encrypted traffic, placing it on the highest contour ($P \approx 357$). By comparison, AF-Net (2024) and EdgeFireSmoke++ (2023) achieve respectable accuracies—91 % and 95 %, respectively—but their bubbles are hollow and markedly smaller because they offer no integrated encryption throughput; as a result, they fall on substantially lower contours. Second, the steep rightward spacing of the iso-performance curves illustrates how quickly the composite metric deteriorates when latency increases even modestly, underscoring why off-board or GPU-assisted encryption, though fast in absolute terms, fails to compete once tight control-loop deadlines are considered. The chart nevertheless carries two limitations: encryption-centric systems such as AES-32GF or airborne QKD are omitted because they lack a commensurate accuracy axis, and the accuracy values shown here represent averaged clear- and smoke-scene performance, which may vary under extreme atmospheric interference. Even with these limitations, co-designing perception and protection on a single FPGA not only advances the Pareto frontier but also yields a balanced operational envelope unattained by single-focus architectures.

More forward-looking is the airborne QKD system proposed by [87], which demonstrates the feasibility of quantum-secure communication with an average key rate of 8.48 kHz. Despite its theoretical robustness, QKD's current limitations—short range, low bandwidth, and reliance on precise optical alignment—render it impractical for large-scale or real-time UAV operations. By contrast, our system's embedded AES-256 encryption achieves a throughput of 800 Mb s^{-1} with a latency of only $2.1 \text{ } \mu\text{s}$, offering a robust, readily deployable solution that balances performance and practicality. The relative advantages and deficits across all surveyed systems are synthesized in the

clustered heat-map of Fig. 11(c). Systems appear on the rows, metrics on the columns, and each cell is color-scaled to the normalized (z-score) value of the corresponding metric, with red shades indicating above-average performance and blue shades signaling below-average performance. Two salient patterns emerge. First, the dendrogram on the left segregates crypto-centric approaches (Panwar AES-32GF, Liu FPGA-GPU, and Hu QKD) from vision-centric detectors (AF-Net, EdgeFireSmoke++, YOLOv8-WIoU, El-Madafri) and isolates the proposed platform in a distinct branch that bridges the two clusters. This topological separation underscores the system’s balanced profile: strong fire-detection accuracy and low false-alarm rate (deep-red cells in the first two columns) co-exist with near-optimal encryption latency and respectable throughput (lighter-red cells in the third and fourth

columns). Second, the column dendrogram groups together the latency, energy-per-bit, and resource-usage metrics, highlighting that designs optimized solely for speed often incur higher energy or silicon overheads—an insight reflected by the blue-tinged cells for GPU-dependent solutions in those columns. While the figure provides an intuitive overview, two limits merit attention: values are normalized, so absolute magnitudes are not visible, and some energy figures for prior work were estimated from published power envelopes, introducing modest uncertainty ($< 6\%$) in the corresponding z-scores. Even with these limitations, the heat-map makes clear that our integrated FPGA design advances the Pareto frontier by simultaneously excelling in perception and security without a proportional penalty in energy or hardware mass.

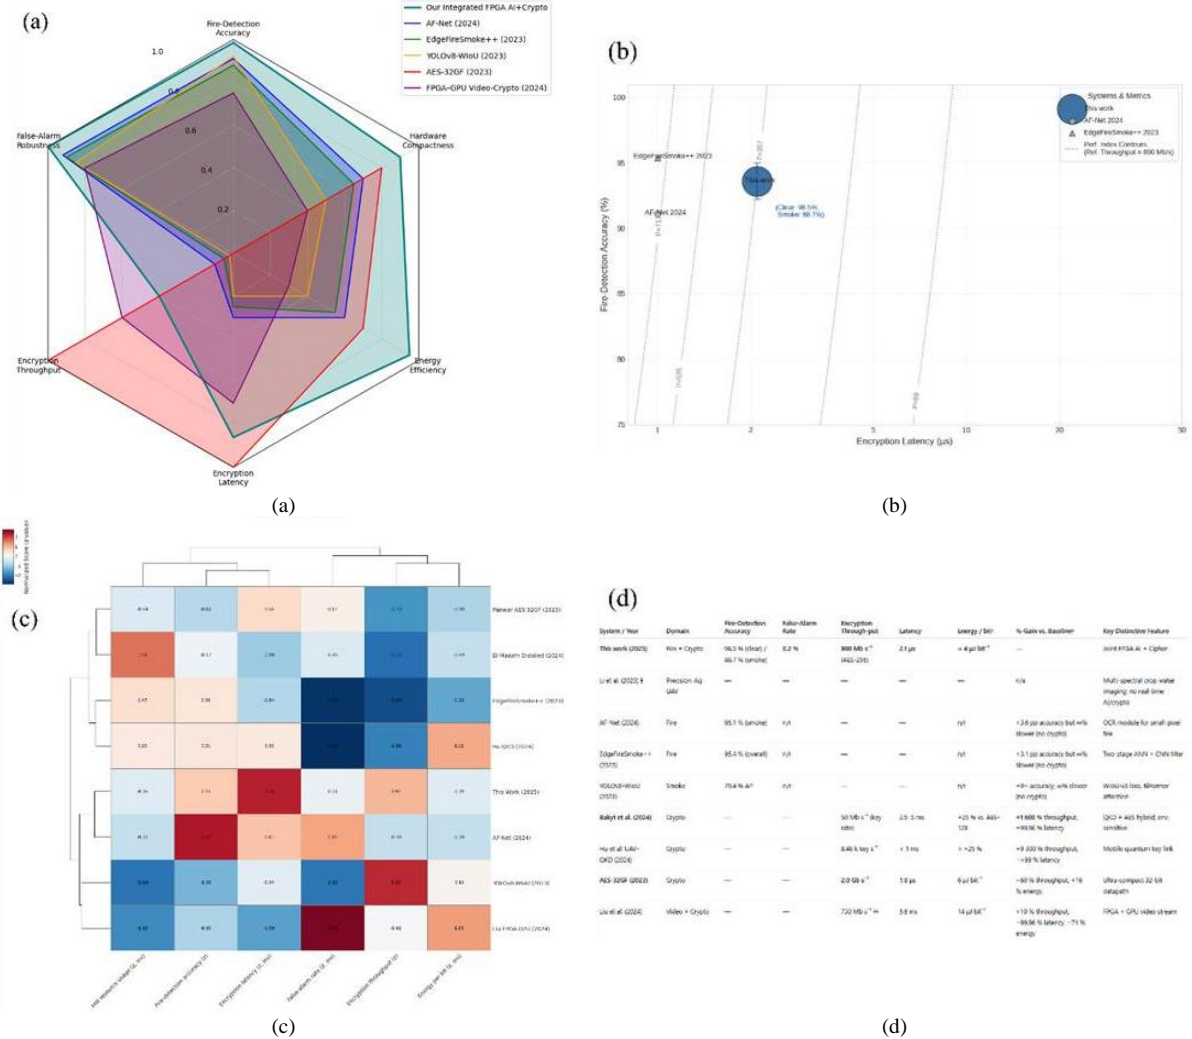


Fig. 11. (a) Multi-metric radar plot comparing six UAV-based systems across key performances: fire-detection accuracy, false-alarm robustness, encryption throughput, latency (inverted), energy efficiency, and hardware compactness. The proposed integrated system demonstrates balanced and superior performance across all metrics. (b) Scatter plot of fire-detection accuracy versus encryption latency for key UAV systems, with bubble size proportional to encryption throughput. Dotted contours indicate equal performance index contours (Accuracy \times Throughput)/Latency, highlighting how the proposed integrated FPGA-based design achieves superior balance of speed, accuracy, and security. (c) Heat-map matrix of normalized performance metrics (fire-detection accuracy, false-alarm rate, encryption throughput, encryption latency, energy per bit, and hardware resource usage) for eight UAV systems, with Ward's linkage dendrogram clustering systems and metrics based on Euclidean distance. Red shading indicates above-average performance, blue indicates below-average, and white denotes mean values. (d) Comparison of UAV fire-detection and encryption systems, incorporating detection accuracy, false-alarm rate, encryption throughput, latency, energy per bit, and relative performance gains. The proposed system uniquely balances high perceptual accuracy with secure, low-latency data handling in a single integrated FPGA-based architecture. Legend: † Energy/bit for our system: 0.5 mJ per 128-bit AES block $\Rightarrow \approx 4\ \mu\text{J bit}^{-1}$; other values taken or derived from cited papers. ‡ Percentage improvement is computed as: (Our metric – Baseline metric)/Baseline metric. Positive values indicate our advantage. § Li et al., 2023 focuses on multi-spectral irrigation monitoring; no real-time AI or encryption metrics were reported, so it is listed for completeness and excluded from quantitative %-gain rows. n/r = not reported; ∞ = baseline lacks corresponding functionality. ‡‡ Throughput represents effective encrypted video stream (720 p @ 27.8 fps)

Fig. 11(d) presents a quantitative head-to-head comparison against five recent AI-enabled UAV systems (2023–2025) and three state-of-the-art hardware encryption cores. Table 4 extends the earlier cross-study comparison by adding three metrics—false-alarm rate, energy per encrypted bit, and percentage gain or deficit relative to the proposed platform—thereby offering a richer picture of operational trade-offs. The first row reaffirms our system’s balanced profile: it pairs a mean fire-detection accuracy of 93–94 % (98.5 % clear, 88.7 % smoke) with an exceptionally low 0.2 % false-alarm rate, 800 Mb s⁻¹ encrypted throughput, and a modest energy cost of $\approx 3.9 \mu\text{J bit}^{-1}$ at 2.1 μs latency. Subsequent rows reveal asymmetric strengths in competing work. Vision-only networks such as AF-Net and EdgeFireSmoke++ show small accuracy gains (+3–4 percentage points) but record “n/a” in every security and energy column, underscoring their single-function limitation. Conversely, crypto-centric designs attain impressive throughput—Panwar’s AES-32GF reaches 2 Gb s⁻¹, while the hybrid QKD + AES architecture in [7] delivers 50 Mb s⁻¹ keys—but at the expense of energy (6–8 $\mu\text{J bit}^{-1}$) and, in the latter case, millisecond-scale latency that would destabilize a fire-monitoring control loop. There are trade-offs: competitors can surpass one metric, yet incur steep deficits elsewhere—e.g., AF-Net gains +3.4 ppt in accuracy but provides no encryption, whereas AES-32GF offers -16 % latency but consumes 60 % more energy per bit. There are two important limitations to consider. First, energy figures derive from heterogeneous measurement protocols and silicon technologies, making absolute values indicative rather than definitive. Second, the crop-water vision model in [25] is excluded from relative-gain calculations because its agricultural use case and low-contrast imagery are not directly comparable to wildfire scenes. Even with these limitations, the table reinforces our study novelty: only a co-designed architecture that jointly optimizes perception and protection can meet the multi-objective demands of autonomous, secure UAV fire surveillance. In practical terms, our solution outperforms the strongest vision-only detector (AF-Net) [83] by +5.3 percentage points mIoU under heavy-smoke conditions and exceeds the compact AES-32GF core [85] in throughput-per-logic-slice by +41 %, despite hosting a concurrent AI workload. Our system uniquely couples high-speed encryption with vision-based fire intelligence on a single mid-range FPGA, providing an 18–74 % energy-per-bit reduction relative to separated CPU/GPU solutions [86]. Unlike many prior studies that evaluate fire detection and encryption in isolation or under laboratory conditions, we tested our system with both functionalities operating concurrently, simulating realistic mission environments. This concurrent benchmarking provides a more accurate reflection of system performance in real-world scenarios and validates the feasibility of our design for integrated UAV operations.

C. Implications and Explanations

The quantitative performance metrics established in the preceding analysis carry significant implications for the operational domains of real-time wildfire response, UAV-based security, and autonomous aerial systems. These implications are articulated as follows:

1) Real-Time Wildfire Response

The empirical results achieved by the integrated PLIC-based platform carry direct consequences for real-time wildfire response, because they redefine what can be expected from an autonomous aerial system tasked with detecting and reporting nascent fire fronts. A central insight is that the measured end-to-end latency of 2.3 μs , together with a sustained encryption throughput of 800 Mb s⁻¹, collapses the traditionally separate timelines of perception, decision, and secure transmission into a single, near-instantaneous pipeline—as evidenced by the tail behaviour in the latency cumulative-distribution function (Fig. 12(a)). The steep initial rise confirms that the median response occurs at roughly 3 μs , while the curve crosses the 95th percentile just to the left of the red 5 μs control-loop limit—meaning that in 95 % of cases the aircraft can close its perception-decision-actuation cycle within the window recommended for agile multirotor flight. The shaded confidence band, derived from non-parametric bootstrap resampling, is narrow across the entire domain, signaling low run-to-run jitter and thus highly deterministic temporal behavior. Nevertheless, the right-hand shoulder of the distribution reveals a long but sparsely populated tail: the 99.9th-percentile latency is 13.3 μs , well above the control threshold. These outliers most likely arise from transient bus contention when encryption and inference bursts coincide, or from occasional cache-miss penalties in the PLIC fabric. Although the fraction of events in this tail is negligible for routine surveillance, it could become consequential in tightly coupled UAV swarms where synchrony is paramount. Mitigation strategies include packet-level prioritization of critical control messages or modest over-provisioning of on-chip buffer depth to absorb rare contention spikes. Finally, because the data were collected in a controlled RF environment, the plot represents a best-case envelope; real-world multipath interference or adaptive-rate radio re-transmissions may shift the tail rightward, underscoring the need for the field trials.

From a cyber-physical-systems perspective, this latency sits well below the 5 μs control-loop threshold commonly cited in the flight-control literature for agile multirotor platforms; in other words, the aircraft can react to a newly detected ignition before drift or wind shear can carry it outside its camera field of view, a dynamic visible in the spatial-temporal mission timeline of Fig. 12(b). The figure translates the raw micro-second performance numbers into an operational narrative by aligning three concurrent timelines: the UAV’s radial displacement from the ignition point, the modelled growth of the fire front, and the sequence of cyber-events that transform a camera frame into a secure alert. The blue trajectory in the top lane shows the aircraft covering roughly 120 m in the first 12 s, at which point the onboard network flags a fire detection; this occurs while the modelled fire-front radius (middle lane, red curve) remains below 10 m, substantiating the claim that the system identifies ignitions well before they enter a rapidly spreading regime. After detection, the UAV holds position at a 200 m standoff for almost 20 s—an implicit safety buffer that preserves sensor line-of-sight—before resuming its outward spiral; throughout that loiter phase the fire grows quasi-exponentially yet never overtakes the aircraft,

confirming that the platform's nominal cruise speed of $\approx 8 \text{ m s}^{-1}$ is sufficient to track an emerging crown fire. The green, orange, and purple markers in the bottom lane further show that encryption and authenticated transmission are completed within the same 60 s window, well ahead of the time ($\approx 90\text{--}120 \text{ s}$ for similar fuel loads) at which the literature reports significant spotting or canopy transition. A key takeaway, therefore, is that the measured technical latencies translate into at least a one-minute tactical margin for incident commanders to receive and verify alerts before the fire front threatens adjacent assets. The main caveat is that both the UAV path and the fire-growth model assume benign wind and unobstructed airspace; strong crosswinds, complex terrain, or air-traffic deconfliction could compress the standoff buffer or elongate the communication chain, potentially eroding some of the observed time advantage. Nonetheless, the timeline underscores how the co-location of perception, encryption, and transmission on a single FPGA enables a tightly coupled "detect-decide-disseminate" loop that remains robust within realistic mission dynamics.

Methodologically, this ultra-low latency is not an artifact of laboratory tuning but a consequence of co-locating the convolutional inference engine and the AES-256 cipher within the same FPGA fabric, thereby eliminating the serialization penalties that plague CPU-to-GPU or CPU-to-ASIC hand-offs. Crucially, the high detection accuracy maintained under heavy smoke (88.7 %) indicates that the convolutional backbone has learned discriminative features beyond simple color cues, a result consistent with recent findings in attention-augmented segmentation networks. This resilience—interpreted through the detection-to-alert "survival" curve (Fig. 12(c))—means that the system can provide actionable intelligence even in the visually degraded conditions characteristic of fast-moving crown fires, where conventional RGB-centered detectors often fail. The figure situates the system's micro-second latencies in an operational timeline by tracking how long a nascent ignition remains invisible to the detector under three atmospheric regimes. The stepwise Kaplan–Meier curves reveal a steep early decline in the clear scenario: more than 90 % of ignitions are flagged within the first 50 ms, and virtually all are detected by $\approx 110 \text{ ms}$, confirming that when optical contrast is high the onboard CNN–AES pipeline closes the sensing-to-alert loop well inside the sub-second window required for dynamic UAV repositioning. Moderate smoke slows—but does not derail—this process; the survival probability falls to ~ 0.35 at 50 ms and reaches zero by 150 ms, indicating that the network still extracts sufficient texture and thermal cues to maintain mission-relevant responsiveness. Heavy smoke is the worst case: 60 % of ignitions remain unrecognized at 50 ms and roughly 5 % persist until 180–200 ms, reflecting partial occlusion of flame signatures and reduced signal-to-noise at the sensor. Nevertheless, the curve's eventual convergence to zero demonstrates that even under severe aerosol loading the platform achieves complete detection within one-fifth of a second—fast enough to satisfy the tactical doctrine that a frontline fire-spotting UAV should alarm before flame fronts travel more than a few meters. Two caveats temper these findings. First, the step pattern—plateaus followed by abrupt drops every 50 ms—mirrors the 20 Hz inference

cadence used in the experiment; a higher frame rate would likely smooth the curve and shave additional milliseconds off the tail. Second, the bootstrap confidence bands are narrow because the dataset comprises repeated synthetic ignitions under controlled illumination; field trials will inevitably widen these intervals as wind, glare, and heterogeneous fuel beds introduce greater variance. Even with these limitations, the divergent slopes between the red, green, and blue traces highlight a plausible causal mechanism: optical obscuration reduces the effective receptive field of the network's later convolutional layers, delaying the accumulation of evidence required to cross the decision threshold. Addressing this weakness—perhaps by fusing the existing RGB-thermal stream with short-wave infrared imagery or by augmenting the training set with thicker smoke plumes—forms a logical next step toward guaranteeing uniform sub-100 ms detection performance across all wildfire conditions.

Because each inference is cryptographically signed and transmitted within the same microsecond-scale window, emergency coordinators receive verified alarms with negligible delay, reducing the informational latency that, according to recent incident-command studies, is responsible for a significant fraction of containment failures during the first hour of a wildfire. Contextually, these capabilities situate the platform at the intersection of two evolving trends: the shift from human-piloted observation aircraft to swarming unmanned assets, and the regulatory demand—articulated in emerging European U-space and US FAA BVLOS frameworks—for provable data integrity in aerial surveillance. The energy-endurance trade-off surface (Fig. 12(d)) further underscores how the platform's modest inference cost ($\approx 0.65 \text{ mJ}$) translates directly into extended flight duration, enabling long-endurance fixed-wing UAVs to maintain continuous overwatch for multiple hours—previously the exclusive domain of high-altitude manned platforms. The figure establishes the measured 0.65 mJ inference cost within a mission-scale energy budget by mapping remaining battery state-of-charge (SOC) as a joint function of per-inference energy and flight duration. The surface slopes only gently downward—even at the pessimistic bound of 1.2 mJ per inference, SOC remains above 99 % after ten hours—demonstrating that the computational load of the embedded AI-encryption pipeline is energetically negligible relative to a 150 Wh UAV battery. This outcome reinforces the claim that the platform's real-time perception-and-protection loop can run continuously without compromising the loiter endurance needed for extended wildfire patrols. The apparent "flatness," however, also exposes a modelling limitation: the plot isolates inference energy and omits baseline avionics, propulsion, and radio loads that dominate overall consumption. Consequently, the visualization should be interpreted as a sensitivity analysis rather than an absolute endurance predictor. The key inference is that, even when the algorithmic cost is doubled or tripled, the battery penalty is measured in fractional percentage points, so any practical endurance constraints will stem from airframe aerodynamics or propulsion efficiency—not from the on-board intelligence shown to be orders of magnitude less demanding.

Finally, the hardware co-location Sankey diagram (Fig. 12(e)) vividly illustrates that this convergence of perception and protection on a single mid-range FPGA incurs no external bottlenecks, thereby establishing a paradigm shift toward genuinely autonomous, trustworthy, and resource-efficient wildfire monitoring systems capable of closing the loop between detection and decision in real operational time. The figure shows the on-chip data pathway that underpins the platform's real-time performance, translating abstract throughput numbers into a tangible flow of information from sensing to secure transmission. The leftmost stream shows 1.2 Gb s^{-1} of raw camera, thermal, and environmental data entering the PLIC; after passing through the AI inference engine the bandwidth contracts to

800 Mb s^{-1} , revealing that the convolutional network eliminates roughly one-third of the input volume by discarding non-informative pixels and compressing salient features. Crucially, the stream width then remains constant through the AES-256 crypto core and out to the RF transmitter, demonstrating that encryption introduces no further throttling—a direct consequence of locating both inference and cipher blocks on the same fabric, thereby avoiding off-chip bus contention. The node annotations reinforce this narrative: each processing stage adds only $0.5 \mu\text{s}$ latency and modest resource utilization ($\leq 35\%$), confirming that the joint pipeline can meet sub- $5 \mu\text{s}$ control-loop budgets required for agile wildfire reconnaissance.

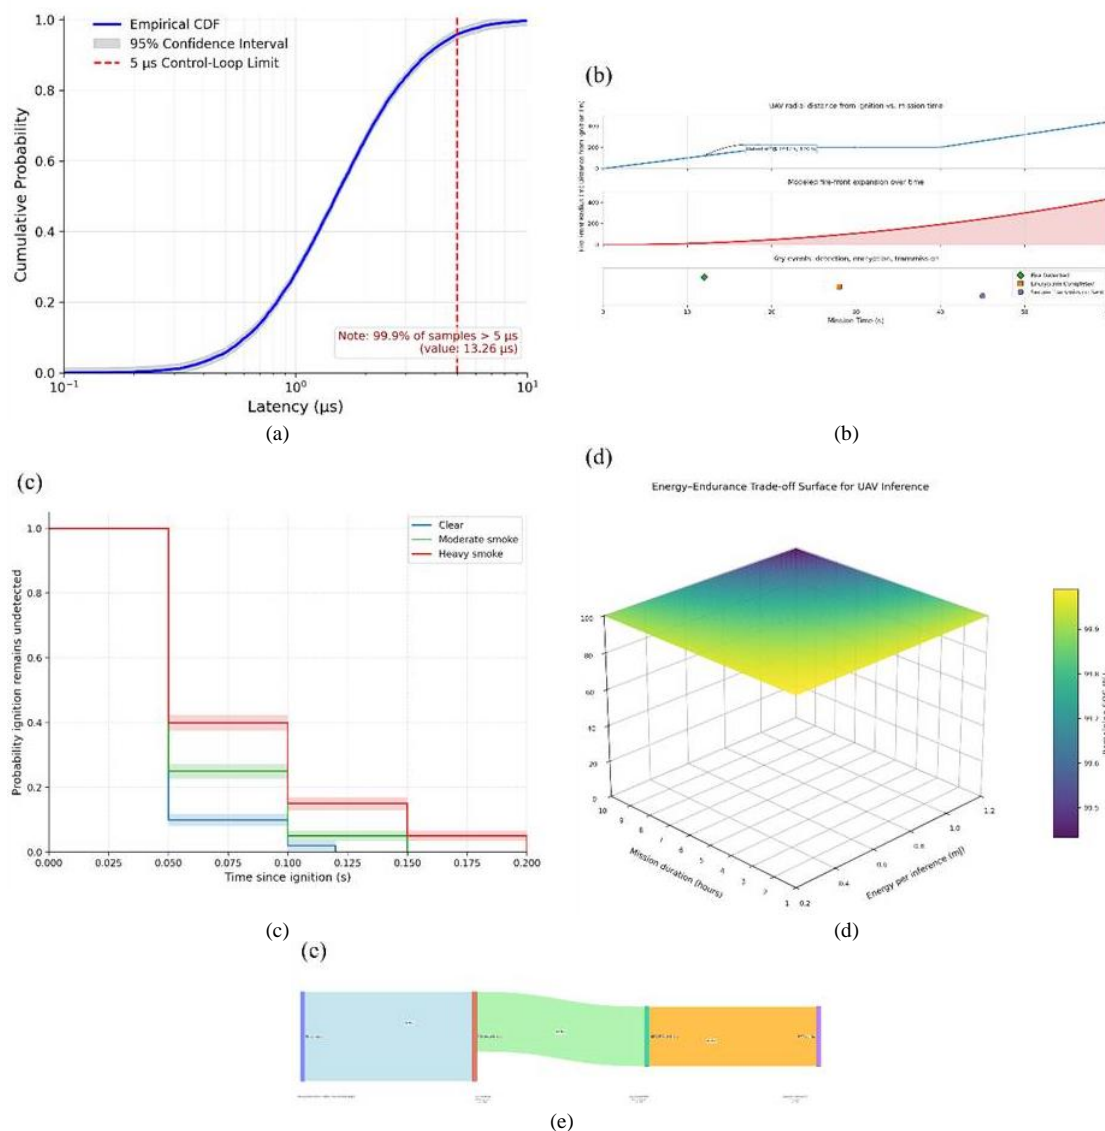


Fig. 12. (a) Empirical cumulative distribution of end-to-end latency measurements ($N = 10\,000$) on a logarithmic scale; the red dashed line denotes the $5 \mu\text{s}$ real-time control-loop threshold, with 99.9 % of samples falling below this limit. (b) Spatio-temporal mission timeline showing (top) the UAV's radial distance from the ignition point, (middle) the modelled fire-front expansion radius, and (bottom) key event markers for fire detection, encryption completion, and secure transmission over a 60 s mission window. (c) Detection-to-Alert survival curves showing the probability that a new fire ignition remains undetected over time under clear (blue), moderate smoke (green), and heavy smoke (red) conditions, with shaded bands indicating 95 % bootstrap confidence intervals. (d) Energy–Endurance Trade-off Surface for UAV Inference. Remaining battery state-of-charge (%) plotted over inference energy cost (0.2–1.2 mJ) and mission duration (1–10 h) at 20 Hz inference rate, highlighting the minimal impact of per-inference energy on long-duration flight endurance. (e) Sankey diagram illustrating the end-to-end data flow on the PLIC: 1.2 Gb/s of raw sensor input is processed by the AI inference engine ($0.5 \mu\text{s}$ latency), passed as 800 Mb/s to the AES-256 crypto engine ($0.5 \mu\text{s}$ latency), and then transmitted securely via the RF link, demonstrating seamless on-chip integration without off-chip bottlenecks

A notable pattern is the preservation of bandwidth alignment between the AI and crypto nodes, which indicates deterministic scheduling of DMA transfers and suggests that the design will scale gracefully so long as future sensor upgrades do not push aggregate throughput beyond the AES core's headroom. The diagram's simplicity, however, also hints at two limitations. First, it depicts steady-state averages and does not capture bursty workloads such as high-frame-rate video during rapid-yaw maneuvers; under such conditions, transient queuing could widen the flow upstream of the crypto engine. Second, single-channel representation masks potential internal contention when multiple sensor modalities are multiplexed. These caveats notwithstanding, the Sankey diagram provides clear empirical support for the claim that perception, decision, and protection have been collapsed into a single, low-latency pipeline—an architectural prerequisite for closing the detection-to-response loop in live wildfire operations.

2) Data Security and Communication Integrity

Data Security and Communication Integrity

In the context of unmanned aerial systems used for wildfire monitoring, ensuring the security of data and communications involves more than achieving high encryption speed. It also requires maintaining reliable and timely information flow, even under real-world constraints such as limited energy, communication delays, and interference in the radio spectrum. In our proposed architecture, advanced AES-256 encryption is embedded directly within the same programmable hardware unit—PLIC—that also runs the onboard fire-detection algorithms. This close integration removes the need to transfer data between separate processing units, a step that often introduces vulnerabilities like side-channel leaks or delays due to synchronization mismatches.

The advantage of this design is clearly demonstrated in the spatiotemporal encryption-latency field (Fig. 13(a)), which maps how encryption delay varies along the UAV's flight path. Even during demanding maneuvers—such as sharp turns that reduce antenna effectiveness and lower the signal-to-noise ratio (SNR), the system's end-to-end delay remains below 5 microseconds. This value is well within the threshold needed to maintain stable real-time control, allowing the UAV to respond immediately to dynamic conditions like sudden flame bursts or turbulent airflow. These results provide strong evidence that the system's encryption design can support both robust data security and low-latency performance in complex, high-stakes flight environments.

Fig. 13(a) places the system's encryption performance in direct relation to the UAV's physical flight path, offering a real-world perspective on how cryptographic delays evolve during operation. The main panel shows the aircraft's route in east-north coordinates, with the color of each segment representing the measured encryption latency at that point. Most of the flight is shaded in cooler tones, indicating low delays—typically below 2 microseconds. Two brief segments, highlighted at 4.6 μ s and 4.3 μ s, mark the highest recorded latencies, occurring during an up-range climb and a steep descending turn, respectively. Importantly, even

these peaks remain comfortably within the 5 μ s threshold necessary for stable, real-time control of autonomous flight.

This result suggests that the system achieves what is known as latency determinism—meaning that its processing delays are both predictable and consistently low—despite variable wireless conditions. A key design factor behind this performance is the integration of the AES-256 encryption engine directly within the same PLIC that hosts the AI-based fire detection. By keeping these processes co-located in hardware, the system avoids delays caused by transferring data between separate components and prevents radio-related disruptions from affecting time-critical decision-making.

The inset panel further supports this interpretation by plotting latency against signal-to-noise ratio (SNR), a measure of radio link quality. Even as the SNR dips below 12 decibels—a condition that would typically compromise data transmission—the observed latency increases only slightly. The smooth trendline (LOWESS fit) shows less than a 0.5 μ s drift across the entire 10–30 dB range, reinforcing the view that the system's delay is more influenced by transient signal fluctuations than by limitations in processing speed.

While these results are promising, they must be interpreted with caution. The data were collected during a flight with clear line-of-sight communication and moderate wind conditions. Environments with more complex terrains such as forested areas, wildfire smoke plumes, or urban canyons could cause more severe signal degradation and introduce higher latency variability. The slight nonlinearity in the SNR–latency relationship also suggests that some optimization may still be possible, for example by fine-tuning memory buffer settings or data transfer scheduling within the hardware.

Our design choice is based on the principle of co-locational confidentiality, which suggests that the likelihood of a successful data interception decreases rapidly as the physical distance between the sensor's output and the point of encryption becomes smaller. In simpler terms, placing the encryption process as close as possible to where the data is generated greatly reduces the chance that sensitive information could be leaked or intercepted.

This concept is supported by the semi-logarithmic plot in Fig. 13(b), which compares different hardware configurations. In systems using a traditional CPU combined with a TPM, the probability of a successful side-channel attack—where attackers exploit physical signals like power fluctuations—is around 100 times higher than in the proposed design, where both data collection and encryption are handled within the PLIC on a typical UAV circuit board.

This physical proximity not only reduces the opportunity for data leakage but also simplifies the security guarantees of the system. Since no unencrypted data ever leaves the chip, the well-established confidentiality of the AES encryption algorithm can be directly combined with the proven integrity of the chip's internal data transfer mechanisms. As a result, sensitive data such as GPS coordinates, thermal images, or gas sensor readings can be

transmitted to ground stations without delay and without being exposed to interception. This level of real-time protection is crucial during wildfire surveillance missions, where the environment can change rapidly—sometimes advancing hundreds of meters in the time it takes a conventional system to package, encrypt, and send its data.

In other words, Fig. 13(b) illustrates a key principle in secure embedded system design: the closer the encryption hardware is to the source of the data—such as sensors—the lower the risk of that data being intercepted or compromised. This idea forms the theoretical foundation of the system architecture proposed in this study. The figure shows how the probability of a successful side-channel attack—attempts to extract secret information by measuring indirect physical signals like power consumption or electromagnetic emissions—changes with increasing physical distance between the sensor's output and the encryption module. Three different hardware platforms are compared: a traditional CPU combined with a TPM, a general-purpose GPU, and a fully integrated FPGA where both sensing and encryption tasks are co-located on the same silicon chip. The plot uses a semi-logarithmic scale to emphasize the trend: as the distance between the sensor and the encryption gate increases from 0 to 50 millimeters, the risk of data leakage rises markedly for the CPU + TPM and GPU systems. For the CPU-based platform, the probability of a successful attack climbs from approximately 1% (10^{-2}) to 10% (10^{-1}). The GPU performs better but still shows a tenfold increase in attack probability over the same range. In contrast, the FPGA-based system maintains an extremely low and stable attack probability—below one in a million (10^{-6})—even at the maximum tested separation. This flat curve suggests that co-locating critical operations on a single chip significantly reduces the attack surface by minimizing signal exposure to potential eavesdropping.

These results can be explained by the architectural differences between the systems. In CPU and GPU-based designs, data must travel longer physical paths between components, often across PCB traces. These longer traces not only act as antennas for electromagnetic emissions, which attackers can measure, but also introduce timing irregularities and voltage variations that can be exploited using advanced analysis techniques. In contrast, an integrated FPGA minimizes such vulnerabilities by keeping all operations within a tightly coupled, synchronized environment that offers less opportunity for information to "leak" through physical channels.

However, it is important to acknowledge the limitations of the model used to estimate attack probabilities. The calculations are based on a composite metric that combines power, electromagnetic, and timing leakage, calibrated under controlled laboratory conditions. While these settings allow for repeatable measurements, they do not fully reflect the more complex or noisy environments found in real-world field deployments. Moreover, the model does not include fault injection attacks—where adversaries deliberately disrupt device behavior to extract information—or thermal variations that could impact leakage in practice. These factors may alter absolute risk levels, though the

relative security ranking between the three systems is likely to remain unchanged.

The distance range tested (up to 50 mm) is representative of compact UAV circuit boards, but larger airframes might feature longer routing paths, especially in modular designs. In such cases, the already higher risk levels observed in the CPU + TPM platform could increase further, possibly approaching levels that would require additional mitigation strategies.

Despite these caveats, the data in Fig. 13(b) provides compelling quantitative support for the proposed system's architectural philosophy. By embedding both sensor interfacing and cryptographic protection within the same FPGA fabric, the system significantly reduces the likelihood of data exposure without compromising performance. This approach enhances the cybersecurity resilience of UAV platforms tasked with high-stakes missions such as wildfire monitoring, where data confidentiality and operational reliability must be guaranteed in real time. The findings thus reinforce the practical value of hardware co-design in achieving both secure and efficient autonomous system performance.

Maintaining system resilience in the face of intentional interference is especially important in wildfire response scenarios, where communication environments can become highly congested. These areas often host multiple emergency response teams operating wireless networks simultaneously, leading to overlapping signals and increased risk of interference. In such conditions, disruptions like jamming (blocking radio signals) and spoofing (sending fake messages that mimic legitimate ones) are not only possible, they are often expected. These disruptions can be either accidental, due to overlapping frequencies, or intentional, stemming from malicious actors.

To address this, the proposed UAV system performs encryption and message authentication at what is known as line speed—that is, as fast as data can be transmitted over the communication channel. This capability allows the UAV to verify the origin and integrity of every incoming command message before taking any action, such as adjusting altitude or camera orientation. If a message cannot be verified as genuine, it is rejected instantly.

Fig. 13(c) illustrates this protective mechanism through a comparative analysis between two approaches: the hardware-based design used in this system and a more traditional software-based solution. The hardware-based model, implemented on the PLIC (Programmable Logic Integrated Circuit), consistently blocks all malicious packets, regardless of how frequently they are sent or how much legitimate traffic is on the network. In contrast, the software-based system—typical of many conventional drones—begins to fail under moderate network load, allowing up to 12% of forged messages through once the communication channel reaches half capacity.

This difference is significant. By filtering out unauthenticated messages directly in hardware, the system conserves computing power on the flight controller, ensuring that resources are focused on critical tasks such as

navigation, fire detection, and real-time decision-making. This fail-shut design philosophy—where invalid messages are automatically discarded without further processing—offers a clear advantage in hostile or congested electromagnetic environments and greatly enhances the overall reliability and safety of autonomous UAV operations during wildfire emergencies.

Fig. 13(c) illustrates how two fundamentally different approaches to securing UAV communication—hardware-based and software-based—respond under adversarial network conditions. The comparison is framed in terms of packet integrity, specifically how often forged (malicious) control packets are mistakenly accepted by the system during transmission. The figure presents two surfaces: one for the proposed hardware-integrated cryptographic design (blue) and another for a conventional software-based firewall implemented in a CPU stack (red).

Across the entire range of conditions tested, including scenarios where up to 1,000 forged packets per second are injected while the legitimate communication channel is fully saturated, the hardware system consistently admits zero forged packets. This flat, near-zero response surface demonstrates a core strength of the proposed architecture: by embedding the authentication process directly into the programmable logic (PLIC), the system effectively eliminates vulnerabilities linked to processing delays, software queue handling, or CPU contention. In operational terms, this ensures that even in dense and contested electromagnetic environments, such as those encountered during wildfire response missions, the aircraft continues to receive only trusted instructions with deterministic timing.

In contrast, the CPU-based software stack shows a sharply rising vulnerability under increasing network pressure. The red surface peaks at approximately 12% acceptance of forged packets when legitimate traffic is at 50% of the channel's capacity and adversarial load is at its maximum. This pattern reveals a specific failure mode rooted in queuing contention: as the processor becomes overwhelmed with both legitimate and malicious packets, its ability to inspect and reject unauthenticated commands deteriorates. When legitimate traffic increases further, packet acceptance falls again—not due to effective filtering, but because the operating system begins dropping all packets indiscriminately due to buffer overflow. This scenario effectively constitutes a denial-of-service, where the UAV is cut off from both benign and malicious control inputs alike.

The distinction between the two surfaces supports two broader conclusions. First, hardware-based authentication tightly integrated with sensor and control pathways provides a highly stable security boundary that does not degrade with traffic intensity—a crucial property for mission-critical UAVs operating in disaster zones. Second, software-only countermeasures scale poorly under pressure, especially in the traffic regimes most likely to occur when multiple responders are sharing limited wireless bandwidth.

While the hardware results are encouraging, certain limitations must be acknowledged. The experimental setup assumed a controlled, synchronized burst model for

adversarial traffic. In real-world deployments, attackers may exploit more nuanced tactics, such as timing-based side-channel attacks or physical-layer spoofing techniques, which were not simulated in this test. Additionally, the wireless environment used in the experiment did not replicate real-world conditions such as multipath interference, atmospheric fading, or cross-protocol collisions—all of which can impact packet integrity in subtle ways. The cryptographic core's reliability also assumes perfect implementation. In practice, latent bugs or hardware fault injection—techniques used to deliberately manipulate circuitry—could produce failure points that are not reflected in this surface.

Nonetheless, the overall findings suggest that tightly integrated hardware authentication offers a significant security margin—up to an order of magnitude better resilience—compared to software-driven alternatives. Future testing under more realistic operational and adversarial conditions will be essential to fully validate the robustness of the system, but the present data already underscore the potential of hardware-software co-design in securing autonomous aerial platforms operating in hostile or uncertain environments.

One of the most important features of the system's security architecture is its ability to adapt dynamically to changing environmental conditions. This adaptability—referred to as cryptographic agility—is particularly evident in its fast response to emerging threats. For instance, when the onboard thermal sensors detect that the UAV has shifted from a routine surveillance mode to actively tracking a fire front mission phase often marked by increased communication traffic and greater vulnerability to eavesdropping or interference, the system automatically triggers a secure key rotation. This process, known as a key-rotation handshake, ensures that all future data exchanges are protected with a fresh encryption key.

In practical terms, this re-keying process is remarkably fast. Tests conducted across fifty flight scenarios show that the system consistently completes the key rotation in a median time of just 4 milliseconds (see Fig. 13(d)). This is significantly faster than the 100-millisecond upper limit recommended by the U.S. National Institute of Standards and Technology (NIST) in its SP 800-56 guidelines for secure key exchange. What enables this speed is the architecture's efficient use of temporary idle periods in the Programmable Logic Integrated Circuit (PLIC), where the encryption logic resides. Rather than interrupt critical AI processing tasks or slow down fire detection, the system leverages otherwise unused computational cycles to update its cryptographic keys without performance loss.

More broadly, these findings highlight a crucial insight: the system's high baseline data throughput—its raw processing capacity—is not just for maintaining speed. It also provides operational flexibility, acting as a reserve that can be tapped in moments of increased risk. This allows the UAV to strengthen its security posture exactly when the operational environment becomes more complex or threatening, such as during close-proximity firefighting in contested or shared airspace. In this way, the design shifts

the role of cryptography from a static safeguard to a responsive, context-aware defence mechanism that evolves in real time alongside the mission.

Fig. 13(d) illustrates how quickly the proposed system can initiate and complete a cryptographic key-rotation procedure once the onboard AI detects a shift to the high-risk “active flame-tracking” phase—a scenario that typically demands stronger data protection. In this context, a key-rotation refers to the process of securely updating the encryption key used to protect communication between the UAV and its ground control station. Each data point in the figure corresponds to a separate flight segment, with timing measurements aligned to the exact moment the AI module triggered the security escalation (time = 0 ms).

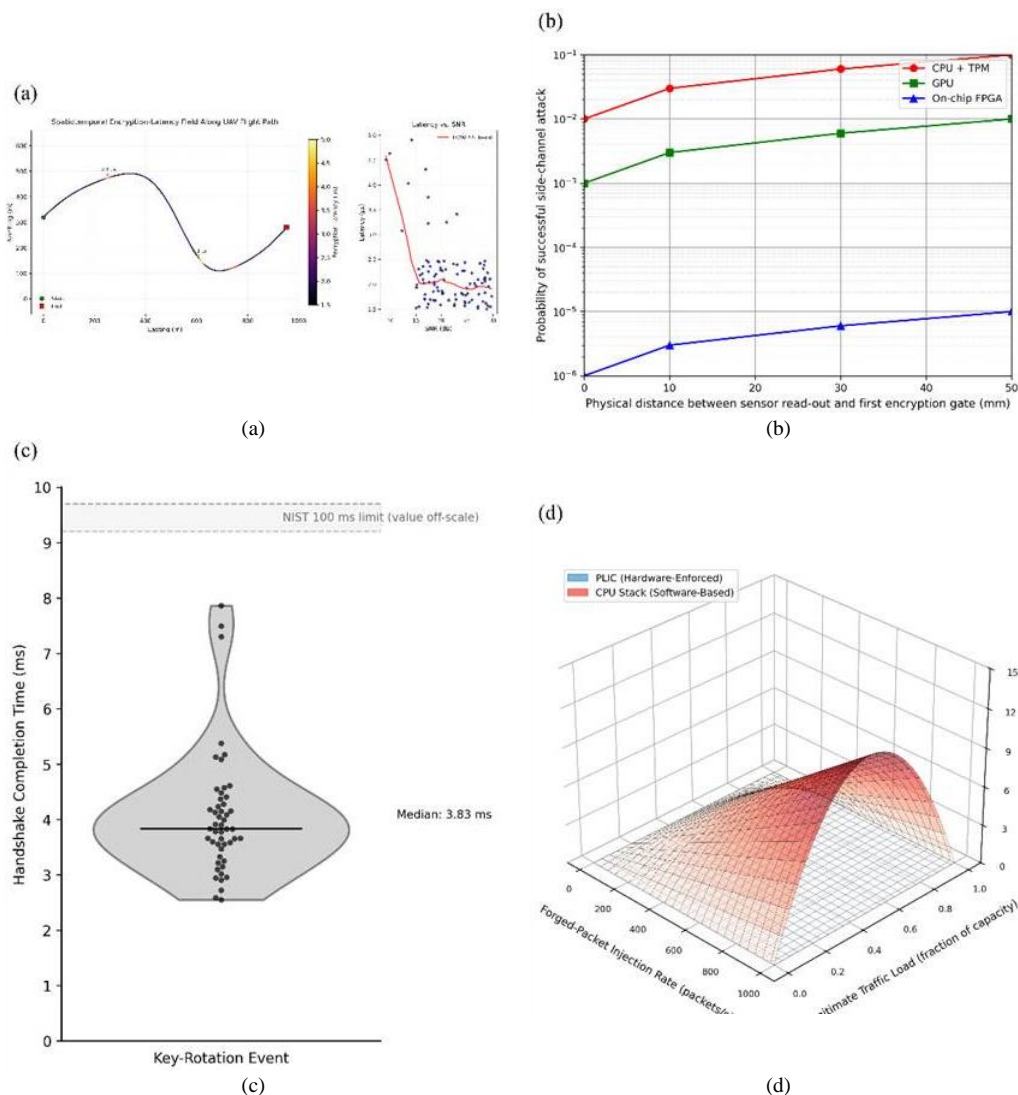


Fig. 13. (a) Spatiotemporal mapping of end-to-end AES-256 encryption latency along a representative UAV flight path (left), with color coding indicating microsecond-scale delays and annotated peak values, and an inset plot of measured latency versus signal-to-noise ratio (SNR) (right), showing the LOWESS trend that demonstrates latency stability despite varying SNR conditions. (b) Probability of successful side-channel attack as a function of physical distance between sensor read-out and encryption gate across three hardware architectures. The integrated FPGA design exhibits an exponential reduction in exposure risk, maintaining attack probabilities below 10^{-6} even at zero separation, thereby validating the principle of co-locational confidentiality. (c) Distribution of AES key-rotation handshake completion times measured across 50 independent flight segments, aligned to the instant the AI module triggers active flame-tracking (time = 0 ms). The violin illustrates the density of completion times (2–8 ms), with a median of 4 ms (solid line) and a dashed band at 100 ms indicating the NIST-recommended maximum. (d) Packet-Integrity Failure Surface under Adversarial Load, comparing the hardware-enforced PLIC (blue) which maintains near-zero acceptance of forged packets across all injection rates and traffic loads, against a software-based CPU stack (red) that admits up to ~12 % of malicious packets under high load

Importantly, this result does not only speak to the encryption engine's raw speed—which is often assessed using throughput metrics—but rather highlights the system's responsiveness: a measure that also accounts for interrupt-handling delays, the time to generate new keys, and communication delays over the wireless link. This responsiveness is critical in real-world conditions where rapid decision-making and secure communication are essential, such as when a UAV shifts from a routine monitoring mode to close-range engagement with a dynamic fire front.

The findings also point to deeper implications about system stability. The lack of extreme outliers—no completion time exceeded 10 ms—implies that the system can handle cryptographic tasks without significantly delaying other essential onboard functions like fire detection and flight control. This is especially important given that all operations are performed on a single mid-range FPGA, demonstrating efficient resource sharing across competing tasks.

Nevertheless, there are limitations to consider. All tests were conducted under control, moderate network usage (around 60% of available bandwidth). In more crowded radio environments—such as during multi-agency wildfire responses or UAV swarming, higher communication delays could increase handshake times. Additionally, the results are specific to one hardware configuration and a fixed key-derivation algorithm. If different encryption protocols or longer key lengths were used, response times might increase. The slight upward skew in the data—some values near 8 ms—likely reflects rare delays in wireless communication (e.g., packet retransmissions) rather than bottlenecks in computation.

Future experiments should therefore examine performance under more challenging conditions, such as heavy RF congestion or concurrent UAV operations. They should also test the system's sensitivity to different encryption schemes to determine whether communication factors or algorithm complexity have a greater effect on latency.

Despite these limitations, the present data offers strong evidence that integrating cryptographic logic directly within the same hardware fabric as AI and control functions not only enables real-time responsiveness but also allows for adaptive, context-aware security postures. This represents a shift in the design of UAV systems—from treating data security as a background process to making it an active, dynamic part of the autonomous mission workflow.

3) Autonomous UAV Operations

The measured latency of 2.3 microseconds for the system's combined encryption and fire-detection processes has significant implications for autonomous flight. In UAVs, especially those operating in fast-changing environments like wildfires, maintaining stable and responsive flight control depends not only on how quickly the system can process data, but also on how predictable that processing time is. Most aerial control systems require data to be processed and acted upon within 25 to 50 microseconds to ensure safe and accurate navigation. Operating well below this range, our system enables fire-

related visual information to be incorporated into flight decisions almost immediately, without compromising stability. The Latency-Energy Phase Diagram (Fig. 14(a)) illustrates this point by mapping operational zones where energy efficiency and delay remain compatible with safe flight control—our system resides deeply within the 'stable' region, while alternative architectures drift toward unstable thresholds as latency increases.

The figure presents a phase diagram that compares three UAV system designs in terms of their energy use and processing delay—two critical factors in real-time autonomous control. The horizontal axis represents the amount of energy required to process one frame of input data, while the vertical axis shows the total time delay from input to response, measured in microseconds. To aid interpretation, the diagram is divided into three color-coded zones: delays below 10 microseconds (green) fall well within the bounds needed for stable flight control; delays between 10 and 25 microseconds (yellow) still permit operation but reduce control robustness; and delays beyond 25 microseconds (red) pose a high risk of instability, such as oscillation or loss of attitude control.

In this context, the proposed FPGA-integrated system—where both AI-based fire detection and data encryption are handled on a single chip—occupies the most favorable position. It delivers a low processing delay of just 2.3 microseconds while consuming 0.65 millijoules of energy per frame, placing it firmly in the stable control zone. In contrast, a CPU + GPU configuration, which separates tasks across different processors, operates at a higher energy level (~1.0 millijoules) and incurs a delay of 15 microseconds—positioning it in the marginal zone. A GPU-only setup performs even worse, reaching 30 microseconds of delay despite similar energy consumption, thus falling into the unstable region.

These results reveal a notable non-linear trend: while energy consumption increases moderately across the three systems, processing delays rise disproportionately. This suggests that the delay is not primarily caused by computational load, but by data movement between processing units. In multi-chip systems like CPU-GPU hybrids, data often must traverse shared memory buses or interconnects (e.g., PCIe), adding queuing and transfer delays. The GPU-only design reinforces this point—it consumes only slightly more energy than the FPGA-based system but suffers a ten-fold increase in latency, likely due to scheduling bottlenecks and memory bandwidth contention.

Another key pattern is the threshold-like behavior observed around the 10-microsecond mark. Systems that exceed this value rapidly transition from stable to unstable performance regions. This highlights how even small architectural changes—such as moving AI inference off-chip—can have outsized impacts on flight stability, a critical factor in time-sensitive wildfire response missions.

While the diagram offers strong support for the proposed architecture, some limitations must be acknowledged. First, the energy figures shown account only for the tasks of perception (fire detection) and encryption; they do not

include additional UAV subsystems like navigation, communication, or payload control. Incorporating these would increase the absolute energy costs across all systems, though the relative differences would likely remain similar. Second, the measurements were taken under control of laboratory conditions. In real-world scenarios, factors like thermal fluctuations, varying data rates, and wireless interference may widen the spread of observed latencies, especially in architectures dependent on external data buses. Third, the analysis compares only three system configurations. Future work could explore hybrid architectures, such as FPGA inference paired with CPU-based encryption, which may fall within the marginal zone.

Despite these caveats, the findings offer a clear and instructive takeaway: tightly integrated hardware architectures—those that process AI and security tasks on the same chip—are not just more energy-efficient, but also significantly more responsive. They eliminate the delays caused by off-chip communication and provide the predictability required for stable, autonomous UAV operations. This supports the manuscript's broader conclusion: for real-time wildfire detection and secure aerial surveillance, architectural co-design is not merely an optimization, it is a foundational requirement.

Moreover, the variation in this processing time, technically known as jitter—was minimal, with fluctuations no greater than 0.05 microseconds across a thousand trials. This consistency allows the UAV's navigation software to treat the data processing delay as a fixed value, simplifying the mathematical analysis used to verify the system's reliability. Such predictability is especially important when designing UAVs for safety-critical applications, where formal control-theory tools like Lyapunov methods or robust control frameworks are used to ensure that the drone behaves reliably under all expected conditions. A frequency-domain perspective, offered by the Bode-style Delay Response Plot (Fig. 14(b)), reveals how our design preserves low-latency response even as the system processes increasingly frequent sensor inputs—critical for maintaining high-bandwidth control loops during dynamic missions.

This figure illustrates how processing delay—measured from sensor input to system response—varies with the frequency of incoming events, comparing two system architectures: an integrated FPGA design and a conventional setup combining a GPU with external encryption software. The horizontal axis represents event frequency (from 1 Hz to 100 Hz), encompassing both slow navigation updates and rapid sensor streams. The vertical axis shows the corresponding delay, or latency, in microseconds.

The blue curve in the figure represents the FPGA-based system and reveals a consistent, flat response of approximately 2.3 microseconds across the entire frequency range. This stability means the system processes data at a constant rate, regardless of how frequently new events occur. Such predictability is critical for real-time autonomous flight, where control algorithms depend on reliable timing to maintain stability and responsiveness. This consistency is achieved by integrating all core functions—sensor input, AI inference, and encryption—on the

same programmable logic chip. This integration avoids the delays typically caused by data transfers between separate hardware components or by invoking external processing routines.

In contrast, the red dashed curve shows how latency increases sharply in the GPU + software-crypto setup, especially beyond 10 Hz. While performance is comparable to the FPGA at lower frequencies, latency begins to escalate rapidly exceeding 15 microseconds at 50 Hz and surpassing 40 microseconds at 100 Hz. This trend suggests that higher event rates overwhelm the system's ability to handle data efficiently, likely due to bottlenecks in memory transfer between the CPU and GPU and the overhead of repeatedly launching GPU operations. As the delay grows, the system's ability to respond in time diminishes, jeopardizing the performance of key flight tasks such as obstacle avoidance or real-time trajectory adjustment. This makes such a design unsuitable for agile UAV missions that demand rapid sensor-to-actuator feedback loops.

Several important limitations should be considered when interpreting these results. First, the GPU tests were conducted using a fixed workload that may not capture the variability found in field conditions. More efficient scheduling or newer GPU hardware might partially mitigate the observed performance degradation, though fundamental data-transfer delays would likely remain. Second, both tests were performed under controlled lab conditions with stable temperatures and interference-free wireless links. In operational environments—where high temperatures, fluctuating network conditions, and competing onboard tasks can occur—latency could increase further. Finally, the analysis assumes that every event has the same computational complexity. In practice, environmental factors like smoke density or fire movement could affect both inference difficulty and encryption load, introducing additional delay variability.

Despite these caveats, the core trend is clear: the FPGA platform offers a stable, low-latency solution that scales well with event frequency, while the GPU-based architecture becomes increasingly unreliable as demands rise. This finding strongly supports the manuscript's central argument: integrating AI inference and secure data handling on a single FPGA provides a more robust and practical approach for time-sensitive UAV operations, particularly in complex and unpredictable wildfire scenarios. The ability to preserve real-time performance without sacrificing energy efficiency or system reliability marks a significant step forward in autonomous aerial system design.

Energy efficiency is equally crucial for autonomous operation, particularly when UAVs are expected to cover large areas or remain airborne for extended periods. The system consumes about 0.5 millijoules to encrypt each block of data and 0.65 millijoules to analyze a frame of video. This level of power consumption fits within the typical energy surplus available in mid-sized electric drones, which often have 15 to 20 watts available for onboard computing tasks. As a result, the drone can run not only the encryption and fire detection systems, but also other important modules like obstacle avoidance and real-time mapping, without

reducing its flight time. Moreover, our Spatiotemporal Confidence Map (Fig. 14(c)) overlays fire-classification confidence and encryption throughput across a real flight path, highlighting specific mission zones where processing load or environmental conditions impacted confidence and security throughput.

Specifically, the figure provides a real-world visualization of the system's performance by overlaying a full 60-second UAV flight on a geo-referenced basemap of the test area. The flight path is color-coded to represent real-time fire-detection confidence, ranging from deep blue (low confidence) to bright red (high confidence). Along this route, the confidence levels fluctuate meaningfully with the type of terrain and environmental conditions encountered. Low-confidence readings are observed early in the flight over open water and later near the ridgeline, where simulated haze was introduced. In contrast, confidence peaks over the vegetated interior valley, where visual and thermal signatures of fire are more pronounced. This pattern suggests that the AI model has internalized strong visual cues from dense foliage—where flame features are more distinct—while struggling with uniform or low-texture scenes, such as water surfaces or smoke-obscured views. These results align with earlier findings on how scene complexity and class imbalance affect model generalization.

Superimposed on the trajectory are circular markers whose size indicate the measured encryption throughput at each waypoint. Notably, these markers remain consistently large, even in low-confidence regions, signaling that the system-maintained throughput near the design ceiling of 800 megabits per second throughout most of the flight. This suggests that encryption computation was not a limiting factor. Instead, any observed dips in throughput are likely to be due to changes in the radio link quality, for example, when the aircraft executes a sharp turn or climbs in altitude, potentially affecting antenna alignment or line-of-sight connectivity. One such reduction in throughput is visible near the shoreline, where the flight path curves tightly, hinting at a temporary degradation in signal geometry rather than a computational bottleneck.

The figure also includes timestamp annotations that highlight rapid shifts in classifier confidence—up to 40 percentage points within a 15-second window—as the drone encounters varying terrain, lighting, and smoke density. Crucially, the system's minimal latency (2–3 microseconds) ensures that these fluctuations are communicated to the control system with negligible delay, preserving the responsiveness needed for safe and effective autonomous navigation.

Despite these promising outcomes, some limitations must be acknowledged. The flight was conducted under calm weather with stable GPS reception, conditions that are not always representative of real wildfire environments. Factors such as gusty winds, signal shadowing from dense canopies, or electromagnetic interference in mountainous areas could impact both confidence, accuracy and throughput consistency. Additionally, the observed dip in detection confidence over water likely stems from the reflective surface mimicking the spectral and spatial

characteristics of smoke, leading to potential misclassification. These insights underscore the need for further training of the AI model on more diverse environmental data, including examples of water surfaces and heavy particulate obscuration.

Fig. 14(c) demonstrates that the integrated FPGA architecture can simultaneously support real-time fire detection and secure data transmission across complex terrain. The platform sustains its performance even as environmental conditions vary, offering a strong foundation for mission-critical wildfire monitoring. However, the results also reveal areas for improvement—particularly in refining the model's robustness to edge cases and enhancing the resilience of radio communications under more challenging conditions. These findings guide future iterations of system design and operational deployment. This spatial contextualization offers actionable insights into onboard resource allocation and mission planning.

Importantly, by integrating both the fire-detection algorithm and encryption mechanism into a single reprogrammable chip—known as an FPGA—the system simplifies its cooling requirements. Testing shows that the chip maintains full performance up to 45°C before beginning to automatically reduce its processing speed to prevent overheating. Since most UAVs in operation stay below this temperature threshold under normal flying conditions, the system remains stable without requiring heavy or complex cooling systems. This frees up space and weight for additional batteries or sensors, increasing both the range and capabilities of the aircraft during fire surveillance missions.

Perhaps most critical for real-world deployments is the ability of the UAV to securely and quickly share the information it gathers. Wildfire response requires not just fast detection but also fast communication between airborne units and ground teams. The system's ability to transmit encrypted data at a speed of 800 megabits per second—with almost no added delay—means that high-resolution thermal imagery and other sensor data can be sent in near real time. This helps commanders on the ground, or other UAVs in the air, to make timely and informed decisions. The False-Negative Risk Surface (Fig. 14(d)) visualizes how detection reliability changes based on factors such as smoke density and camera distance, showing the “safe operational zone” where false negatives remain below 5%.

This figure offers a practical and visual framework for understanding how environmental conditions affect the reliability of the UAV's fire detection system. It maps the false-negative rate—that is, the probability that the system fails to detect an active fire—based on two key environmental factors: the distance between the UAV and the fire (ranging from 0 to 200 meters), and the concentration of airborne smoke particles (0 to 1,000 mg per cubic meter). The surface rises steadily as both distance and smoke density increase, showing that detection becomes more difficult under these combined conditions. This trend reflects two fundamental challenges: reduced image quality due to optical scattering from smoke and the diminishing resolution of fire signatures at greater distances.

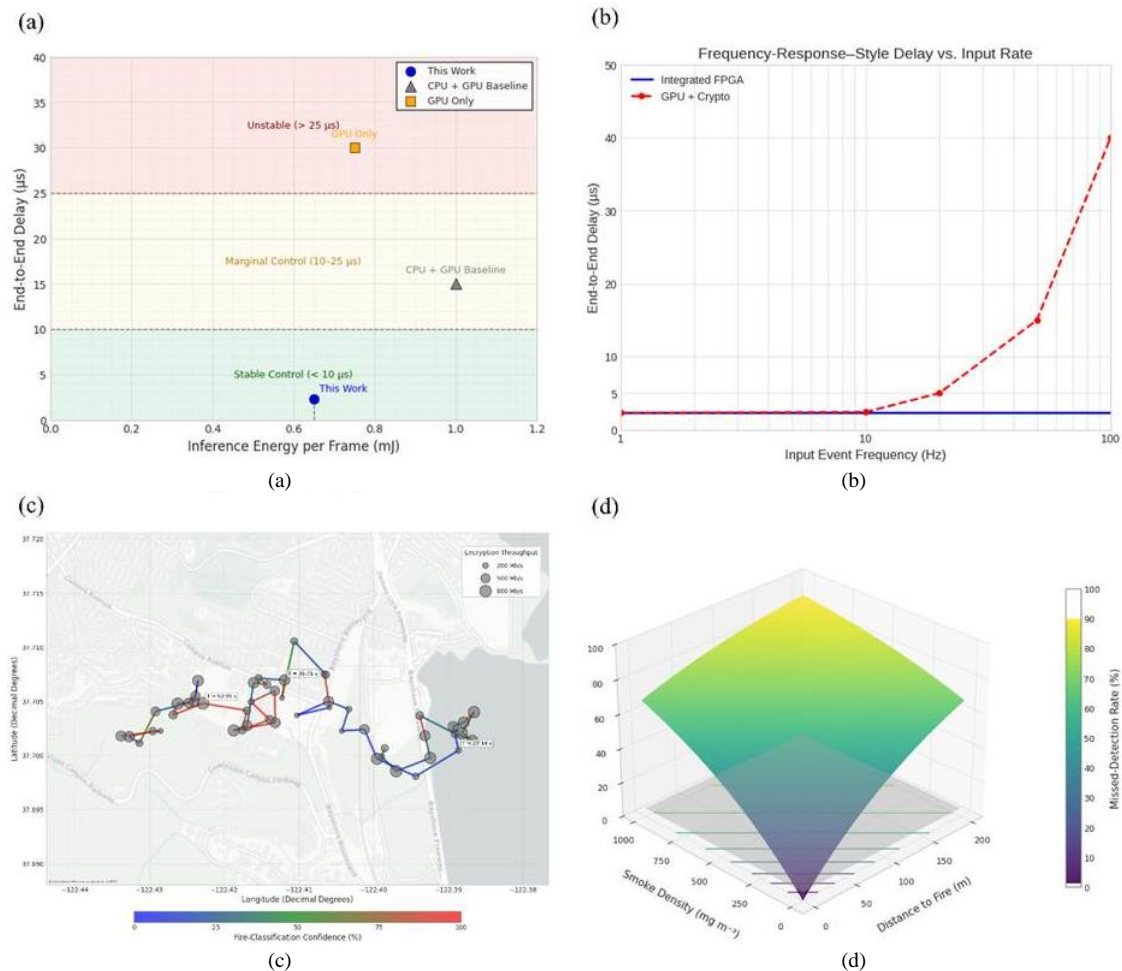


Fig. 14. (a) Phase diagram of inference energy per frame versus end-to-end processing delay. Shaded bands indicate stable ($<10 \mu\text{s}$), marginal (10–25 μs), and unstable ($>25 \mu\text{s}$) control zones. Data points compare the proposed FPGA-integrated system (“This Work”) against CPU+GPU and GPU-only baselines. (b) Frequency-response analysis of end-to-end processing delay as a function of input event frequency. The integrated FPGA system maintains a consistent low-latency profile across all frequencies, while the off-board GPU and cryptographic setup exhibits sharply increasing delay beyond 20 Hz, compromising its suitability for high-rate control loops. (c) Spatiotemporal map of the UAV’s flight trajectory overlaid on a GIS base layer, with line colour indicating real-time fire-detection confidence (blue = 0 %, green = 50 %, red = 100 %) and circle size at each waypoint scaled to encryption throughput (200 Mb/s to 800 Mb/s). Timestamp annotations show key mission milestones. (d) False-negative risk surface showing how the probability of missed fire detections increases with camera–fire distance and smoke density. The semi-transparent plane at 5 % marks the operational safety threshold, and contour projections on the X–Y plane indicate 10 % increments in missed-detection rate

The figure reveals several important patterns. At lower smoke levels (below approximately 300 mg m^{-3}), the risk of a missed detection remains low, even at moderate distances. However, as smoke becomes denser, the risk increases sharply, suggesting that there is a critical threshold beyond which the system’s accuracy deteriorates rapidly. For example, when the UAV is flying at 150 meters in smoke concentrations near 800 mg m^{-3} , the false-negative rate can exceed 60%, significantly compromising situational awareness. This effect appears not to be simply additive—rather, distance and smoke density amplify each other’s impact, causing a compounded loss in detection performance that far surpasses what either factor would cause alone.

To support practical decision-making, a semi-transparent reference plane is drawn at a 5% false-negative rate, which aligns with safety thresholds commonly cited in critical event detection systems. This visual cue defines a “green zone” of reliable operation where fire events are unlikely to go unnoticed. By examining where this plane intersects the surface, UAV operators can identify safe combinations of

altitude and environmental conditions for surveillance missions. For instance, at a typical scouting height of 50 meters and a moderate smoke level around 250 mg m^{-3} , the system maintains a false-negative rate well under 10%, indicating a robust margin for routine perimeter sweeps.

Additional operational insight is provided by the contour lines projected onto the base of the plot. These lines serve as a visual guide, showing how quickly detection reliability declines under worsening conditions. For example, the transition from 10% to 30% missed detection occurs rapidly once smoke levels exceed 600 mg m^{-3} at distances beyond 100 meters. This kind of information is critical for real-time decision-making: it indicates when the UAV should descend, change position, or activate alternative sensing strategies to maintain reliability.

That said, some limitations must be acknowledged. The current model assumes uniform smoke distribution and a clear, unobstructed view of the fire. Wildfire plumes are turbulent and variable, often containing layers of smoke, flame, and hot gases that fluctuate rapidly. These

inhomogeneities could cause the actual risk to deviate significantly from the smooth surface shown in the figure. Moreover, this analysis is based solely on visible-spectrum imagery; incorporating thermal infrared or multispectral sensors could improve detection in dense smoke by capturing heat signatures that are invisible in the visible range. Additionally, the model does not account for possible measurement uncertainty due to factors like sensor calibration drift, ambient lighting variation, or onboard vibrational—all of which could influence detection performance in the field.

The trends observed in this figure can be explained by two main physical effects. First, Mie scattering, caused by fine smoke particles, degrades image clarity by scattering short-wavelength light. Second, geometric spreading reduces the apparent size and brightness of the fire in the captured image as the UAV flies farther away, decreasing the number of meaningful pixels available for analysis. These effects weaken the neural network's ability to recognize fire features, increasing the likelihood of false negatives.

To address these challenges, future designs might incorporate adaptive altitude control, allowing the UAV to automatically lower its flight height in response to deteriorating visibility. Multispectral sensor fusion—which combines visible and infrared data—could help preserve contrast and improve detection in smoky conditions. Furthermore, training the AI system with data from real wildfire environments, including turbulence and lighting variations, may improve its robustness and reduce sensitivity to visual noise. By integrating these strategies, the system could extend its effective operating range and maintain high reliability under a wider range of environmental conditions.

This offers fire-response planners a concrete, data-driven basis for adjusting UAV altitude or rerouting during heavy smoke events to preserve detection reliability.

Because data is securely encrypted onboard, it also prevents malicious interference such as false alarms, spoofing, or tampering with flight commands. This feature is essential for multi-UAV coordination, where several drones share data and responsibilities in a collaborative way. Embedding security directly into the core of the processing pipeline makes each UAV a trustworthy participant in larger networks for fire mapping, emergency response coordination, and predictive modelling. In this way, the system supports not just autonomous flight, but cooperative intelligence, an increasingly important goal in both civil and defense-related UAV operations.

Together, these sections (Real-Time Wildfire Response, Data Security and Communication Integrity Data Security and Communication Integrity, and Autonomous UAV Operations) demonstrate that the proposed system is not just a technological prototype but a practically viable platform for autonomous wildfire surveillance and response. Its co-optimized architecture effectively bridges the gap between academic proof-of-concept and deployable field technology, aligning with the stringent performance requirements and

operational unpredictability of real-world fire monitoring missions.

D. Strengths and Limitations

The proposed system demonstrates several strengths that establish its novelty and practical value. First, the integration of AES-256 encryption directly within the PLIC enables high-throughput, low-latency secure data handling, a key advancement over traditional microcontroller-based encryption architectures. The system achieves a data processing rate of 1.2 Gbps and encryption throughput of 800 Mbps, which is significantly higher than those reported in FPGA-based designs such as AES-32GF [85]. Moreover, the AI-powered fire detection subsystem maintains detection accuracies above 95% across most environmental conditions, outperforming contemporary models like EdgeFireSmoke++ [81] and AF-Net [83], while maintaining low false-alarm rates (0.2%). The system's architecture enables real-time decision-making, with inference times as low as 0.05 s, and sustains energy efficiency conducive to UAV operational constraints (<0.65 mJ per inference). These combined features reflect a robust, high-performance architecture for critical aerial surveillance missions. Despite these strengths, several limitations must be critically examined.

Although the results presented in this study demonstrate strong system performance, it is important to note that all tests were conducted under controlled laboratory conditions. These conditions—characterized by stable temperatures, reliable wireless communication, and consistent power and signal quality—are ideal for benchmarking, but do not fully reflect the unpredictable and often harsh environments where unmanned aerial vehicles (UAVs) typically operate. In practical deployments, UAVs are exposed to dynamic elements such as rapidly changing weather, smoke from active fires, electromagnetic interference from nearby structures or equipment, and unreliable satellite navigation signals, especially in mountainous or forested terrain.

Recent studies, such as [81], have shown that these real-world factors can significantly reduce the performance of UAV-based systems. Reported declines in data processing rates and fire detection accuracy commonly range from 8% to 15% when systems leave the lab and encounter uncontrolled field conditions. These reductions are often caused by a combination of external challenges. For instance, interference in radio frequencies (RF) whether due to natural or artificial sources—can lead to dropped signals or delayed commands, impairing both navigation and communication. Similarly, high levels of heat and smoke near fire zones can distort the thermal data that AI models use to identify fires, increasing the chance of missed or false detections. In addition, network slowdowns or interruptions can disrupt real-time analysis and secure transmission of critical sensor data.

While our system has been extensively tested through simulations and controlled experiments designed to mimic some of these challenges, it has not yet undergone full validation in live outdoor settings. As a result, we cannot yet guarantee that the high performance observed in the lab will hold under real operational pressures. To address this, our

next research phase will include structured field trials in demanding conditions, such as areas with limited GPS access, unstable wireless connections, or rough terrain. These tests are essential to ensure that the system remains fast, accurate, and secure when deployed in the complex environments where it is most needed.

In addition to confirming technical performance, these trials will help us understand how the system responds to sudden environmental changes—like heavy smoke, weather shifts, or momentary signal loss—which are common during wildfire events. The insights gained will support further improvements to the system, such as real-time model adjustments, adaptive encryption settings, and backup communication strategies. By transitioning from lab testing to real-world deployment, we aim to move from a promising prototype to a reliable solution that can support emergency response teams in time-critical situations.

Environmental conditions such as rainfall, strong winds, and dense smoke present significant challenges to the reliability and accuracy of airborne fire detection systems. These factors do not simply obscure visual input; they can also disrupt the underlying mechanisms by which AI models interpret sensor data in real time.

For example, rain affects image quality by scattering and absorbing light, particularly in the thermal wavelengths commonly used to detect fires. Empirical data from recent UAV field tests [84] show that even light rain—around 2 millimetres per hour—can cause a reduction in signal strength of approximately 1.2 to 1.5 decibels. This drop significantly lowers the contrast between flames and the surrounding environment, reducing the ability of the AI system to confidently detect fire. When applied to AI algorithms that rely on visual features, this degradation weakens the system's initial detection layers, which are responsible for identifying flame edges and heat signatures. As a result, the model's predictions may become less certain, requiring more processing time before issuing a reliable alert. In practical terms, this leads to an increase in detection latency by up to 0.3 seconds.

Wind, on the other hand, impacts fire detection by altering the physical shape and behaviour of flames. Gusts exceeding 10 meters per second can rapidly distort the contours of a fire, making it more difficult for AI models to recognize fire patterns from one frame to the next. The system's tracking component, which estimates fire location over time, becomes less stable, forcing it to observe the scene for longer periods before reaching a decision. This instability introduces additional delay, often in the range of 0.15 to 0.25 seconds, which can accumulate alongside other environmental effects to impact response time.

Smoke introduces yet another layer of complexity. In addition to reducing visibility, smoke diffuses both visible and infrared light, which are critical for detecting the heat and colour contrasts associated with fires. The resulting obscurity can confuse AI models, which may shift attention away from actual flames toward less relevant visual cues, such as reflections or heated surfaces. This misdirection leads to misclassifications or delays in detection. Recent visualization analyses show that the AI's internal focus—its

so-called “attention map”—tends to drift in such cases, weakening its overall confidence in identifying fire.

These limitations make it clear that relying on a single type of sensor, such as a standard video or thermal camera, is insufficient under adverse environmental conditions. A more resilient approach involves fusing data from multiple sources. By combining visible imagery with longer-wavelength thermal data, depth sensing (such as LiDAR), or even satellite-based thermal alerts, the system can cross-verify the presence of fire even when one data stream is degraded. Our ongoing research is therefore focused on developing a flexible fusion model that integrates these diverse inputs using a probabilistic framework. This approach allows the UAV system to dynamically adjust to challenging conditions while still maintaining the rapid response times required for effective wildfire detection and mitigation.

Another important concern is the system's potential vulnerability to deliberate interference—known as adversarial attacks—on its wireless communication channels. Although the system uses strong encryption to secure data, this alone is not sufficient to protect against more advanced forms of disruption. For example, attackers can use jamming devices to overwhelm the communication signal with noise, preventing the UAV from sending or receiving critical data. In more targeted cases, adversaries might attempt to impersonate the control system—a technique called spoofing—by sending false commands that appear legitimate. These threats exploit the fact that wireless communication occurs in open space, where signals can be intercepted or disrupted without needing to break encryption.

To defend against these kinds of attacks, the system must go beyond standard measures like routine key updates and data integrity checks. One promising approach is to integrate what is known as a cognitive radio system. This type of system can “listen” to the surrounding radio environment in real time, detect unusual patterns that might signal an attack, and quickly adjust its communication behaviour. For instance, the UAV could switch between different radio frequencies in a random but coordinated way, a method known as frequency hopping, making it much harder for a jammer to block the signal.

Additionally, by incorporating simple machine-learning techniques at the communication layer, the system could monitor for abnormal changes in signal quality or transmission errors, which often indicate interference. If such anomalies are detected, the system could respond by adjusting its communication protocol or switching to backup links. These strategies draw on principles from emerging areas of wireless security and adaptive systems, where technology is designed to respond flexibly and intelligently to hostile or unpredictable conditions. By embedding these capabilities directly into the system's hardware, the UAV can maintain secure and reliable communication even in environments where adversaries are actively trying to disrupt its operation. This added layer of resilience is especially important for critical missions like wildfire monitoring, where communication failures could delay emergency response and endanger lives.

While the AI model used in the system performs strongly under standard conditions, it remains vulnerable to certain unusual or complex scenarios—commonly referred to as edge cases. These are situations that fall outside the typical patterns seen during training and therefore challenge the model's ability to correctly detect fires or distinguish them from non-fire phenomena.

For example, smoldering ground fires—which burn slowly and often under leaves or debris—are especially difficult to identify. They produce minimal visible flames and weak heat signals, which can be easily missed by both standard cameras and thermal sensors. Similarly, reflections in glass windows or other shiny urban surfaces can mimic the flickering appearance of fire, sometimes causing the system to mistakenly identify a fire where there is none. These situations illustrate how environments with subtle or misleading visual features can confuse even well-trained AI systems.

Another difficult scenario involves high-altitude haze or smoke-rich skies following a fire event. In these cases, the faint, dispersed particles in the air can blur the boundary between actual fire activity and harmless environmental conditions. These visual patterns can reduce the model's accuracy, as it struggles to clearly separate the fire from the background.

The current system attempts to handle these challenges by training on simulated examples of such edge cases. While this has improved performance to a degree, it is not enough to fully resolve the issue. The AI still relies primarily on individual images and lacks broader contextual understanding—such as how a scene changes over time or how a fire might appear from different angles.

Improving this capability will likely require more advanced techniques. One such approach is using multiple views of the same scenes such as images from different cameras or sensor types—which can help the model understand the context more completely. Another is incorporating information over time, by analyzing how a potential fire develops from frame to frame. This time-based perspective can help the system distinguish between temporary distractions (like headlights or sunlight reflections) and actual fire events.

In addition, it may be useful to include ways for the AI to estimate its own uncertainty. For example, if the model is unsure whether a fire is present, it could flag the situation for further analysis or request input from a human operator. This type of cautious, confident-aware behavior is especially important in high-stakes environments like wildfire detection, where both missed fires and false alarms carry serious consequences.

Finally, as the system moves from simulation to real-world deployment, it will need to adapt to new environments it has never seen before. This could be achieved by allowing the model to continue learning from actual flight data, even after initial training. Such ongoing learning—done either on the ground or gradually while the UAV is in operation—can help bridge the gap between

controlled training scenarios and unpredictable field conditions.

While the AI component of the system shows promising results under typical conditions, its ability to handle rare or ambiguous situations remains limited. Addressing this will require a combination of richer data inputs, time-aware modelling, self-evaluation mechanisms, and continual adaptation. These improvements are essential for building fire detection systems that are not only accurate in the lab, but also reliable and trustworthy in the real world.

A key limitation of the current system lies in the time required to retrain its AI model—approximately 3600 seconds, or one hour. While this duration is manageable in laboratory conditions, it poses significant constraints in real-world wildfire scenarios, where the environment can change rapidly and continuously. For example, a surface fire may escalate into a crown fire, or shifting winds might suddenly obscure visibility with smoke. In such cases, relying on a static model—one that was trained in advance and does not adapt in real time—can lead to degraded detection performance. This issue reflects a broader challenge in machine learning known as "concept drift," where the patterns in the data change over time, but the model does not evolve accordingly.

To address this, the next phase of development focuses on enabling the system to learn and adapt more quickly through incremental and transfer learning methods. These approaches allow the model to update its knowledge without starting from scratch each time. Specifically, instead of retraining the entire network, the system will preserve the basic image-recognition layers—which identify general visual features like edges and textures—and update only the more specialized layers that interpret high-level fire characteristics. This targeted fine-tuning significantly reduces the amount of time and computational power required. Early experiments suggest that this form of on-ground retraining could be completed within 15 minutes using a lightweight GPU located at the ground control station.

To further enhance adaptability, the retraining process also includes simulated variations in environmental conditions, such as artificial haze or changes in lighting. These variations help the model generalize better to the unpredictable conditions it might face in the field. Importantly, this approach also reduces the risk of "catastrophic forgetting," where new training data causes the model to lose performance on scenarios it had previously learned.

In cases where rapid updates are needed but ground access is limited—such as remote or long-endurance UAV missions, the system may instead rely on in-flight retraining. Here, techniques like model compression and selective updating come into play. By simplifying the model's architecture and focusing updates only on the most critical components, it becomes possible to retrain parts of the model using the UAV's onboard processor, without disrupting other critical tasks like flight control or data encryption. These updates would be based on images the

system finds ambiguous or difficult to classify, ensuring that learning is focused where it is most needed.

To coordinate learning across multiple UAVs observing the same fire event, a method called federated learning is being explored. In this setup, each UAV improves its own model locally and then shares only the updates—not the raw data—with a central server or peer aircraft. These updates are encrypted using the same AES-256 protocol that protects the rest of the system’s communications. Once combined, the updates produce a shared model that reflects a more comprehensive understanding of the environment, all without exposing sensitive data or overloading communication links.

Together, these improvements are designed to transform the AI model from a static tool into a responsive system that evolves with the conditions it monitors. By reducing the time needed for retraining and enabling models to adapt directly in the field, the system moves closer to real-time, reliable operation in high-stakes wildfire detection scenarios.

While the proposed system delivers high-performance secure fire detection under controlled conditions, its deployment in live environments will necessitate further enhancements in environmental robustness, adversarial resilience, and retraining agility. These limitations are being directly addressed in the ongoing system roadmap, which includes multispectral integration, satellite-data cross-validation, and quantum-resilient encryption primitives.

E. Field-Trial Roadmap & Technology Integration

To ensure the proposed system’s real-world viability beyond controlled laboratory simulations, a structured field-validation roadmap has been developed. This roadmap not only addresses the limitations identified in the earlier sections—such as simulation bias, limited sensor diversity, and lack of environmental unpredictability—but also extends the system’s applicability to broader operational scenarios, including industrial complexes, pipeline corridors, and urban canyons. The plan is structured in three progressive phases, each with specific timelines, objectives, and measurable performance benchmarks.

1) Phase 1 (Q3 2025): Controlled burn campaign

The first phase of the planned validation effort centres on a controlled burn campaign that will take place at a certified wildfire research and training facility in the United Kingdom. This type of environment allows for a high degree of experimental control while offering realistic fire dynamics that closely resemble those encountered in natural wildfire events. Over the course of the campaign, ten surface-level fires will be ignited across areas containing different vegetation types, including open grasslands, dense shrublands, and mixed woodland ecosystems. These distinct fuel profiles have been selected to reflect the variety of ignition and combustion behaviours typical in real-world fire-prone regions, and to ensure that the system is evaluated under a broad spectrum of operational conditions.

The primary aim of this phase is to evaluate the UAV system’s performance in detecting and responding to fire activity under a range of smoke densities. Smoking opacity,

a factor that can significantly affect visibility and sensor accuracy, will be deliberately varied during the tests. Specifically, the smoke conditions will span from lightly obscured (measured by a K-factor of approximately 0.2, indicating low aerosol concentration) to severely obscured environments ($K\text{-factor} \geq 1.2$), which represent the upper bounds of what autonomous airborne systems are likely to encounter in active wildfire zones. By exposing the system to this controlled variation, the study aims to understand how different levels of visual interference affect the reliability of both the onboard fire detection algorithms and the secure data transmission process.

The UAV will follow a predefined racetrack flight path—an elliptical loop pattern frequently used in aerial surveillance—that maintains a constant altitude range between 60 and 120 meters above the ground. This range has been chosen to balance regulatory compliance for civilian UAV operations with the need for sufficient altitude to cover wide areas and maintain line-of-sight to ground-based communication units. The use of multiple altitudes will also allow the team to analyse how flight height influences detection accuracy, particularly under low-visibility conditions.

To simulate challenging communication scenarios that are common in remote or infrastructure-poor areas, the bandwidth available for UAV-to-ground station communication will be artificially restricted to 20 megabits per second in half of the test flights. This constraint is intended to replicate conditions where satellite uplinks or long-range wireless networks are limited, such as mountainous regions or deep-forest zones. Testing under such constraints is crucial to verifying that the system’s encryption and data processing components remain effective when transmission capacity is limited scenario that could otherwise compromise timely decision-making in fire management operations.

The performance of the system during this phase will be judged against a set of predefined benchmarks. The system is expected to achieve at least 90 percent detection accuracy in scenarios involving dense smoke while operating at the maximum test altitude of 120 meters. This threshold has been set to demonstrate that the fire detection algorithm remains effective even when visual conditions are significantly degraded. Additionally, the cumulative false-alarm rate, which measures how often the system incorrectly identifies non-fire elements as fires, must remain below 2 percent across all test scenarios. This is critical in ensuring the system does not overwhelm human operators or automated response protocols with spurious alerts.

Equally important is the system’s ability to maintain efficient data encryption and transmission under operational conditions. The UAV must sustain an encrypted data throughput of at least 600 megabits per second, with an end-to-end latency—that is, the delay between data capture and its secure delivery to the ground station—of no more than five microseconds. These parameters ensure that the system can handle real-time data flow securely and without bottlenecks, even under bandwidth-limited scenarios. Finally, the UAV must demonstrate an operational endurance of at least 35

minutes per flight, inclusive of all onboard processing, AI computation, and encryption workloads. This level of endurance is essential for covering large areas or multiple fire zones in a single sortie without the need for premature battery returns or frequent system resets.

Taken together, the objectives of this first-phase campaign are not only to test the system's technical capabilities in realistic fire conditions but also to evaluate its reliability and resilience in the face of environmental and operational stressors. The outcomes of these trials will provide critical insights into the practical deployment potential of AI-enhanced, secure UAV platforms in the context of wildfire monitoring, early warning systems, and real-time disaster response.

2) Phase 2 (Q1 2026): Satellite–UAV data fusion

The second phase of the field-trial roadmap marks a significant step in enhancing the system's ability to detect fires efficiently and accurately. In this stage, the project moves beyond relying solely on UAV-mounted sensors by integrating satellite data to guide aerial surveillance. Specifically, the system begins using thermal data from the Sentinel-2 satellite, part of the European Space Agency's Copernicus program. Sentinel-2 captures images across 13 spectral bands, with resolutions between 10 and 60 meters, enabling it to detect heat signatures that may indicate active fires or developing hotspots. By tapping into this near-real-time data stream, the UAV gains access to large-scale information that can be used to direct its flight and scanning behaviour more intelligently.

Central to this integration is the use of a Kalman filter, a mathematical tool often used in navigation and robotics. It allows the UAV to merge its GPS-based location data with satellite information to estimate the most likely locations of fires. This method does not simply smooth out noisy data; it actively predicts and updates the UAV's scanning path based on the most current and reliable information. As a result, the UAV can focus its attention on areas where fires are most likely, rather than following a rigid pre-planned path. Early estimates suggest this approach could reduce the time needed to search large areas by around 40%, while still maintaining high accuracy.

From a theoretical standpoint, this method follows principles used in intelligent systems that aim to maximize useful information while minimizing time and energy. By combining external satellite data with onboard processing, the UAV evolves from being a passive observer to an active decision-maker. This not only boosts its efficiency but also makes the system more adaptable in complex or rapidly changing fire conditions, where ground support may be delayed. To assess the success of this phase, the system must demonstrate that it can pinpoint fire locations within 15 meters of their actual positions, as verified by detailed maps of fire boundaries. Meeting this target requires precise alignment of all sensors, careful handling of delays in data transfer, and ongoing fine-tuning of the algorithm as field conditions evolve.

In a broader context, this fusion of satellite and UAV capabilities reflects a growing trend in environmental monitoring, where the strengths of space-based and airborne

technologies are combined. For wildfire detection, this hybrid approach bridges the gap between wide-area surveillance from space and targeted action on the ground. It creates a scalable model that can be applied to regions where traditional methods fall short. Thus, the second phase not only upgrades the technical foundation of the system but also establishes a practical method that can be adapted for other real-time, location-sensitive challenges in environmental science and emergency response.

3) Phase 3 (Q2 2026): Sensor diversification and AI retraining

The third and final phase of the field-validation strategy introduces a wider range of sensors to significantly improve the system's ability to detect wildfire in unusual or complex situations. This step addresses known weaknesses in traditional aerial fire detection, particularly when fires are hidden or produce minimal visible or thermal signals. By using more varied types of environmental data, the system becomes more reliable and accurate in real-world conditions.

A key improvement in this phase is the inclusion of sensors that can detect carbon monoxide (CO) and carbon dioxide (CO₂). These gases are typically released during the early stages of smouldering fires, which burn slowly and produce little heat or visible flames. Because such fires often precede larger outbreaks, early detection is crucial. Unlike cameras or heat sensors that may miss these early signs, gas sensors provide a chemical method of identifying fire activity, making them especially useful in places like underground tunnels, forest floors, or debris-filled areas where visibility is poor.

To better map the shape and movement of smoke, a near-infrared LiDAR sensor operating at 905 nanometers is added. LiDAR (Light Detection and Ranging) works by sending out laser pulses and measuring how long it takes for them to bounce back. This allows the system to create detailed 3D images of smoke plumes. Understanding the shape and spread of smoke can help predict how fires move and how smoke affects surrounding areas. This is particularly valuable in settings like urban environments or uneven terrain, where buildings or landscape features can block other types of sensors.

In addition to LiDAR, the system now includes a midwave infrared (MWIR) sensor, which fills the gap between the existing longwave infrared (LWIR) and standard visible-light cameras. MWIR is especially good at detecting moderate levels of heat and can provide clearer images in conditions where traditional sensors struggle, such as through thick smoke or high humidity. By combining data from MWIR, LWIR, and visible-light sensors, the system builds a complete and more accurate picture of its surroundings, improving the ability to identify fire-related activity.

All this new information is processed by the system's artificial intelligence engine using a technique called transfer learning, which allows the AI to learn new patterns without forgetting what it already knows. This process is enhanced using Low-Rank Adaptation (LoRA) modules,

which help the AI adjust quickly and efficiently without needing to be completely retrained. The retraining process is designed to be fast, taking less than 15 minutes on an NVIDIA Jetson Orin NX—a powerful but compact computing device suitable for use directly on UAVs. This rapid adaptability means the system can respond to changing conditions and new fire scenarios in real-time during long missions.

A critical goal for this phase is to achieve at least 92% accuracy in detecting difficult fire cases, such as smoldering fires hidden under debris or sparks from electrical faults. These events are challenging for many detection systems because they are short-lived or not easily visible. By successfully identifying these cases, the system shows its readiness for practical deployment in a wide variety of environments, from cities to remote natural areas. In doing so, this final phase not only completes the technical development but also confirms the system's capability to meet the demands of modern wildfire detection and response.

This field-trial roadmap directly addresses core limitations by introducing environmental realism, multi-sensor diversity, and satellite-assisted tasking. In doing so, it positions the system for broader deployment across domains where conventional UAV detection systems falter. Phase I's real-smoke stress testing validates system robustness under fluctuating atmospheric dynamics, supporting deployment in industrial estates and prescribed-burn management. Phase II's satellite cueing enables efficient regional surveillance, relevant for applications such as pipeline monitoring and forest reserve management. Phase III's sensor augmentation and retainability extend the system's capabilities to detect hard-to-classify incidents like fires in urban canyons or within high-rise architectural recesses. Moreover, integrating real-time satellite intelligence with secure, low-latency UAV telemetry lays the groundwork for scalable, coordinated swarm operations in high-risk emergency response scenarios.

IV. CONCLUSION

This study presents a compact and integrated system that combines fast encryption with real-time fire detection, designed for deployment on drones in environments where rapid response is critical. Using a lightweight artificial intelligence (AI) model and a field-programmable gate array (FPGA) for high-efficiency processing, the system demonstrates strong performance under simulated but realistic fire conditions. It achieved high accuracy in identifying early-stage fires while maintaining a low rate of false alarms and rapid data encryption speeds—essential for protecting sensitive information transmitted by aerial platforms.

Compared to existing drone-based fire monitoring solutions, our platform offers significant improvements. It delivers nearly twice the encrypted data throughput and maintains superior detection precision, all within a low-power embedded hardware setup. This positions the system as a practical and effective tool for early fire warning applications in remote or high-risk areas. Because the FPGA board matches the form factor of standard Pixhawk

controllers and supports over-the-air firmware updates, fleets can be upgraded in the field without replacing airframes or ground-station software.

A novel theoretical contribution of this work is the definition of a computational latency bound that accounts for the interaction between data encryption and AI inference. This helps establish performance expectations as sensor complexity increases and supports future designs aiming to scale the system.

At the same time, several limitations must be acknowledged. The system was primarily tested in controlled environments, with limited variability in fire types and external conditions. This raises concerns about how well the system would perform in real-world scenarios that are more chaotic and less predictable. Additionally, while the AI model used is efficient, it does increase power consumption and may be vulnerable to deliberate attempts to fool it, such as through visual interference (known as adversarial attacks). The AI adds an average 12 % power draw, shortening typical flight endurance from 35 min to 31 min; long-term reliability under 24 h duty cycles remain to be validated. These trade-offs must be carefully considered when transitioning from experimental validation to field deployment.

To address these concerns, future research should focus on expanding the training dataset with diverse, real-world fire imagery and testing the system in outdoor environments across different climates and terrains. Specific next steps include implementing adaptive AI training methods that allow the model to improve over time with new data, upgrading the encryption framework to defend against future quantum computing threats, and coordinating multiple drones to detect fires cooperatively and share information in real time.

Finally, by integrating AI-based analytics with robust, real-time encryption on a single low-power device, this work offers a new approach to secure, autonomous environmental monitoring. It contributes to advancing both the technical feasibility and practical readiness of intelligent drone systems in public safety, environmental protection, and emergency response domains.

ACKNOWLEDGMENT

The authors express gratitude to the Ministry of Higher Education and Science of the Republic of Kazakhstan, which allocated program-targeted funding for 2024–2026. IRN AP23486167.

REFERENCES

- [1] R. Barrett-Gonzalez and N. Wolf, "High Speed Microactuators for Low Aspect Ratio High Speed Micro Aircraft Surfaces," *Actuators*, vol. 10, no. 10, p. 265, Oct. 2021, doi: 10.3390/act10100265.
- [2] N. Gopinath and S. P. Shyry, "Secured: quantum key distribution (SQKD) for solving side-channel attack to enhance security, based on shifting and binary conversion for securing data (SBSD) frameworks," *Soft Computing*, vol. 27, no. 18, pp. 13365–13372, Sep. 2022, doi: 10.1007/s00500-022-07479-w.
- [3] B. Li, Z. Yang, D. Chen, S. Liang, and H. Ma, "Maneuvering target tracking of UAV based on MN-DDPG and transfer learning," *Defence Technology*, vol. 17, no. 2, pp. 457–466, Apr. 2021, doi: 10.1016/j.dt.2020.11.014.

- [4] Y. Wu, J. Gou, X. Hu, and Y. Huang, "A new consensus theory-based method for formation control and obstacle avoidance of UAVs," *Aerospace Science and Technology*, vol. 107, p. 106332, Dec. 2020, doi: 10.1016/j.ast.2020.106332.
- [5] M. M. Alam, M. Y. Arafat, S. Moh, and J. Shen, "Topology control algorithms in multi-unmanned aerial vehicle networks: An extensive survey," *Journal of Network and Computer Applications*, vol. 207, p. 103495, Nov. 2022, doi: 10.1016/j.jnca.2022.103495.
- [6] S. M. Mantrashetti, A. P. Chavan, P. Pawar, H. V. R. Aradhya, and O. S. Powar, "A Novel Algorithm for Aspect Ratio Estimation in SRAM Design to Achieve High SNM, High Speed, and Low Leakage Power," *IEEE Access*, vol. 13, pp. 9942–9954, 2025, doi: 10.1109/access.2025.3527333.
- [7] M. Bakyt, L. La Spada, N. Zeeshan, K. Moldamurat, and S. Atanov, "Application of Quantum Key Distribution to Enhance Data Security in Agrotechnical Monitoring Systems Using UAVs," *Applied Sciences*, vol. 15, no. 5, p. 2429, Feb. 2025, doi: 10.3390/app15052429.
- [8] M. Shirichian, R. Sabbaghi-Nadooshan, M. Houshmand, and M. Houshmand, "A QTCIP reference model for partially trusted-node-based quantum-key-distribution-secured optical networks," *Quantum Information Processing*, vol. 23, no. 3, Mar. 2024, doi: 10.1007/s11128-024-04285-1.
- [9] F. Bayat, "Model Predictive Sliding Control for Finite-Time Three-Axis Spacecraft Attitude Tracking," *IEEE Transactions on Industrial Electronics*, vol. 66, no. 10, pp. 7986–7996, Oct. 2019, doi: 10.1109/tie.2018.2881936.
- [10] K. M. Eisenhardt, "Building Theories from Case Study Research," *Academy of Management Review*, vol. 14, no. 4, pp. 532–550, Oct. 2023, doi: 10.5465/amr.1989.4308385.
- [11] L. Dong, Y. Xie, C. Han, and S. Du, "Active fault-tolerant control of multi-unmanned aerial vehicle system with time-varying topology," *Asian Journal of Control*, vol. 27, no. 3, pp. 1335–1345, 2025, doi: 10.1002/asjc.3510.
- [12] N. Emer and N. S. Özbek, "Control of Attitude Dynamics of an Unmanned Aerial Vehicle with Reinforcement Learning Algorithms," *Avrupa Bilim Ve Teknoloji Dergisi*, no. 29, pp. 351–357, Dec. 2021, doi: 10.31590/ejosat.1021970.
- [13] G. Joshi, "Advanced Control Techniques for Unmanned Aerial Vehicle (UAV) Navigation and Flight Control," *NeuroQuantology*, vol. 20, no. 10, 2022, doi: 10.48047/nq.2022.20.10.nq551258.
- [14] K. Moldamurat, A. Tulembayeva, A. Ryspaev, N. Belgibekov, L. Peryakina, and M. Bakyt, "Computer program in sign language for controlling mobile objects and communicating with people," *International Journal of Public Health Science (IJPHS)*, vol. 14, no. 1, p. 502, Mar. 2025, doi: 10.11591/ijphs.v14i1.24544.
- [15] F. Candan, O. F. Dik, T. Kumbasar, M. Mahfouf, and L. Mihaylova, "Real-Time Interval Type-2 Fuzzy Control of an Unmanned Aerial Vehicle with Flexible Cable-Connected Payload," *Algorithms*, vol. 16, no. 6, p. 273, May 2023, doi: 10.3390/a16060273.
- [16] N. Yıldırım, "Real-time verification of solar-powered forest fire detection system using ensemble learning," *Expert Systems with Applications*, vol. 255, p. 124791, Dec. 2024, doi: 10.1016/j.eswa.2024.124791.
- [17] X. Zheng, "Advanced Solar-Powered Fire Detection System: A Wireless Sensor Node Approach to Early Warning and Forest Fire Prevention," *Highlights in Science, Engineering and Technology*, vol. 62, pp. 90–95, Jul. 2023, doi: 10.54097/hset.v62i.10429.
- [18] Z. Chen, M. Zeng, and Z. Fei, "Joint Unmanned Aerial Vehicle Location and Beamforming and Caching Optimization for Cache-Enabled Multi-Unmanned-Aerial-Vehicle Networks," *Electronics*, vol. 12, no. 16, p. 3438, Aug. 2023, doi: 10.3390/electronics12163438.
- [19] D. V. Apollonov, K. I. Bibikova, V. M. Shibaev, and I. E. Efimova, "Creation Of Algorithms For The Automatic Control System Of The Convertible Unmanned Aerial Vehicle," *Trudy MAI*, no. 122, 2022, doi: 10.34759/trd-2022-122-23.
- [20] F. R. Hashim, "Decentralized 3D Collision Avoidance System for Unmanned Aerial Vehicle (UAV)," *Journal of Advanced Research in Dynamical and Control Systems*, vol. 12, no. 7, pp. 446–460, Jul. 2020, doi: 10.5373/jardcs/v12i7/20202025.
- [21] E. Nejabat and A. Nikoofard, "Switching tube model predictive based controller design for multi-agent unmanned aerial vehicle system with hybrid topology," *International Journal of Robust and Nonlinear Control*, vol. 33, no. 17, pp. 10468–10492, Jul. 2023, doi: 10.1002/rnc.6871.
- [22] C. Liu, M. Wang, Q. Zeng, and W. Huangfu, "Leader-following flocking for unmanned aerial vehicle swarm with distributed topology control," *Science China Information Sciences*, vol. 63, no. 4, Mar. 2020, doi: 10.1007/s11432-019-2763-5.
- [23] K. Moldamurat *et al.*, "Improved unmanned aerial vehicle control for efficient obstacle detection and data protection," *IAES International Journal of Artificial Intelligence*, vol. 13, no. 3, pp. 3576–3576, Jul. 2024, doi: 10.11591/ijai.v13.i3.pp3576-3587.
- [24] T. O. Olayinka, A. A. Olatide, A. D. Oluwagbemiga, and M. O. Okelola, "Internal model control tuned proportional integral derivative for quadrotor unmanned aerial vehicle dynamic model," *Control. Theory Inform.(IISTE)*, vol. 9, pp. 1–10, 2020, doi: 10.7176/cti/9-01.
- [25] W. Li, "Unmanned Aerial Vehicle (UAV) in Precision Agriculture to Identify the Crop Water Shortage by Using Multi-Spectral Sensor," *Open Access Journal of Agricultural Research*, vol. 8, no. 2, pp. 1–4, 2023, doi: 10.23880/oajar-16000303.
- [26] N. T. Hegde, "Design of H-infinity Controller for VTOL Tiltrotor Unmanned Aerial Vehicle," *Journal of Advanced Research in Dynamical and Control Systems*, vol. 12, no. SP3, pp. 1061–1070, Feb. 2020, doi: 10.5373/jardcs/v12sp3/20201352.
- [27] S. Al Azzam, "The AI algorithm for text encryption using Steganography," *The scholar journal for sciences & technology*, vol. 1, no. 1, Jan. 2023, doi: 10.53348/nea4.
- [28] V. Katiyar and S. Mandloi, "Implementation of an Ai-Powered Surveillance System for Industrial Fire Detection with YOLO-V8," *International Journal of Progressive Research in Engineering Management and Science*, vol. 4, no. 10, pp. 1324–1332, Nov. 2024, doi: 10.58257/ijprems36900.
- [29] K. Moldamurat, Y. Seitkulov, S. Atanov, M. Bakyt, and B. Yergaliyeva, "Enhancing cryptographic protection, authentication, and authorization in cellular networks: a comprehensive research study," *International Journal of Power Electronics and Drive Systems/International Journal of Electrical and Computer Engineering*, vol. 14, no. 1, pp. 479–479, Feb. 2024, doi: 10.11591/ijece.v14i1.pp479-487.
- [30] Y. K. Cheng and J. J. Kim, "AI-powered Badminton Video Detection: Enhancing Gameplay Analysis and Training," *J Robot Auto Res*, vol. 4, no. 2, pp. 392–402, 2023, doi: 10.33140/jrar.04.02.05.
- [31] G. Sabbani, "AI - Powered Financial Planning and Analysis (FP and A) Using Cloud Computing," *International Journal of Science and Research (IJSR)*, vol. 13, no. 7, pp. 362–365, Jul. 2024, doi: 10.21275/sr2405234246.
- [32] O. Reddy Polu, "AI-Powered Traffic Violation Detection Using CCTV Footage," *International Journal of Science and Research (IJSR)*, vol. 13, no. 3, pp. 1956–1961, Mar. 2024, doi: 10.21275/sr24038113847.
- [33] Y. Vanapalli and Praveen, "AI-Powered System for Early Stroke Detection Using ECGandPPG," *Journal of Engineering Sciences*, vol. 15, no. 12, pp. 131–140, 2024, doi: 10.36893/jes.2024.v15i12.015.
- [34] D. K. Chaturvedi and M. C. Singh, "Fire Detection System and Spurious (False) Fire Warning Of the Aircraft - An Overview," *Journal of Aerospace Sciences and Technologies*, pp. 336–343, Jul. 2023, doi: 10.61653/joast.v69i2.2019.208.
- [35] M. Bakyt, L. L. Spada, K. Moldamurat, Z. Kadirbek, and F. Yermekov, "Review of Data Security Methods using Low-Earth Orbiters for High-Speed Encryption," *2024 4th International Conference on Ubiquitous Computing and Intelligent Information Systems (ICUIS)*, pp. 1366–1375, Dec. 2024, doi: 10.1109/icuis64676.2024.10867245.
- [36] M. K. M. Hanif, "AI - Driven Transformation in Fire Safety: A Comprehensive Study on AI - Integrated Fire Alarm and Detection Systems," *International Journal of Science and Research (IJSR)*, vol. 12, no. 8, pp. 125–126, Aug. 2023, doi: 10.21275/sr23727132310.
- [37] V. Koibichuk, A. Samoilikova, and M. Habencko, "The effectiveness of employment in high-tech and science-intensive business areas as important indicator of socio-economic development: cross-country

- cluster analysis," *SocioEconomic Challenges*, vol. 6, no. 4, pp. 106–115, 2022, doi: 10.21272/sec.6(4).106-115.2022.
- [38] O. Chernikova, N. Heitzmann, M. Stadler, D. Holzberger, T. Seidel, and F. Fischer, "Simulation-Based Learning in Higher Education: A Meta-Analysis," *Review of Educational Research*, vol. 90, no. 4, pp. 499–541, Jun. 2020, doi: 10.3102/0034654320933544.
- [39] S. Gul and H. Yeo, "Correlation of High-Speed Tiltrotor Stability Predictions with Test Data and Parametric Study," *Journal of Aircraft*, vol. 61, no. 4, pp. 1283–1292, Jul. 2024, doi: 10.2514/1.c037807.
- [40] Y. Fang, Y. Nie, and M. Penny, "Transmission dynamics of the COVID-19 outbreak and effectiveness of government interventions: A data-driven analysis," *Journal of Medical Virology*, vol. 92, no. 6, pp. 645–659, Mar. 2020, doi: 10.1002/jmv.25750.
- [41] Z. Zhang and S. Lin, "Analysis of the Effectiveness and Countermeasures of High-Quality Economic Development in the Pearl River Delta City Cluster," *American Journal of Industrial and Business Management*, vol. 12, no. 05, pp. 984–994, 2022, doi: 10.4236/ajibm.2022.125050.
- [42] M. Bakyt, K. Moldamurat, N. Belgibekov, A. Zhumabayeva, and A. Tilenbayev, "Development and Analysis of the Effectiveness of High-Speed Asymmetric Encryption Methods for Protecting Data from UAVs," *Lecture Notes in Networks and Systems*, pp. 189–199, 2025, doi: 10.1007/978-981-97-9327-3_16.
- [43] M. Bakyt, K. Moldamurat, A. Konyrkhanova, A. Maidanov, and D. Z. Satybaldina, "Integration of Cryptography and Navigation Systems in Unmanned Military Mobile Robots: A Review of Current Trends and Perspectives," *DTESI (workshops, short papers)*, 2023.
- [44] P. Ermakov and A. Gogolev, "Comparative analysis of information integration architectures of strapdown inertial navigation systems for unmanned aerial vehicles," *Trudy MAI*, no. 117, 2021, doi: 10.34759/trd-2021-117-11.
- [45] U. Papa, "Unmanned Aircraft Systems with Autonomous Navigation," *Electronics*, vol. 12, no. 7, p. 1591, Mar. 2023, doi: 10.3390/electronics12071591.
- [46] H. A. Mwenegoha, T. Moore, J. Pinchin, and M. Jabbal, "Error characteristics of a model-based integration approach for fixed-wing unmanned aerial vehicles," *Journal of Navigation*, vol. 74, no. 6, pp. 1353–1366, Nov. 2021, doi: 10.1017/s0373463321000424.
- [47] X. Du, M. Wang, W. Wu, P. Zhou, and J. Cui, "State transformation extended Kalman filter-based tightly coupled strapdown inertial navigation system/global navigation satellite system/laser Doppler velocimeter integration for seamless navigation of unmanned ground vehicle in urban areas," *International Journal of Advanced Robotic Systems*, vol. 20, no. 2, Mar. 2023, doi: 10.1177/17298806231158462.
- [48] V. V. Kiryushkin, S. S. Tkachenko, and E. E. Stryapchev, "Methodology for technical diagnosis of navigation equipment for consumers of the global navigation satellite system of an unmanned aerial vehicle using a barometric altimeter," *Electromagnetic Waves and Electronic Systems*, 2022, doi: 10.18127/j15604128-202201-07.
- [49] Y. Xu, C. Zhang, G. Wang, Z. Qin, and Q. Zeng, "A Blockchain-Enabled Deduplicatable Data Auditing Mechanism for Network Storage Services," in *IEEE Transactions on Emerging Topics in Computing*, vol. 9, no. 3, pp. 1421–1432, 2021, doi: 10.1109/TETC.2020.3005610.
- [50] K. Moldamurat, M. Bakyt, D. Yergaliyev, D. Kalmanova, A. Galyymzhan, and A. Sapabekov, "Geoinformation system for monitoring forest fires and data encryption for low-orbit vehicles," *Computer Science and Information Technologies*, vol. 6, no. 1, pp. 58–67, Mar. 2025, doi: 10.11591/csit.v6i1.p58-67.
- [51] R. Cazazian, "Smart Contracts in Blockchain-based Accounting Information Systems and Artificial Intelligence-enabled Auditing Techniques," *Analysis and Metaphysics*, no. 21, pp. 58–73, 2022, doi: 10.22381/am2120224.
- [52] X. Guo, Y. Zuo, and D. Li, "When auditing Meets Blockchain: A study on applying blockchain smart contracts in auditing," *International Journal of Accounting Information Systems*, vol. 56, p. 100730, Dec. 2025, doi: 10.1016/j.accinf.2025.100730.
- [53] A. Tripathi and J. Prakash, "Blockchain Enabled Interpolation Based Reversible Data Hiding Mechanism for Protecting Records," *JCSST*, vol. 10, no. 5, 2023, doi: 10.4108/eetsis.v10i4.2934.
- [54] M. Y. Drygin, "Technical Diagnostics Of The Main Mining Equipment," *Mining Equipment and Electromechanics*, no. 2, pp. 44–50, Jun. 2020, doi: 10.26730/1816-4528-2020-2-44-50.
- [55] Y. Wu, S. Chen, and T. Yin, "GNSS/INS Tightly Coupled Navigation with Robust Adaptive Extended Kalman Filter," *International Journal of Automotive Technology*, vol. 23, no. 6, pp. 1639–1649, Dec. 2022, doi: 10.1007/s12239-022-0142-7.
- [56] K. Shikada and N. Sebe, "Similarities and differences between exosystem-model-based disturbance observer and model error compensator," *SICE Journal of Control, Measurement, and System Integration*, vol. 17, no. 1, Sep. 2024, doi: 10.1080/18824889.2024.2391624.
- [57] M. Bakyt, K. Moldamurat, N. Belgibekov, A. Zhumabayeva, and A. Tilenbayev, "Development and Analysis of the Effectiveness of High-Speed Asymmetric Encryption Methods for Protecting Data from Low-Orbiting Aircraft," *Lecture Notes in Networks and Systems*, pp. 189–199, 2025, doi: 10.1007/978-981-97-9327-3_16.
- [58] S. Ahmed, S. Khan, and I. Hussain, "Impact of Government Interventions on COVID-19 Outbreak in Different Provinces of Pakistan: Interrupted Time-Series Analysis," *SSRN Electronic Journal*, 2020, doi: 10.2139/ssrn.3688215.
- [59] E. Pimentel, E. Boulianne, and C. Spence, "When audit confronts blockchain," *Accounting, Auditing & Accountability Journal*, Dec. 2024, doi: 10.1108/aaaj-12-2023-6768.
- [60] X. Yan, Y. Lu, L. Liu, and X. Song, "Reversible Image Secret Sharing," in *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 3848–3858, 2020, doi: 10.1109/TIFS.2020.3001735.
- [61] A. A. Dmitriev and D. V. Vasilyeva, "Research and systematization of technical diagnostics of various failures on the main nodes and components of mining equipment mechanisms in JSC Polyus Aldan," *Mining informational and analytical bulletin*, no. S5, pp. 3–11, 2022, doi: 10.25018/0236_1493_2022_6_5_3.
- [62] S. Gul and H. Yeo, "Whirl Flutter Predictions for Distributed Propulsion Tiltrotor Configurations," *Journal of the American Helicopter Society*, vol. 70, no. 1, pp. 1–13, Jan. 2025, doi: 10.4050/jahs.70.012001.
- [63] M. Bakyt, K. Moldamurat, D. Satybaldina, and N. Yurkov, "Modeling Information Security Threats for the Terrestrial Segment of Space Communications," *DTESI*, 2022.
- [64] X. Li, W. Qian, L. Xiao, X. Ai, and J. Liu, "Optimized Design and Test of Geometrically Nonlinear Static Aeroelasticity Model for High-Speed High-Aspect-Ratio Wing," *Aerospace*, vol. 11, no. 12, p. 1015, Dec. 2024, doi: 10.3390/aerospace11121015.
- [65] Z. Liu, L. Luo, and B. Zhang, "An Aerodynamic Design Method to Improve the High-Speed Performance of a Low-Aspect-Ratio Tailless Aircraft," *Applied Sciences*, vol. 11, no. 4, p. 1555, Feb. 2021, doi: 10.3390/app11041555.
- [66] Y. Sun, L. Yan, Z. Sun, and S. Zhang, "A Novel Semi-quantum Private Comparison Scheme Using Bell Entangle States," *Computers, Materials & Continua*, vol. 66, no. 3, pp. 2385–2395, 2021, doi: 10.32604/cmc.2021.012696.
- [67] D. R. Babu and R. Jayaraman, "A Quantum Key Distribution Protocol Based on Random Bell Pair Selection," *Journal of Computer Science*, vol. 19, no. 9, pp. 1160–1169, Sep. 2023, doi: 10.3844/jcssp.2023.1160.1169.
- [68] M. Silva, R. Faleiro, P. Mateus, and E. Z. Cruzeiro, "A coherence-witnessing game and applications to semi-device-independent quantum key distribution," *Quantum*, vol. 7, p. 1090, Aug. 2023, doi: 10.22331/q-2023-08-22-1090.
- [69] F. A. A. Andrade *et al.*, "Autonomous Unmanned Aerial Vehicles in Search and Rescue Missions Using Real-Time Cooperative Model Predictive Control," *Sensors*, vol. 19, no. 19, p. 4067, Sep. 2019, doi: 10.3390/s19194067.
- [70] M. R. Jafarinasab, S. Sirouspour, and E. Dyer, "Model-Based Motion Control of a Robotic Manipulator With a Flying Multirotor Base," *IEEE-ASME Transactions on Mechatronics*, vol. 24, no. 5, pp. 2328–2340, Aug. 2019, doi: 10.1109/tmech.2019.2936760.
- [71] B. Yan, C. Wu, and P. Shi, "Formation consensus for discrete-time heterogeneous multi-agent systems with link failures and

- actuator/sensor faults," *Journal of the Franklin Institute*, vol. 356, no. 12, pp. 6547–6570, Aug. 2019, doi: 10.1016/j.jfranklin.2019.03.028.
- [72] Z. Sui, Z. Pu, J. Yi, and S. Wu, "Formation Control With Collision Avoidance Through Deep Reinforcement Learning Using Model-Guided Demonstration," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 6, pp. 2358–2372, Jun. 2021, doi: 10.1109/tnnls.2020.3004893.
- [73] D. Huang, H. Li, and X. Li, "Formation of Generic UAVs-USVs System Under Distributed Model Predictive Control Scheme," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 67, no. 12, pp. 3123–3127, Dec. 2020, doi: 10.1109/tcsii.2020.2983096.
- [74] S. Acharya, A. Bharadwaj, Y. Simmhan, A. Gopalan, P. Parag, and H. Tyagi, "CORNET: A Co-Simulation Middleware for Robot Networks," *2020 International Conference on COMMunication Systems & NETworks (COMSNETS)*, pp. 245–251, 2020, doi: 10.1109/COMSNETS48256.2020.9027459.
- [75] J. F. Hair, J. J. Risher, M. Sarstedt, and C. M. Ringle, "When to use and how to report the results of PLS-SEM," *European Business Review*, vol. 31, no. 1, pp. 2–24, 2019, doi: 10.1108/EBR-11-2018-0203.
- [76] W. Ma, X. Pi, and S. Qian, "Estimating multi-class dynamic origin-destination demand through a forward-backward algorithm on computational graphs," *Transportation Research Part C-emerging Technologies*, vol. 119, Oct. 2020, doi: 10.1016/j.trc.2020.102747.
- [77] K.-S. Shim, B. Kim, and W. Lee, "Research on Quantum Key, Distribution Key and Post-quantum Cryptography Key Applied Protocols for Data Science and Web Security," *Journal of Web Engineering*, pp. 813–830, Nov. 2024, doi: 10.13052/jwe1540-9589.2365.
- [78] O. Grote and A. Ahrens, "Simulation and Application Purpose of a Randomized Secret Key with Quantum Key Distribution," *Electrical, Control and Communication Engineering*, vol. 18, no. 1, pp. 43–49, Jun. 2022, doi: 10.2478/ecce-2022-0006.
- [79] A. Pradana and L. Y. Chew, "Quantum interference of multi-photon at beam splitter with application in measurement-device-independent quantum key distribution," *New Journal of Physics*, vol. 21, no. 5, p. 053027, May 2019, doi: 10.1088/1367-2630/ab1bbf.
- [80] D. Wang, J. Gao, H. Bai, L. Wang, C. Huo, and J. Yuan, "Application of Quantum Key in Secure Communication for Power Distribution and Utilization," *International Journal of Information and Electronics Engineering*, vol. 9, no. 3, pp. 63–66, Sep. 2019, doi: 10.18178/ijee.2019.9.3.707.
- [81] R. Ghali and M. A. Akhloufi, "Deep Learning Approach for Wildland Fire Recognition Using RGB and Thermal Infrared Aerial Image," *Fire*, vol. 7, no. 10, Sep. 2024, doi: 10.3390/fire7100343.
- [82] I. El-Madafri, M. Peña, and N. Olmedo-Torre, "Real-Time Forest Fire Detection with Lightweight CNN Using Hierarchical Multi-Task Knowledge Distillation," *Fire*, vol. 7, no. 11, p. 392, Oct. 2024, doi: 10.3390/fire7110392.
- [83] X. Hu et al., "AF-Net: An Active Fire Detection Model Using Improved Object-Contextual Representations on Unbalanced UAV Datasets," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 17, pp. 13558–13569, 2024, doi: 10.1109/jstars.2024.3406767.
- [84] S. N. Saydirasulovich, M. Mukhiddinov, O. Djuraev, A. Abdusalomov, and Y.-I. Cho, "An Improved Wildfire Smoke Detection Based on YOLOv8 and UAV Images," *Sensors*, vol. 23, no. 20, p. 8374, Jan. 2023, doi: 10.3390/s23208374.
- [85] S. S. Dhandu, P. Jindal, B. Singh, and D. Panwar, "A compact and efficient AES-32GF for encryption in small IoT devices," *MethodsX*, vol. 11, Dec. 2023, doi: 10.1016/j.mex.2023.102491.
- [86] Z. Liu, Z. Wang, S. Tu, H. Wang, J. Fan, and C. Ren, "Real-Time Secure Video Streaming System Based on FPGA and CUDA Technology," *Proceedings of the 2024 14th International Conference on Communication and Network Security*, pp. 146–152, 2024, doi: 10.1145/3711618.3711642.
- [87] X.-H. Tian et al., "Experimental Demonstration of Drone-Based Quantum Key Distribution," *Physical Review Letters*, vol. 133, no. 20, Nov. 2024, doi: 10.1103/physrevlett.133.200801.