Emergency Management
Science and Technology

REVIEW

Open Access

https://doi.org/10.48130/emst-0026-0009

Emergency Management Science and Technology **2026**, 6: e009

State-of-the-art electronic systems and decision-making architectures for wildfire detection and suppression: a comprehensive review

Osman Acar* and Poyraz Baydemir

Mechanical Engineering Department, Selçuk University, Konya, Türkiye
* Correspondence: osmanacar@selcuk.edu.tr (Acar O)

Abstract

Wildfires represent an escalating global hazard intensified by climate change and land-use change, rendering traditional detection approaches such as satellite monitoring and manual ground patrols insufficient because of high latency and vulnerability to adverse weather. This review critically synthesizes the cited peer-reviewed and technical literature to evaluate the state of the art in electronic sensing systems, embedded artificial intelligence (AI), and autonomous platforms for wildfire management. The study organizes technologies into an engineering taxonomy covering wireless sensor networks, thermal infrared imaging systems, computer vision models, geospatial scaling methodologies, autonomous unmanned aerial vehicles (UAVs) and unmanned ground vehicles (UGVs), and multispectral/hyperspectral fuel condition sensing used for risk assessment before a fire occurs. Multispectral and hyperspectral sensing is therefore treated as a risk assessment and mission prioritization layer rather than as a primary modality for active fire detection. Beyond component-level performance metrics, the review evaluates the physical failure mechanisms and methodological limitations, including thermal crossover, nonfire hot surface confusion, smoke-induced optical degradation, radiative heat exposure, and the limited realism of laboratory or controlled burn validation protocols. It further distinguishes component connectivity from decision-level integration by emphasizing multirate temporal fusion, confidence arbitration, closed-loop suppression feedback, communication-denied edge autonomy, near-fire thermal constraints, and mission-level decision governance. The reviewed evidence suggests that thermal infrared imaging generally provides stronger smoke penetration capability than visible spectrum sensors under the reviewed conditions, whereas edge-optimized deep learning architectures can achieve the real-time detection performance essential for a rapid response. However, individual component maturity does not by itself imply operational readiness; reliable deployment requires standardized benchmarks, robust sensor fusion, degradation-aware autonomy, and validated end-to-end intervention pipelines under extreme environmental conditions. Beyond fully autonomous execution, the review frames wildfire robotics as supervised human–robot teaming in which operators retain mission intent, approval, override, and safety authority while autonomous platforms manage perception, local planning, and exception reporting. Compared with recent integrative surveys that already connect remote sensing, AI, UAV platforms, and wildfire management workflows, the distinctive contribution of this review is a deployment-oriented cross-layer analysis of how these components interact, fail, and must be governed as part of a safety-critical detection and suppression architecture.

Keywords: Wildfire detection, Climate change, Electronic sensing systems, Computer vision in wildfire, Unmanned aerial vehicles

Citation: Acar O, Baydemir P. 2026. State-of-the-art electronic systems and decision-making architectures for wildfire detection and suppression: a comprehensive review. *Emergency Management Science and Technology* 6: e009 https://doi.org/10.48130/emst-0026-0009

Introduction

The management of wildfires has evolved from a seasonal operational challenge into a critical year-round global priority. As climate change accelerates the drying of vegetation and extends fire seasons, traditional reactive suppression strategies are becoming increasingly insufficient. The intersection of robotics, embedded artificial intelligence (AI), and advanced sensing technologies offers a transformative opportunity to shift this paradigm from reactive suppression to proactive detection and autonomous intervention. This review article aims to bridge the gap between isolated academic research findings and the practical engineering requirements of field-deployable wildfire systems. By examining the 'electronic ecosystem' of wildfire management—spanning from the physical sensing layer to the cognitive decision-making algorithms—we provide a comprehensive roadmap for the next generation of autonomous firefighting technologies. This section outlines the motivation behind this transition and defines the scope of the technological survey presented herein. Accordingly, system-level integration is treated here not merely as the connection of sensors, robots, and communication protocols, but as the maintenance of a continuously updated decision state that links perception, uncertainty management, action selection, and post-action verification.

Background and motivation

Wildfire incidents have increased significantly in frequency and intensity over the past two decades, with annual global economic losses exceeding tens of billions of dollars, driven by climate-induced drought, elevated temperatures, and anthropogenic land-use changes[1−3]. Direct property damage, ecological degradation, and human health impacts from smoke inhalation present compounding challenges. Traditional detection infrastructure relies on three primary modalities: satellite-based remote sensing with revisit times ranging from tens of minutes to hours, fixed-position visible-spectrum camera towers with limited spatial coverage, and ground-based patrols constrained by terrain accessibility. These approaches share common limitations: significant detection latency;

https://doi.org/10.48130/emst-0026-0009

frequent false alarms under cloud, fog, dust, glare, water vapor, and other smoke-like visual backgrounds; and severe performance degradation in smoke-obscured or night-time scenarios[1,4–6].

Recent technological advances in electronic systems offer potential improvements across the wildfire monitoring chain. Low-cost wireless sensor networks enable distributed point-based detection with subminute alert latency, thermal infrared imaging provides smoke-penetrating visibility unavailable to visible-spectrum systems, deep learning inference accelerators permit low-latency inference-level on-device fire detection without cloud connectivity dependencies, and autonomous unmanned aerial vehicles (UAVs) extend the surveillance range while maintaining geolocation of suspected ignitions[1–3,7,8]. Despite these capabilities, system-level integration—multimodal sensor fusion, edge–cloud coordination, and multiagent task allocation—remains a fragmented research domain lacking unified architectural frameworks.

A further operational constraint is that wildfire autonomy cannot assume stable infrastructure. During large fires, the quality of cellular links, local gateways, and even global navigation satellite systems (GNSSs) may degrade simultaneously; therefore, edge computing should be interpreted not only as a latency reduction strategy but as the minimum autonomy layer required for disconnected operation.

Similarly, thermal management must be framed as a coupled near-fire survivability problem rather than as ordinary outdoor thermal throttling. UAVs' electronics generate heat internally when the platform is exposed to solar loading, hot plume convection, and distance-dependent radiant heat from flames, making the standoff distance and dwell time part of the decision architecture itself[9–11].

Although recent work has addressed the mechanical and structural aspects of hybrid aerial–ground platforms for wildfire intervention—including morphing airframes, transformation mechanisms, and dual-mode propulsion architectures—the electronic intelligence layer enabling autonomous perception, decision-making, and coordinated action remains underexplored. This review complements those mechanical foundations by synthesizing the electronic systems, embedded AI architectures, and communication frameworks that constitute the 'brain' to the mechanical 'body' of next-generation wildfire response platforms.

A final autonomy-related distinction is also required. In this review, autonomy is not limited to automating perception, path planning, and low-level control. A wildfire robot must also decide what task should be attempted, whether the expected risk reduction justifies the platform's exposure, when the decision should be escalated to a human operator, and which emergency behavior is mandatory when navigation, propulsion, thermal margin, or the reliability of communication collapses. For safety-critical wildfire operations, this decision layer should be framed as supervised autonomy rather than human replacement: Humans retain responsibility for the mission's intent, approval, override, and accountability for high-consequence choices, whereas robots execute bounded perception, navigation, suppression assessment, and fallback behaviors[12–24].

Literature review scope and contribution

Building on recent surveys that already integrate wildfire detection, monitoring, UAV-based remote sensing, and AI-enabled wildfire-management workflows[1–3], this review focuses on a narrower but deployment-critical problem: how the sensing hardware, edge inference, geospatial transformation, decision support, autonomy, and human supervision interact as an interdependent engineering architecture. Rather than treating sensing, AI, and robotics as adjacent technology categories, the review evaluates their cross-layer dependencies under smoke, radiative heat, communication loss, asynchronous data rates, conflicting sensor evidence, limited energy, and safety-critical intervention requirements.

The distinctive synthesis offered by this review is a cross-layer, field-deployment-oriented analysis of wildfire detection and suppression systems. The cited literature is organized into an engineering capability taxonomy spanning five layers, and the analysis evaluates not only the components' performance but also their failure mechanisms, validation realism, temporal integration, safety constraints, mission-level autonomy, and human-supervised operation.

Sensing devices and hardware specifications: Wireless sensor networks, thermal infrared cameras, active-fire imaging systems, environmental ruggedization considerations, and multispectral/hyperspectral fuel condition sensing used for risk assessment before a fire occurs.

I. **Perception algorithms and inference performance:** Object detection models, semantic segmentation architectures, accuracy metrics, and frame latency characteristics for edge deployment;

II. **Geospatial scaling and coordinate transformation:** Pixel-to-meter conversion methods, orthomosaic generation, camera calibration, and GNSS-denied localization techniques;

III. **Alerting frameworks and decision support architectures:** Communication protocols, event-driven messaging systems, multi-source data fusion, and human–machine interface design.

IV. **Autonomous platforms and multiagent coordination:** UAV and unmanned ground vehicle (UGV) integration strategies, task allocation algorithms, and collaborative perception under limited communication.

For each layer, we synthesize the demonstrated capabilities, extract the quantitative performance metrics where available, identify the integration points between layers, and highlight deployment gaps preventing operational maturity.

Beyond component cataloguing, this review evaluates integration as a decision architecture: how asynchronous data streams are temporally aligned, how conflicting sensor evidence is arbitrated, how source confidence is managed under degraded field conditions, and how suppression outcomes feed back into replanning.

This evaluation also covers degradation-aware engineering constraints, including link loss, local-only mission replanning, computing-aware endurance, thermal derating, and near-fire standoff limits. It further incorporates mission-level autonomy and supervised human–robot teaming criteria, including limited-resource priority arbitration, risk-governed path acceptance, emergency state handling beyond ordinary replanning, management of the operator's cognitive load, explainable recommendations, the intervention's granularity, and arbitration of authority between human intent and autonomous safety constraints.

Accordingly, the contribution is not merely an updated bibliography but a critical engineering synthesis that connects component-level capabilities with operational decision risks. This lens foregrounds issues that are often treated separately in prior surveys: thermal crossover and nonfire hot surface confusion, controlled burn versus wildfire validation gaps, asynchronous multisource fusion, source-confidence arbitration, closed-loop suppression assessment, communication-denied edge autonomy, near-fire thermal standoff limits, mission priority governance, and human–robot teaming.

Review organization and evaluation criteria

The review is organized around the five-layer engineering capability taxonomy summarized in Fig. 1. The taxonomy links sensing

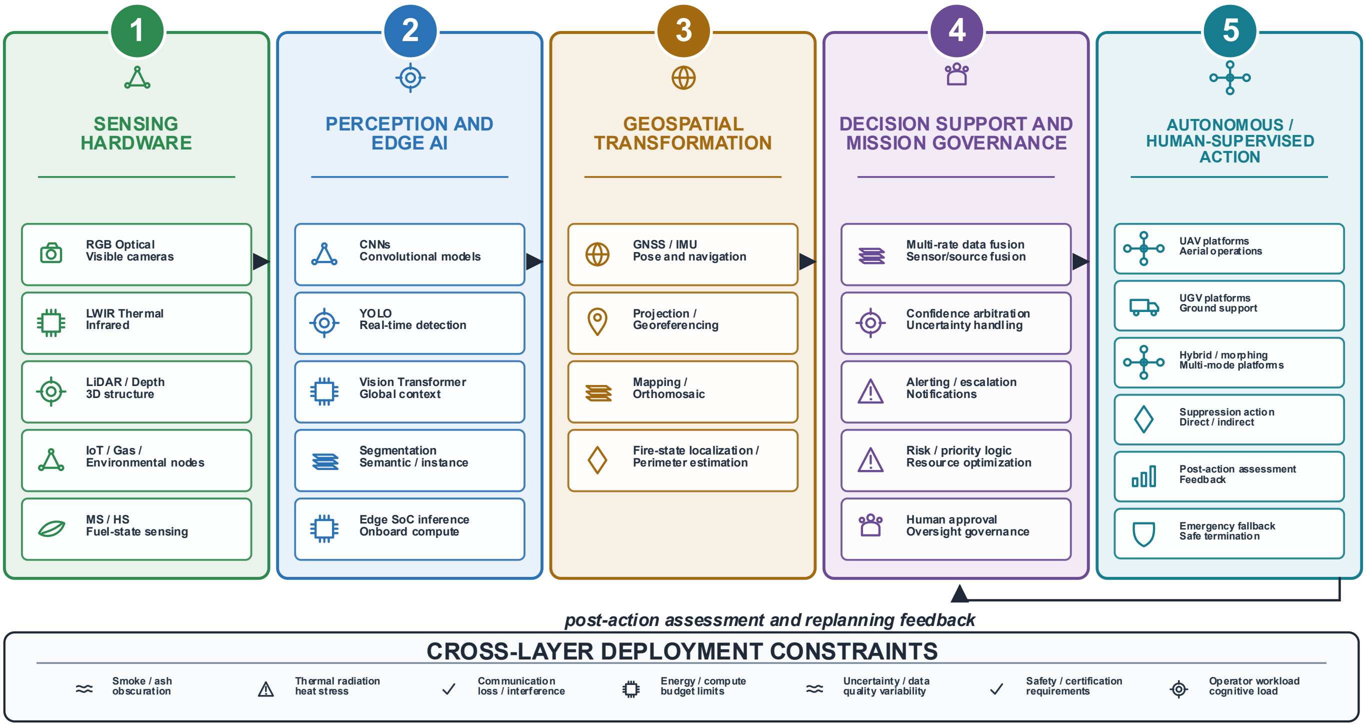

Fig. 1 Five-layer wildfire decision taxonomy.

hardware, perception and edge inference, geospatial transformation, decision support logic, and autonomous or human-supervised action into a single deployment-oriented framework. It is used to structure the synthesis that follows rather than to imply a fixed implementation sequence.

For each technology category, the review records reported performance indicators where available, including detection accuracy, inference latency, operational range, power consumption, and validation settings. These metrics are interpreted cautiously because study conditions, datasets, sensor altitudes, fire scales, and environmental disturbances vary substantially across the literature.

To avoid treating published accuracy or latency values as direct indicators of operational readiness, the evidence base is also assessed qualitatively. The synthesis distinguishes laboratory demonstrations, prescribed fire tests, outdoor field trials, and operational deployments; it also considers whether studies report temporal alignment, how missing data are handled, the weighting of the sources' reliability, conflict arbitration, postaction feedback, communication-denied operation, thermal constraints, mission-level autonomy, and human–robot teaming mechanisms.

For latency terminology, this review distinguishes three uses of the term 'real-time.' Inference-level real-time refers to the per-frame execution by the neural network on embedded hardware, typically measured in tens of milliseconds. Onboard event-level real-time refers to image acquisition, preprocessing, inference, and local event triggering, typically measured from subseconds to a few seconds. Operational end-to-end real-time refers to the complete capture-to-notification workflow, including communication, database updates, coordination logic, and human or operator notification, which may range from seconds to tens of seconds depending on the connectivity. Reported latency values are therefore interpreted according to the latency level explicitly stated.

Technology landscape and engineering synthesis

The review synthesis reveals a diverse technological ecosystem organized into five capability layers. This section synthesizes the literature for each layer, emphasizing system-level characteristics, operational trade-offs, and representative deployment patterns rather than presenting experimental or numerical results from a new dataset. Where quantitative performance data are reported consistently across multiple studies, characteristic ranges are highlighted as engineering benchmarks. However, direct numerical comparison remains constrained by heterogeneous test conditions, dataset differences, and evaluation methodologies.

The reviewed technologies span a maturity spectrum from laboratory proofs-of-concept to field-validated systems. Sensing modalities and computer vision algorithms demonstrate substantial operational deployment, whereas multiagent coordination and autonomous intervention remain predominantly research-focused. Integration across layers, fusing diverse sensor streams, coordinating distributed platforms, and closing the loop from detection to suppression, represents the primary frontier separating component-level maturity from end-to-end system capability.

Accordingly, this synthesis is interpreted in two stages: First, as evidence of the components' capability within each layer; and second, as evidence of whether those components can participate in a shared, time-aware, uncertainty-aware, and feedback-driven operational architecture.

Figure 1 summarizes the five-layer engineering capability taxonomy used to organize the review. The figure emphasizes the logical relationships among capability layers rather than merely being an illustrative technology tree: physical observations are first acquired by the sensing hardware, then interpreted by the edge perception algorithms, transformed into geospatial information on the fire's

https://doi.org/10.48130/emst-0026-0009

state, converted into decision support and mission governance outputs, and finally executed through autonomous or human-supervised platforms. Feedback from suppression or monitoring actions is returned to the decision layer for reassessment and replanning.

In this context, the commonly used abbreviations are introduced in the surrounding text rather than concentrated in the figure caption. Red–green–blue optical imaging (RGB), long-wave infrared imaging (LWIR), light detection and ranging (LiDAR), convolutional neural networks (CNNs), you only look once detectors (YOLO), vision transformers (ViT), system-on-chip edge processors (SoC), UAVs, UGVs, GNSS/global positioning systems (GPS), and simultaneous localization and mapping (SLAM) are treated as layer-specific examples within the taxonomy.

Figure 2 translates this taxonomy into a component-level reference architecture for a deployable UAV/UGV wildfire platform. It shows how multimodal payloads, time synchronization, edge inference, state estimation, mission planning, telemetry, and low-level actuation connect in an executable system. In this review, the figure is used as a bridge between the five-layer taxonomy and the later discussion of closed-loop decision governance; it should not be interpreted as a complete decision architecture by itself.

Table 1 summarizes representative studies and technology categories across the five capability layers introduced in Fig. 1. It functions as a cross-layer capability matrix linking the sensing hardware, perception and edge AI, physical platforms, geospatial and autonomy functions, and the associated engineering trade-offs. Multispectral and hyperspectral entries should be interpreted as support for assessing fuel condition and risk before a fire, whereas RGB, thermal infrared, gas/smoke, and the related active fire modalities support detection and confirmation.

Representative studies supporting the technology categories in Table 1 include those examining wildfire monitoring surveys, UAV-based AI systems, edge-based fire detection, long-range wide-area network (LoRaWAN)/Internet of Things (IoT)-based early detection, data-driven fire-spread forecasting, fuel and susceptibility mapping, UAV/UGV firefighting platforms, hybrid morphing wildfire platforms, examples of fixed-wing/multirotor/UGV platforms, and transformer-based fire segmentation[1–3,5,25–42].

Sensing devices and hardware architectures

Early wildfire detection relies on converting fire's physical signatures into measurable electronic signals. The reviewed literature uses three dominant sensing paradigms: distributed point-based networks monitoring environmental parameters, imaging systems capturing thermal or visual signatures, and spectral sensors analyzing vegetation health indicators. Each modality presents characteristic trade-offs between spatial coverage, detection range, false alarm susceptibility, and deployment complexity.

Distributed sensor networks deploy autonomous nodes across forested terrain, monitoring temperature, humidity, gas concentrations, and particulate levels[62–65]. These systems leverage low-power wide-area communication protocols including LoRaWAN and Zigbee to transmit alerts over distances extending several kilometers while operating for months on battery power[64,66]. Representative implementations integrate commercially available gas sensors for detecting carbon monoxide (CO) and carbon dioxide (such as MQ-series electrochemical cells), humidity probes (DHT22 capacitive sensors), and thermistors or thermocouples for monitoring the ambient temperature[62,67]. Alert thresholds are typically configured to trigger on temperature elevations exceeding 50–60 °C above the background ambient temperature or smoke density surpassing a CO concentration of 300 parts per million (ppm), representing empirically determined values balancing early detection against false alarm suppression[65,66].

For ground robotic support, artificial perception systems for identifying live flammable material can complement fixed IoT nodes by linking local fuel recognition with mobile inspection and suppression planning[42].

The primary advantage of network-based architectures lies in their spatial coverage with minimal infrastructure[64]. Nodes deployed at regular intervals of 100–500 m blanket areas inaccessible to fixed installations, with wireless mesh topologies providing redundancy against individual node failures. However, environmental false alarms from nonfire heat sources (vehicle exhausts, industrial equipment, solar heating of enclosures), vulnerability to weather exposure (moisture ingress degrading electronics), and wildlife damage (rodent nesting or insect colonization of sensor enclosures), and lack of visual confirmation represent persistent challenges[62,63]. Many reported systems address these limitations by implementing edge processing for local anomaly detection before transmitting alerts, using threshold hysteresis and temporal correlation across neighboring nodes to reduce both network traffic and end-to-end latency from detection to notification[64,67]. Multiparameter fusion strategies requiring coincident temperature, smoke, and humidity anomalies further improve specificity, though at the cost

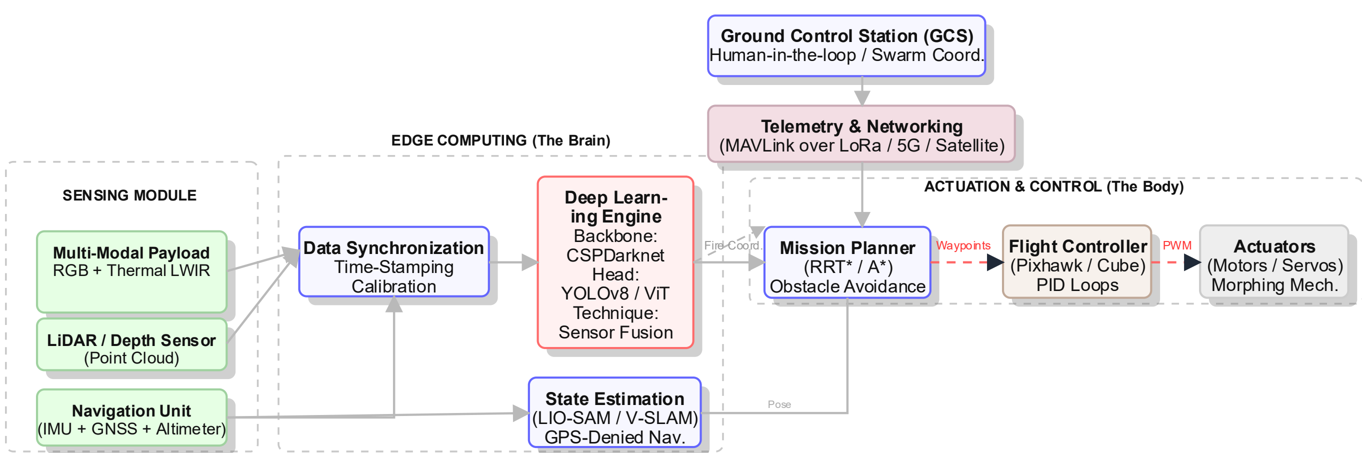

Fig. 2 The deployable UAV/UGV wildfire system's architecture.

Table 1. Technology landscape matrix.

	Subdomain	Key technologies/models	Engineering trade-offs	Ref.
Sensing	Optical (RGB)	Sony RX1R; Zenmuse H20; GoPro; 4K global-shutter cameras	⊕ High spatial resolution (< 1 cm/px) ⊖ Ineffective in smoke and at night	[43,44]
	Thermal (infrared)	FLIR Vue Pro R; Tau 2; Zenmuse XT2 (LWIR)	⊕ Smoke penetration in the 8–14-µm LWIR window ⊖ Thermal crossover and lower spatial resolution	[45,46]
	LiDAR	Velodyne VLP-16; Livox Mid-40; Ouster OS1	⊕ Precise three-dimensional (3D) fuel/structure mapping ⊖ Beam scattering in ash/smoke and high payload cost	[47,48]
	Multispectral	MicaSense RedEdge; Parrot Sequoia; normalized difference vegetation index (NDVI)/fuel-state sensing	⊕ Vegetation stress and fuel–moisture analysis ⊖ High computational and calibration payload	[49,50]
	Internet of Things (IoT)/gas	BME680 carbon monoxide/volatile organic compound sensing; temperature/humidity nodes; long-range wide-area networks	⊕ Preignition and smoldering detection ⊖ Sparse spatial resolution and drift sensitivity	[51,52]
AI and processing	CNN detectors	YOLOv8; EfficientDet; Faster R-CNN	⊕ Low inference latency on embedded edge devices ⊖ False positives under clouds, haze, and hard negatives	[53,54]
	Transformers	ViT; Swin transformer; DeTR	⊕ Global contextual awareness ⊖ High power draw and token-level computation cost	[26,55]
	Segmentation	U-Net; DeepLabV3+; Mask R-CNN	⊕ Pixel-level area and perimeter estimation ⊖ Slow inference for high-resolution frames	[27,56]
	Smoke models	DarkDarkNet; 3D-CNN; motion features	⊕ Early warning potential ⊖ Confusion with fog, steam, and low-contrast clouds	[28,57]
	Edge hardware	NVIDIA Jetson Orin/NX; Raspberry Pi 4; FPGA accelerators	⊕ Onboard inference and autonomous operation ⊖ Energy, memory, and thermal throttling constraints	[4,58]
	Fusion	Early fusion at pixel/feature level; late fusion at decision level	⊕ Cross-sensor reliability and uncertainty reduction ⊖ Calibration, synchronization, and arbitration complexity	[29,59]
Platforms	Multirotor	DJI Matrice 300; custom quadrotors	⊕ Vertical takeoff and landing (VTOL), agility, and static hover ⊖ Limited endurance under continuous flight	[30,31]
	Fixed-wing	eBee X; Wingtra VTOL; solar UAVs	⊕ Large area coverage ⊖ Cannot hover or verify a spot without loitering trade-offs	[32,33]
	Ground (UGV)	Husky; Summit-XL; firefighting tanks	⊕ High payload capacity and mission duration ⊖ Terrain-related mobility constraints	[34,35]
	Hybrid	Morphing TQTV; amphibious or mode-switching platforms	⊕ Mode-switching persistence ⊖ Mechanical complexity and reliability burden	[36]
Autonomy	Localization	GNSS/real-time kinematic (RTK); visual odometry; LiDAR–inertial odometry-SAM	⊕ Absolute geolocation and pose estimation ⊖ Signal loss under canopy and smoke-degraded vision	[37,58]
	Planning	A*; RRT*; artificial potential fields	⊕ Dynamic obstacle avoidance and replanning ⊖ Computing overhead and risk of local minima	[38,60]
	Swarm	Leader–follower control; mesh networking; multi-UAV cooperation	⊕ Scalability and redundancy ⊖ Bandwidth saturation and coordination complexity	[34,61]

* Note: ⊕ = advantage, ⊖ = disadvantage. RRT*: Rapidly-exploring Random Tree Star (an optimized sampling-based path planner); FPGA: Field Programmable Gate Array; TQTV: Transformable Quadcopter-to-Terrain-Vehicle; SAM: Smoothing and Mapping; A*: A-Star Algorithm (a graph-based search algorithm used for dynamic path planning).

of an increased minimum detectable fire size because of the stricter triggering criteria[65].

Thermal imaging systems are frequently reported in the literature because of their smoke-penetrating capability and continuous day and night operation[2,7,67]. LWIR microbolometer arrays operating in the 8–14-µm atmospheric window enable detection through smoke densities that completely obscure visible-spectrum cameras[68,69]. The reviewed implementations span uncooled microbolometer arrays suitable for deployment via handheld or small UAVs to higher-resolution cooled midwave infrared (MWIR) detectors carried by larger aerial platforms. The radiometric calibration accuracy varies significantly between controlled laboratory conditions and field deployments, with atmospheric attenuation and variations in emissivity across vegetation types degrading the absolute temperature measurements. Operational studies report temperature measurement uncertainties ranging from ± 2 °C under ideal conditions to ± 5 °C or greater in high-humidity or smoke-laden atmospheres[70,71]. Recent UAV thermal imaging work further shows that the observation height can alter thermal readings through atmospheric attenuation, surface emissivity, and atmospheric volume emission effects[72].

Long-duration field deployment also requires drift-aware calibration rather than one-time laboratory sensitivity checks. Low-cost MQ-series and broader metal oxide-based gas sensors can experience baseline resistance drift, altered gas response, humidity/temperature cross-sensitivity, heater aging, and contamination from volatile organic compounds or particulates. For example, a 1-year metal oxide sensor dataset reported long-term baseline coefficients of variation of 25.74%–40.77% compared with short-term values of 0.38%–1.21%, illustrating how slowly varying drift can overwhelm fixed alarm thresholds[73]. Similarly, uncooled microbolometer arrays require nonuniformity correction because their pixel responsivity, gain/offset coefficients, self-heating of the cameras, and optical contamination introduce fixed pattern noise and radiometric drift over time. Compact microbolometer non-uniformity correction (NUC) studies report corections of the operating temperature of

https://doi.org/10.48130/emst-0026-0009

over −20 to 50 °C and residual nonuniformity below 0.12% at approximately 14.3-Hz calibration update rates, but such correction must be treated as an explicit maintenance and autonomy requirement rather than an implicit camera property[74]. For deployment during wildfires, practical countermeasures include periodic clean air or blackbody reference checks, temperature/humidity compensation, drift-aware baseline tracking, redundant sensor voting, inspections of the lens/optics after ash exposure, and diagnostic flags that downweight aged or unstable sensors during fusion.

Multispectral fusion approaches combining thermal with visible imagery mitigate false positives from solar-heated surfaces by correlating thermal anomalies with visual flame signatures[67,75,76]. This dual-modality strategy reduces false alarm rates by 40%–60% compared with thermal-only systems while maintaining detection sensitivity, as thermal signatures confirm fires obscured by smoke, whereas visible imagery validates anomalies caused by nonfire heat sources such as reflective metal surfaces or sun-warmed rock formations[64,77]. Satellite-based thermal monitoring from platforms such as MODIS and VIIRS complements ground and aerial sensing by providing global coverage at the cost of temporal resolution, with typical revisit intervals of several hours limiting their utility for early detection of a fire's ignition but enabling large-scale fire progression tracking[53,68].

Gimbal-stabilized mounting on aerial platforms enables target tracking during the vehicle's motion despite adding mechanical complexity and payload mass[1,45]. Power consumption scales with the resolution and frame rate but generally remains compatible with typical UAV power budgets, with representative thermal imaging systems drawing 2–8 W for compact modules and 15–30 W for high-resolution gimbaled units[57,78]. The trade-off between detection range and spatial resolution constrains the system design: longer-range detection sacrifices pixel count, whereas higher resolution reduces the effective range because of the fundamental physical limitations of the sensors governed by the detector element's pitch and the optical aperture's constraints[2,67].

Because assessing the fire risk differs from active fire detection, multispectral cameras are treated here as risk assessment and fuel condition sensors acting before a fire begins rather than primary detectors of active fires. Multispectral cameras capturing discrete spectral bands enable an analysis of vegetation stress and fuel moisture estimation as predictive indicators for fire risk assessment[76,79–81]. Representative systems integrate 4–10 spectral channels spanning wavelengths from visible through to near-infrared and short-wave infrared (SWIR) wavelengths, enabling the calculation of vegetation indices such as normalized difference vegetation index (NDVI) and the moisture stress index (MSI)[3,80]. Healthy vegetation exhibits NDVI values exceeding 0.6, while stressed or senescent vegetation associated with an elevated fire risk demonstrates values below 0.3, providing a quantitative assessment of fuel conditions[79]. Hyperspectral systems with 100–200 contiguous spectral bands provide stronger material classification capability, enabling discrimination of the composition of plant species and precise estimations of the fuel load, but impose substantial computational and data management burdens with data rates exceeding 100 MB/s requiring onboard compression or selective transmission strategies[80,82].

Beyond active detection of ignition, remote sensing and forest parameter extraction studies support the use of UAV, mobile laser scanning, and susceptibility mapping methods for fuel characterization before a fire, tree-level structural assessment, and predictions of the fire danger[83–86]. These sources should be interpreted as support for risk mapping and fuel state rather than direct evidence of rapid operational suppression capability.

SWIR bands operating in the 1.4–2.5-μm atmospheric transmission window penetrate thin smoke more effectively than visible wavelengths, supporting scene interpretation and confirmation under partially obscured conditions where standard RGB cameras may fail[67,68]. Comparative evaluations have demonstrated that SWIR imaging maintains usable target contrast at smoke optical depths exceeding 1.5, corresponding to a reduction in visibility below 100 m, whereas visible-band imagery becomes unusable at optical depths above 0.8[68,69]. However, hyperspectral platforms remain primarily research tools because of the high acquisition costs (exceeding $\$50,000$ for scientific-grade instruments) and computational complexity; operational deployments favor simpler visible–thermal fusion as these have a lower cost (under $\$5,000$ for integrated systems) and reduced computational requirements supporting onboard event-level processing within latency and computational constraints[76,79,80].

Multispectral and hyperspectral imaging for prefire risk assessment

Unlike multispectral sensors that capture discrete bands, hyperspectral imaging (HSI) devices acquire data across hundreds of contiguous narrow spectral bands (typically 400–2,500 nm), enabling the construction of a continuous spectral signature for each pixel. In this review, HSI is therefore positioned primarily as an assessment modality of fuel condition and pre-ignition risk rather than as a direct substitute for RGB-based flame detection, LWIR-based hotspot detection, or gas/smoke sensing. This capability is critical for precise fuel moisture content (FMC) estimations and pre-ignition vegetation stress analyses, which rely on subtle absorption features in the SWIR region that are often missed by standard RGB or multispectral cameras[69,79].

Lightweight UAV spectral sensing for fuel-condition assessment spans true hyperspectral payloads and lower-dimensional multispectral systems. Push-broom HSI devices require precise GNSS–inertial measurement unit (IMU) synchronization because the platform's motion determines the line-to-line accuracy of georectification; however, wildfire-oriented field studies often favor multispectral UAV imagery for predictin dead fuel moisture because the data rates, calibration burden, and edge processing complexity are lower than in full hyperspectral cubes[25,87]. Commercial HSI payloads such as the Headwall Nano-Hyperspec, Resonon Pika L, and compact visible and near-infrared (VNIR) systems remain relevant research options, but their deployment claims should be separated from peer-reviewed evidence on fuel moisture retrieval and operational wildfire sensing.

However, the deployment of HSI for wildfire management faces significant engineering trade-offs. The high dimensionality of hyperspectral data (often > 1 GB per minute) necessitates robust onboard storage or high-bandwidth transmission links. Furthermore, processing these high-dimensional cubes requires dimensionality reduction techniques (e.g., principal component analysis) or specialized one-dimensional (1D) CNN architectures to extract actionable fire indices (like the normalized burn ratio) within onboard latency and computational constraints without overwhelming the edge computing node[25,56].

Physical failure mechanisms of thermal and optical sensing in wildfire environments

Thermal and optical sensors do not detect 'fire' directly; they infer fire's presence from radiance, texture, color, motion, and contextual correlations. This distinction is critical in wildfire environments because the same physical observables can be produced by nonfire phenomena. A radiometric thermal camera measures infrared radiance shaped by the target temperature, surface emissivity, reflected background radiation, atmospheric attenuation, the detector's response, and calibration drift. Therefore, an apparent hot region in the 8–14 μm LWIR band may correspond to active flame, smoldering vegetation, sun-heated rock, exposed soil, asphalt, metallic debris, vehicle exhaust, or recently burned ash. Treating thermal intensity as a direct proxy for flames' presence is consequently a major source of false alarms.

The most important example is thermal crossover. During high-risk fire days, ambient air and ground surface temperatures may exceed 40–50 °C, whereas dark soil, rocks, and metallic surfaces exposed to solar radiation can become substantially hotter than the surrounding vegetation. Under these conditions, the thermal contrast between early flame, smoldering material, and the background decreases, and absolute temperature thresholds become unreliable. In addition, differences in emissivity make objects at similar physical temperatures appear different radiometrically, and low-emissivity metals can reflect the sky's or flames' radiation and create apparent thermal anomalies. Smoke and ash further modify the measurement chain by attenuating, scattering, and spatially blurring the signal, whereas hot gases and turbulence introduce shimmering boundaries that may be misinterpreted as motion cues.

Consequently, robust wildfire perception requires algorithms to discriminate fire from hot surfaces using mechanism-level features rather than isolated frame-level temperature maxima. Useful discriminants include temporal flicker and growth patterns, local thermal gradients, plume motion, spatial morphology, co-occurrence with smoke, multiband spectral response, gas/particulate anomalies, wind direction, the fuel map context, and persistence over time. A stable hot rock or vehicle exhaust may remain spatially fixed and thermally smooth, whereas a flame front typically exhibits turbulent boundaries, rapid temporal variation, coupled smoke generation, and directional propagation influenced by wind and fuel continuity. These distinctions motivate physics-aware sensor fusion and hard-negative evaluation rather than purely appearance-based classification. The deployment implications of these mechanisms are summarized in Table 2.

Computer vision and machine learning architectures

The efficacy of an autonomous wildfire suppression system is fundamentally determined by the computational pipeline responsible for detecting ignition sources, quantifying a fire's intensity, and planning suppression trajectories under strict time constraints. Our analysis of the literature reveals a distinct bifurcation in algorithmic approaches: lightweight CNNs optimized for onboard edge inference, and complex vision transformers (ViTs) requiring ground-based processing.

Deep learning architectures: Accuracy vs. latency trade-offs

State-of-the-art detection frameworks predominantly utilize one-stage object detectors because frame-level detection must support rapid aerial maneuvering. The YOLO series is one of the most commonly reported choices for UAV-based fire detection[90–92]. Comparative evaluations demonstrate YOLO architectures achieving a mAP@0.5 of 0.82–0.90 with inference-level latency of 15–35 ms on an NVIDIA Jetson Xavier NX. These values indicate favorable trade-offs for onboard frame-level detection but should not be interpreted as operational end-to-end alert latency, which also includes image acquisition, preprocessing, communication, geospatial database updates, coordination logic, and operator notification. Two-stage detectors such as Faster R-CNN can report slightly higher accuracy (mean average precision [mAP]@0.5 of 0.88–0.92) but typically require an inference latency of 80–150 ms, making them less suitable for embedded frame-level detection under strict onboard computational budgets[90–93].

Evolution across YOLO versions reflects architectural refinements addressing wildfire-specific challenges, and recent edge-focused wildfire studies have also motivated lightweight detector selection for UAV-based deployment[5,90,91]. YOLOv3 introduced multiscale feature pyramids detecting objects across three resolution levels, improving the detection of small fires at long range with mAP gains of 8%–12% over YOLOv2 but increasing the inference latency to 35–40 ms on Jetson Xavier[90]. YOLOv5 adopted the CSPDarknet53 backbone and path aggregation network (PANet) for bidirectional feature fusion, reducing the parameters from 62 million to 7–47 million (depending on the variant: YOLOv5s/m/x) while maintaining accuracy within 2% of YOLOv3 and achieving an inference of 20–28 ms[90,91]. YOLOv7 incorporated efficient layer aggregation networks (ELAN), demonstrating an inference of 12–18 ms for the

Table 2. Wildfire sensing failure mechanisms.

Failure mechanism	Physical cause	Observed effect	Required countermeasure	Ref.
Thermal crossover	High ambient temperature, sun-heated soil/rock/metal, reduced contrast with the flame background	Nonfire hot surfaces resemble ignition points; fixed thresholds fail	Adaptive background normalization; temporal change detection; hard-negative training	[7,9,11,72,78,88]
Emissivity and reflection bias	Different materials emit and reflect LWIR/MWIR radiation differently	Apparent temperature differs from physical temperature; reflective objects create false hotspots	Radiometric calibration; material-aware priors; RGB/SWIR/LWIR fusion	[7,72,88]
Smoke, ash, and turbulence	Particulate attenuation, scattering, hot gases, convective motion	Blurred boundaries, reduced contrast, unstable apparent object shapes	Multiframe filtering; plume motion analysis; uncertainty-aware detection	[7,11,68,78]
Sensor contamination and heat stress	Dust/ash deposition on optics, vibration, radiative heat flux, electronics heating	Loss of sharpness, calibration drift, thermal throttling, intermittent detections	Ruggedized enclosures; lens protection; thermal management; field reliability tests	[9–11,68]
Contextual ambiguity	Hot engines, campfires, industrial sources, sun glint, fog/cloud/smoke-like distractors	High false positives despite high benchmark mAP	Geospatial context, fuel maps, gas/particulate matter sensors, meteorological constraints	[3,7,8,78,89]
Calibration drift and sensor aging	MQ/metal oxide baseline shift, humidity/temperature cycling, microbolometer gain/offset drift, nonuniform pixel response	Drifting thresholds, fixed-pattern noise, false alarms, missed detections, and inconsistent thermal readings during long deployment	Drift-aware baseline tracking; periodic clean-air/blackbody checks; nonuniformity correction; redundant sensor voting; diagnostic downweighting	[73,74]

https://doi.org/10.48130/emst-0026-0009

YOLOv7-tiny variant (6 million parameters) with a mAP@0.5 of 0.84–0.88[90,92]. YOLOv8 adopted anchor-free detection that eliminated hand-tuned anchor boxes, achieving an inference of 15–22 ms with a mAP@0.5 of 0.87–0.92 on recent wildfire datasets[91,92].

Mathematically, the optimization objective for these models in wildfire scenarios differs from standard benchmarks. The loss function is often modified to prioritize recall (minimizing missed detections) over precision, using weighted local loss to address the class imbalance between large background areas (forest) and small fire targets[92,94]:

$$\mathrm{FL}(p_t) = -\alpha_t(1-p_t)^{\gamma}\log(p_t) \tag{1}$$

In this expression, *p* represents the predicted probability for the positive class, *y* denotes the ground truth label (1 for fire, 0 for background), α_t provides class-specific weighting (typically larger for positive/fire samples), and γ (commonly 2.0) modulates the contribution of easy versus hard examples, forcing the model to learn harder, smoke-occluded samples[92,94]. Training with focal loss improves the detection of small fires by 15%–22% as measured by recall at a precision threshold of 0.75 compared with the cross-entropy baseline[91,92,94].

Although CNNs (e.g., YOLOv8, EfficientDet) excel in local feature extraction, recent studies have explored ViTs and transformer-based segmentation models for their global receptive fields, which can improve the discrimination of smoke plumes, the flame region, and complex backgrounds[26,40,41,47,59]. However, ViTs exhibit O(N^2) computational complexity with respect to the image token count, often rendering them unsuitable for embedded hardware without aggressive quantization or task-specific compression[26,47,59]. Hybrid CNN-transformer architectures such as YOLOv8 with C2f modules integrate transformer-style cross-stage connections while maintaining convolutional backbones; these partially mitigate the computational costs, achieving an inference of 25–35 ms while improving small object detection by 5%–8% mAP compared with purely convolutional variants[91,92].

Segmentation-specific networks remain relevant where planning suppression requires estimates of the fire perimeter, active flame boundary, or hotspot areas rather than only bounding boxes; residual segmentation architectures and fire-specific temporal transformer models provide complementary evidence for this pixel-level requirement[40,95].

Edge computing constraints and thermal throttling

Deploying these algorithms on UAVs introduces severe size, weight, and power (SWaP) constraints. Our review of the hardware implementations indicates that the NVIDIA Jetson series (Nano, Xavier, Orin) is a frequently used embedded platform[91,96,97]. However, running heavy inference models like YOLOv8-Large continuously at over 30 FPS can push the junction temperature of these system-on-chip (SoC) devices above 85 °C, triggering thermal throttling and reducing the inference speed by up to 40%[91,96].

Neural network quantization reduces numerical precision from 32-bit floating-point (FP32) to 8-bit integer (INT8) representation, decreasing the memory footprint by fourfold and enabling specialized integer arithmetic units achieving two- to fourfold improvements in throughput on Jetson platforms[90,91,96]. Quantization-aware training (QAT) fine-tuning for 10–30 epochs demonstrates stronger accuracy retention than post-training quantization, maintaining the mAP within 1%–2% of the FP32 baseline for wildfire detection networks[90,91,96].

The TensorRT optimization framework applies graph-level transformations maximizing throughput: layer fusion combines sequential operations (convolution–batch normalization–activation) into single Compute Unified Device Architecture (CUDA) kernels, reducing dynamic random-access memory (DRAM) bandwidth bottleneck by 30%–40% and improving latency by 15%–25%[90,91]. Benchmark evaluations demonstrated that YOLOv5s optimized with TensorRT INT8 achieved an inference time of 22 ms (45 FPS) on the Jetson Xavier NX compared with 38 ms (26 FPS) for unoptimized PyTorch FP32 implementation, reducing latency by 42% with mAP degradation under 1% (from 0.884 to 0.877)[90,91]. Similar optimizations on the Jetson Orin Nano (20 Trillion Operations Per Second [TOPS] INT8, 15-W thermal design power [TDP]) enable YOLOv8n to achieve an inference speed of 15 ms (66 FPS) for 640 × 640 inputs while consuming 8–12 W, which is suitable for extended UAV deployment on 20,000 mAh batteries providing 4–6 h of continuous operation[91].

Thermal management constrains sustained inference throughput on fanless embedded platforms deployed in outdoor environments where ambient temperatures may exceed 40 °C[96,97]. Dynamic voltage and frequency scaling (DVFS) throttles graphical processing unit (GPU) clock speeds from a nominal 1.1 GHz to 800–900 MHz when the junction temperatures approach the thermal limits of 85 °C, degrading the inference throughput by 20%–30%[96]. Comparative evaluations demonstrated that the Jetson Xavier NX sustains 30 FPS YOLOv5s inference indefinitely at an ambient temperature of 25 °C, degrading to 22–24 FPS at an ambient temperature of 40 °C ambient through thermal throttling[91,96].

Near-fire thermal runaway and operational standoff constraints

The thermal limit of an edge-AI payload should not be reduced to ambient temperature throttling. In wildfire approach missions, the onboard computer, batteries, electronic stability controls (ESCs), sensors, and enclosed avionics already generate heat, whereas the platform simultaneously receives external heat from the flames' radiation, hot plume convection, solar loading, and degraded cooling surfaces contaminated by ash or dust. Fire-resilient aerial robot studies and small unmanned aircraft systems (UAS)-based cyber-physical optimization work show that thermal survivability and mission planning must be co-designed rather than treated as independent subsystems[9,10].

A practical first-order thermal budget can be expressed as follows:

$$\begin{aligned} \mathrm{Q}_{\text{total}} = {} & \mathrm{Q}_{\text{compute}} + \mathrm{Q}_{\text{propulsion/electronics}} + \mathrm{Q}_{\text{solar}} + \mathrm{Q}_{\text{fire radiation}} \\ & + \mathrm{Q}_{\text{plume convection}} - \mathrm{Q}_{\text{effective cooling}} \end{aligned} \tag{2}$$

The critical point is that the fire radiation term changes with the relative position, line of sight, smoke/vegetation attenuation, and the flame front's geometry. Studies of radiant heat modeling have shown that this quantity is inherently spatially and temporally variable, so the assumption of a single ambient temperature is insufficient for near-fire operation[11].

Instead of reporting only whether an embedded board throttles at 85 °C, field-oriented studies should report a distance–time–computation envelope: The permitted dwell time at a given standoff zone while running a specified perception, SLAM, planning, and communication workload. As the platform approaches the flame front, the autonomy stack should progressively derate noncritical tasks, sugh as reducing video streaming, lowering the inference frequency, switching from large to compact models, limiting LiDAR/SLAM update rates where safe, prioritizing thermal and battery telemetry, and enforcing abort or retreat rules before the junction, battery, or enclosure temperatures approach unsafe limits. Table 3 summarizes the resulting thermal operating zones and allowable autonomy behavior.

Dataset dependencies and generalization

A critical limitation identified across the cited literature is the reliance on synthetic or limited-scope datasets (e.g., FLAME [a specific aerial radiometric thermal/RGB dataset used for wildfire management benchmark evaluation], Corsican Fire)[43,67,94]. Representative datasets range from 2,000 to 10,000 annotated images, orders of magnitude smaller than Common Objects in Context (COCO) (118,000 training images) or ImageNet (1.28 million images), necessitating aggressive augmentation[90,91].

Data augmentation strategies tailored for wildfire scenarios include photometric augmentations (brightness adjustment ± 30%, contrast modification with gamma correction $\gamma \in [0.7, 1.5]$, hue–saturation–value jittering ± 10° hue/ ± 40% saturation) and geometric augmentations (horizontal flipping, random cropping 0.7–1.0× scale, random rotation ± 15°)[90,91,97]. Advanced techniques integrate synthetic smoke overlays and Generative Adversarial Network (GAN)-generated fire samples, improving robustness to smoke occlusion by 12%–18% as measured by the detection rate in partially obscured scenarios[43,67,94].

Transfer learning initializes networks with weights pretrained on COCO or ImageNet before fine-tuning on fire-specific datasets for 50–200 epochs, enabling convergence with limited training data while achieving a mAP within 5%–8% of models trained from scratch on a hypothetical 100,000-image wildfire corpus[43,91,94].

Models trained on these datasets often suffer from a significant domain shift when deployed in real-world scenarios with different vegetation types, sensor viewpoints, illumination levels, smoke density, and platform motions[43,91,94]. The severity is visible in reported benchmark gaps rather than only in qualitative claims: a recent segmentation study reported best-model mean intersection over union (mIoU)/F1/overall accuracy (OA) values of 79.4%/76.6%/96.9% on the ground camera-based Corsican Fire dataset versus 84.4%/81.6%/99.9% on the UAV-based FLAME dataset, corresponding to the mIoU and F1 being about five percentage points lower under the Corsican Fire evaluation setting[98]. A separate RGB-T detection benchmark further reported that the low-light, low-contrast Corsican Fire conditions produced combined fire/smoke recall of only 51.8% and a mAP@0.5 of 69.1%, indicating that smoke/background ambiguity can cause severe hard-condition degradation even when precision remains comparatively higher[99]. These values should not be interpreted as universal cross-dataset transfer losses, because training protocols and model families differ, but they quantify the magnitude of dataset sensitivity and justify the need for domain adaptation, hard-negative, and multisite field validation. Recent domain adaptation and large heterogeneous benchmark efforts also show that geographic, platform, and sensor-related shifts should be evaluated explicitly rather than assumed away[41,88,100]. Future research must focus on unsupervised domain adaptation and synthetic-to-real transfer learning to ensure robust operation in diverse wildfire environments[41,43,67,94,100].

Physics-aware fire discrimination and hard-negative evaluation

The abovementioned dataset limitations imply that perception of wildfires should move beyond single-frame pattern recognition toward physics-aware discrimination. A detector trained only to map image texture to a fire label can achieve high accuracy in curated datasets while failing against nonfire hot surfaces, sun glint, dust clouds, fog, steam, or smoke-like cloud formations. This failure mode is particularly dangerous for autonomous suppression systems because false positives may trigger unnecessary dispatch, while false negatives may delay the initial attack during the short window when suppression is most effective.

A field-oriented algorithmic pipeline should therefore combine at least four categories of evidence. First, temporal evidence should capture flames' flicker, the growth rate, evolution of the smoke plume, and the derivative of thermal intensity rather than relying on absolute temperature alone. Second, multimodal evidence should combine LWIR/MWIR thermal imagery with RGB, SWIR, gas, particulate, wind, and fuel condition data to reject hot but noncombusting objects. Third, spatial evidence should analyze the irregular flame front morphology, plume attachment, and the continuity of the anomaly with combustible vegetation. Fourth, contextual evidence should incorporate geospatial priors such as the land-cover class, distance to roads or industrial facilities, topography, wind direction, and historical likelihood of ignition.

Hard-negative mining is central to this transition. Datasets and evaluation protocols should deliberately include sun-heated rocks, asphalt, exposed soil, metallic objects, engines, campfires, prescribed burns, fog, low clouds, steam, dust, and ash-filled scenes that are visually or thermally similar to wildfires' signatures. Performance should then be reported separately for ordinary positives, smoke-occluded positives, small early ignitions, and hard negatives. Such stratified reporting would reveal whether a system is genuinely discriminating wildfires' mechanisms or merely exploiting dataset-specific appearance cues.

Geospatial technologies and estimating the size of fires

Accurate estimation of a fire's localization and size require the transformation of image-space detections into world coordinates—a process integrating camera calibration, platform positioning, attitude determination, and terrain modeling. This transformation chain accumulates errors from multiple sources, with the total localization uncertainty depending on the altitude, sensor quality, and availability of correction services.

Table 3. Thermal operating zones for edge autonomy.

Operating zone	Dominant thermal stress	Compute/autonomy policy	Allowed behavior	Ref.
Monitoring/far field	Ambient plus onboard electronics' heat	Full perception stack, mapping, logging, and normal communication	Wide-area surveillance, hotspot tracking, and cloud synchronization	[8,10,91,96]
Warm approach	Increasing radiative and convective load	Optimized model, reduced streaming, active thermal monitoring	Fire confirmation, local replanning, short approach windows	[9–11,91,96]
Hot standoff	High external heat plus peak onboard computation	Low-power inference, safety supervisor priority, strict dwell timer	Suppression drop, rapid reassessment, immediate exit corridor	[9–11]
Critical/overload	Thermal runaway risk for SoC, battery, or enclosure	Noncritical tasks disabled; abort logic dominates	Retreat, return to base, emergency landing, or handoff to another platform	[9–11]

https://doi.org/10.48130/emst-0026-0009

Error budgets accumulate from multiple sources: GPS-based horizontal accuracy can range from meter-level precision (a 1–3 m circular error is probable) with standard satellite navigation to centimeter accuracy (2–5 cm) with real-time kinematic correction. Additional uncertainty arises from inertial heading error (typically 0.5–2° for low-cost inertial measurement units [IMUs]), camera calibration residuals (a 0.2–1-pixel reprojection error), and the terrain model's vertical accuracy (1–5 m for standard digital elevation models)[15,16]. For a UAV detecting a fire hotspot from 100 m altitude with a 60° field-of-view camera, a combined geolocation uncertainty of 3–8 m is typical without ground control points, improving to 0.5–2 m with RTK–GNSS and surveyed reference targets. This uncertainty directly affects the accuracy of suppression attempts: Water or retardant drops require positioning errors below 5 m for effective targeting, whereas alerting and perimeter mapping can tolerate uncertainty of 10–50 m, depending on the fire's size[16,19,20].

Uncertainty propagation studies reported in the literature estimate total localization error ranging from ± 1.5 m at 50 m altitude using RTK GPS and calibrated cameras to ± 5 m at 200 m altitude with standard GNSS, with angular uncertainties in the IMU's attitude amplifying through the projection geometry as the altitude increases[44,57]. For estimating a fire's size, orthomosaic generation via structure-from-motion stitches the overlapping aerial imagery captured from multiple viewpoints at 70%–80% overlap ratios, achieving ground sampling distances of 2–5 cm per pixel, which is suitable for accurate perimeter mapping and area calculation with uncertainties under 5% for fires exceeding 100 m^2[7,44,76]. This photogrammetric approach enables post-event damage assessment and tracing of a fire's progression across observation epochs separated by minutes to hours.

Smoke occlusion and GPS signal denial in canyon or urban environments necessitate vision-based localization alternatives[45,57,69,78]. Visual–inertial odometry fuses camera-based feature tracking with integrated inertial measurements, enabling relative pose estimation without an external positioning reference, achieving trajectory estimation with drift rates of 0.5%–2.0% of the distance traveled under favorable illumination conditions and in feature-rich environments[44,57]. However, accumulated drift over trajectories exceeding several hundreds of meters requires periodic global position corrections via landmark recognition or loop closure detection when returning to previously mapped areas, with loop closure algorithms based on visual bag-of-words or deep learned embeddings enabling position correction when the accumulated drift exceeds 5–10 m[45,78].

LiDAR-based SLAM provides robust three-dimensional (3D) environment reconstruction in GPS-denied conditions, generating point cloud maps with centimeter-level spatial accuracy over trajectories of 100–200 m[7,45,69,78]. Representative implementations employ Velodyne Puck or Livox mid-range sensors (a range of 15–100 m) integrated with iterative closest point or normal distributions transform registration algorithms for scan matching[45,69]. The computational load and power requirements of LiDAR SLAM, typically consuming 30–80 W including the sensors' power and processing, constrain deployment to larger aerial platforms exceeding a 5-kg maximum takeoff weight or ground vehicles with multihour battery capacity[69,78]. Hybrid visual–LiDAR approaches combining visual feature tracking for short-term motion estimation with LiDAR-based loop closure and mapping reduce the computational requirements while keeping the localization accuracy suitable for autonomous navigation in smoke-obscured or texture-poor environments[45,69,78].

Alerting and decision support systems

Operational fire responses require the integration of detections from heterogeneous sensors into unified situational awareness frameworks that enable spatial query, temporal tracking, and predictive modeling[1,2,48,67,68]. Geographic information systems aggregate observations from satellite-based thermal anomaly products (MODIS, VIIRS, Sentinel-3), aerial surveillance platforms, and ground-based sensor networks into common spatial databases supporting real-time visualization and decision support[48,68,76]. Open-source geospatial platforms including QGIS and GeoServer facilitate spatial operations such as identifying threatened infrastructure within buffer zones (typically a 5-km radius from active fire perimeters), calculating evacuation routes avoiding the predicted fire spread zones according to the wind direction and topographic channeling, or estimating suppression resource deployment times accounting for road network constraints[1,48,67].

Temporal fusion algorithms track a fire's evolution across multiple observation epochs[1,2,7,101]. Kalman filter and particle filter implementations estimate the fire's front position and propagation velocity on the basis of wind speed (sustained and gusts), topographic slope (rise over run), and fuel type classification from land cover databases, updating state estimates as new observations arrive at intervals ranging from minutes (aerial surveillance) to hours (satellite revisits)[2,7,48]. Tracking accuracy depends on the fire's rate of advance (1–3 m/min for grass fires, 0.3–1 m/min for forest crown fires), observation frequency, and the model's fidelity to actual combustion physics, including spotting behavior and ember transport, which introduce stochastic elements that are not captured in deterministic models of spread[7,68]. Monte Carlo-based fire simulation ensembles that capture uncertainty in wind forecasts and fuel moisture improve probabilistic predictions of threatened areas compared with simple deterministic model runs[48,68].

The system's architecture partitions computation between the edge devices and the centralized cloud infrastructure according to the latency level and the reliability of communication[49,65,67,96]. Edge processing on embedded platforms aboard aerial vehicles supports inference-level real-time execution, with per-frame model latency below 30 ms in representative optimized cases, thereby avoiding network round-trip delays approaching 500–2,000 ms over 4G/long-term evolution (LTE) cellular links or longer over satellite connections. Cloud resources remain appropriate for non-immediate tasks, including model training, analysis of historical data, and multivehicle coordination requiring global situational awareness[49,65,96]. Representative data pipelines capture imagery on mobile platforms, perform onboard detection using optimized neural networks, transmit positive alerts via cellular or satellite links, update centralized geospatial databases implemented in PostGIS or similar spatially enabled database systems, and dispatch notifications to the response personnel via mobile applications implementing push notification protocols[49,62,65,67]. At the operational level, end-to-end alert latency from capture to notification ranges from under 1 s for edge-processed detection systems with cellular connectivity to 5–15 s when including satellite link delays and cloud processing overhead[49,65]. These end-to-end values describe the complete alerting pipeline rather than the neural network's inference time alone.

Bandwidth optimization strategies transmit compressed detection metadata including the bounding coordinates, confidence scores, and thumbnail imagery (typically 256 × 256 pixels, JPEG compressed) rather than full-resolution imagery (4–8 megapixeld for raw frames), reducing the upload requirements from 5–10 MB per detected frame to under 100 KB while preserving actionable information content[49,67,96]. This approach enables operation over the bandwidth-constrained links characteristic of remote wildland environments where terrestrial infrastructure is limited or absent, with LoRaWAN and satellite IoT protocols providing last-resort connectivity at data rates of 0.3–50 kbps, which is sufficient for transmitting metadata but inadequate for uploading images[65–67].

This edge–cloud split should be treated as a best-case operating mode rather than the primary safety assumption. In large wildfires, the cellular infrastructure can be damaged, local gateways may lose power, radio frequency (RF) links may be intermittent, and smoke/topography can reduce the links' quality. Therefore, cloud connectivity should be considered an opportunistic resource for global coordination, historical analysis, and post-mission synchronization, whereas safety-critical perception, localization, thermal supervision, local estimation of a fire's state, and short-horizon replanning must remain executable onboard or within a local UAV/UGV mesh, drawing on distributed processing, local routing, and UAV–IoT wildfire architectures rather than assuming persistent cloud availability[8,102,103].

A degradation-aware architecture should define explicit autonomy modes. In networked mode, the cloud or edge gateway can maintain the global fire map and optimize multiplatform allocation. In local mesh mode, UAVs and UGVs exchange compact state summaries and coordinate without cloud access. In isolated mode, a single platform continues with onboard belief updates, conservative replanning, and store-and-forward logging. In survival mode, triggered by combined link loss, GNSS degradation, low battery, or thermal overload, suppression is suspended, and abort, loiter, return to base, or safe landing behavior takes priority.

Communication loss may reduce offloading and increase onboard computation exactly when navigation, thermal monitoring, and replanning become more difficult. Thus, the relevant mission budget should not only consider propulsion endurance but also the following:

$$\begin{aligned} E_{\mathrm{mission}} = {} & E_{\mathrm{propulsion}} + E_{\mathrm{compute}} + E_{\mathrm{sensing}} + E_{\mathrm{communication}} \\ & + E_{\mathrm{cooling}} + E_{\mathrm{actuation}} + E_{\mathrm{reserve}} \end{aligned} \tag{3}$$

Compute-aware planning studies for small UAS systems support treating the cyber workload, thermal degradation, and physical actuation as coupled mission costs rather than independent design variables[10].

Multirate and asynchronous data fusion requirements

A practical wildfire decision-support system cannot assume that all observations arrive simultaneously or with equal reliability. Satellite-sourced thermal anomaly products may be delayed by tens of minutes to hours and have a coarse spatial resolution; UAV and UGV payloads generate frame-level detections at second-scale rates; fixed ground nodes may report temperature, gas, humidity, or particulate anomalies every few seconds; and meteorological feeds update with yet another rhythm. Treating these streams as a single synchronized input risks delayed decisions, duplicated alarms, or suppression plans based on inconsistent estimates of the fire's status.

A more realistic architecture should therefore use a timestamped observation layer. Each incoming detection instance should be presented with its source identity, acquisition time, reception time, geolocation, uncertainty bounds, sensor health, environmental suitability, and confidence. A temporal buffer can then align observations within task-dependent windows: Second-scale windows for local UAV/UGV control, minute-scale windows for incident verification and re-tasking, and longer windows for satellite-based regional correction. Missing observations should not be treated as negative evidence by default; instead, the uncertainty of the fire's estimated status should expand, and the planner should actively request reobservation when the risk remains high.

In this view, fusion is not a one-time operation but a multirate state estimation problem. Fast local sensors update the near-real-time belief state, whereas less frequent satellite and weather products act as delayed, wide-area constraints. Kalman filter, particle filter, Bayesian, or factor graph formulations can be used, depending on the nonlinearity of the fire front's dynamics and the required uncertainty representation. The key architectural requirement is that every decision must be tied to a time-stamped belief state rather than to an isolated detection confidence.

These requirements form the transition from decision support-focused visualization to autonomous navigation: A planner should not adopt the raw detection instances, but a temporally aligned belief state whose uncertainty grows when data are late, missing, or mutually inconsistent. Table 4 summarizes the observation streams that must be fused at different update rates.

Autonomous platforms and navigation strategies

Operating in wildfire environments presents a unique robotic challenge: The simultaneous denial of GNSS signals by the dense canopy and the occlusion of visual features by particulate matter (smoke). Consequently, autonomous platforms must possess navigation capabilities that are independent of the external infrastructure. This section analyzes the state-of-the-art in GPS-denied state estimation, dynamic path planning, and the energetic implications of platform morphology[45,50,51,69,78].

State estimation in GPS-denied environments

The dense forest canopy and thick smoke plumes degrade the reception of GNSS signals, with carrier-to-noise density ratios

Table 4. Multirate observation streams.

Observation source	Temporal behavior	Architectural role	Primary integration risk	Ref.
Satellite thermal hotspots	Tens of minutes to hours; delayed products	Wide-area situational awareness; regional correction	Late evidence, coarse pixels, cloud/smoke contamination	[8,70,104]
UAV-based RGB/LWIR detection	Frame- to second-scale; mission dependent	Rapid verification, hotspot localization, suppression assessment	Motion blur, smoke occlusion, battery-limited coverage	[7,88,89]
UGV/LiDAR/local payloads	Second-scale	Near-ground mapping, local obstacle/fuel assessment	Terrain constraints, ash/dust ghost returns	[7,78,105]
Ground sensor nodes	Seconds to minutes	Early local anomaly detection and persistence checking	Sparse coverage, high sensitivity to false positives	[51,52,57,106]
Meteorological feeds	Minutes to hours	Predicting the fire's spread and risk-weighted planning	Forecast uncertainty and microclimate mismatch	[7,65,66]

https://doi.org/10.48130/emst-0026-0009

dropping below 35 dB-Hz, rendering standard positioning solutions unreliable or completely unavailable[45,69,78]. This necessitates alternative localization strategies based on onboard perception.

Standard visual SLAM (V-SLAM) algorithms, which rely on feature point matching (e.g., oriented features from accelerated segment test [FAST] and rotated BRIEF SLAM3, ORB-SLAM3), can undergo severe tracking degradation in dense smoke because of the Tyndall effect, where suspended particles scatter visible light[45,57,78]. Feature-based approaches extract and match visual keypoints (FAST corners with oriented FAST and rotated BRIEF [ORB] descriptors) across sequential frames, triangulating 3D landmark positions via bundle adjustment, minimizing the reprojection errors[45]. However, V-SLAM exhibits tracking loss rates of 30%–60% in dense smoke compared with < 5% in clear conditions, with failures occurring when smoke occlusion exceeds 50% of the image's area or in texture-poor environments (uniform vegetation, featureless terrain)[45,57].

Our review indicates a shift towards LiDAR–inertial odometry (LIO), which offers improved robustness in visually degraded environments by leveraging the 3D point cloud geometry independent of the lighting or atmospheric transparency[7,45,69,78]. The state estimation problem is formulated as a nonlinear optimization of the robot's pose graph. The objective is to minimize the residual errors from IMUs' pre-integration and LiDAR-based point-to-plane matching:

$$X^* = \underset{X}{\operatorname{argmin}} \left(\sum_{k \in \mathrm{B}} \| r_B(\hat{z}_k, X) \|^2_{\Sigma_B^{-1}} + \sum_{j \in \mathrm{L}} \| r_L\left(\hat{z}_j, X\right) \|^2_{\Sigma_L^{-1}} \right) \tag{4}$$

where, r_B represents the IMU's residual before integration, r_L is the LiDAR's point-to-plane residual, and X denotes the full state trajectory to be optimized[69,78]. Comparative field evaluations demonstrate LIO achieving a position drift of 0.3%–0.8% of traveled distance in smoke-obscured forest environments versus 2%–5% for visual methods, with a LiDAR point density (typically 100,000–300,000 points/s for Velodyne Puck) providing geometric richness robust to smoke attenuation up to a 70%–80% reduction in intensity[7,45,69,78]. Recent studies indicate that LIO systems utilizing 905-nm lasers can maintain < 2% drift over 1 km trajectories in moderate smoke, whereas RGB-based systems diverge within meters[69,78]. UAVs's and UGVs' relative positioning and infrared stereo datasets also show that GPS-denied collaboration and depth perception under visually degraded conditions require explicit multisensor fallback rather than a vision-only assumption[107,108]. The computational and power penalty remains significant: LiDAR SLAM implementations consume 30–80 W (including 10–15 W for the sensor and 20–65 W for processing), limiting deployment to platforms exceeding 5 kg MTOW or ground vehicles with an extended battery capacity[69,78].

Table 5 synthesizes the navigation modality trade-offs under wildfire conditions. Although GNSS provides a low computational overhead, the forest canopy blocks satellite signals. V-SLAM degrades severely in smoke, whereas LiDAR-based systems maintain functionality despite partial atmospheric opacity. Thermal imaging offers smoke penetration but suffers from thermal crossover when the ambient and target temperatures converge.

Representative studies supporting the navigation comparison in Table 5 include work on the relative positioning of GPS-denied UAVs/UGVs, infrared stereo perception in degraded visibility, LiDAR/visual navigation, and autonomous UAV–UGV cooperation[44,45,78,107−110].

Dynamic path planning and obstacle avoidance

In addition to classical planners, recent studies on wildfire and aerial robotics have increasingly examined reinforcement learning trajectory generation, online routing of multiagent UAVs/UGVs, and high-precision UAV-based firefighting as mission-level constraints become coupled with navigation, suppression accuracy, and the platform's survival[39,103,111,112].

Wildfire environments present nonstationary obstacles: Fire fronts advancing at 1–5 m/min, falling trees, and dynamic smoke plumes obscuring perception[51,52,68,92]. Effective path planning must account for the temporal evolution of traversable space while optimizing objectives such as the minimal path length, collision probability, and energy expenditure.

Graph-based search algorithms construct discrete representations of configuration space[51,52,92]. A* (A-star) explores state-space graphs using the heuristic-guided cost function f(n) = $g(n)$ + $h(n)$, where $g(n)$ represents the accumulated cost from the start node to node n, and $h(n)$ estimates the remaining cost to reach the goal, typically using Euclidean or Manhattan distance. Given admissible heuristics, A* can identify an optimal path within the chosen graph representation[51,92]. Implementation on occupancy grid maps (typical resolution: 0.5–2 m) enables real-time replanning at frequencies of 1–10 Hz that are suitable for dynamic environment updates, with eight-connected grid neighborhoods producing smoother trajectories than four-connected variants at modest computational cost increase (25% more evaluated nodes)[51].

Sampling-based planners including rapidly-exploring random trees (RRT and the optimized variant RRT*) construct search trees via randomized sampling, achieving probabilistic completeness without explicit discretization of high-dimensional configuration spaces[51,52,92,111]. RRT* iteratively improves the solution cost by rewiring operations connecting nodes along lower-cost paths:

$$c_{\text{best}}(x_{\text{new}}) = \min_{x_i \in X_{\text{near}}} [c(x_i) + c(x_i, x_{\text{new}})] \tag{5}$$

where, X_{near} identifies nodes within a radius $r = \gamma(\log N/N)^{(1/d)}$, N denotes the tree size, d represents state-space dimensionality, γ is a problem-specific scaling constant, and $c(\cdot)$ computes the edge cost accounting for the distance and collision risk[51,52]. RRT* converges

Table 5. Navigation modality trade-offs.

Navigation modality	Smoke penetration	Canopy penetration	Comp. load	Primary failure mode
GNSS (GPS/Galileo)	Unaffected (RF waves penetrate)	Blocked (multipath signal)	Negligible	Loss of locking under dense trees
Visual SLAM (RGB)	Fails (Tyndall scattering)	Variable (needs texture)	High (feature matching)	White-out in smoke/motion blur
Thermal SLAM (LWIR)	High (8–14-µm window)	Medium (leaf occlusions)	Medium	Thermal crossover (uniform heat)
LiDAR SLAM (LIO)	Wavelength dependent* (905 nm vs 1,550 nm)	High (multiecho return)	Very high (3D registration)	Ghost obstacles from ash particles

* Note: 1,550-nm LiDAR penetrates moderate smoke significantly better than 905-nm LiDAR because of reduced particulate absorption.

asymptotically to the optimal solution as sampling iterations approach infinity, exhibiting a 10%–30% improvement in the path cost over standard RRT in obstacle-rich environments at double to quadruple the computational expense[51,52]. Dynamic replanning variants, including RRT-Connect and Anytime-RRT, adapt plans upon detecting environmental changes, maintaining the trajectory's feasibility at 5–15 Hz update rates[52,92].

Artificial potential fields (APF) generate continuous control via virtual forces: Attractive potential pulling toward goals and repulsive potentials pushing away from obstacles[51,92]:

$$U(q) = U_{\mathrm{att}}(q) + \sum_i U_{\mathrm{rep},i}(q) \tag{6}$$

$$U_{\mathrm{att}}(q) = \frac{1}{2} k_{\mathrm{att}} \| q - q_{\mathrm{goal}} \|^2 \tag{7}$$

$$U_{\mathrm{rep},i}(q) = \begin{cases} \frac{1}{2} k_{\mathrm{rep}} \left(\frac{1}{d_i(q)} - \frac{1}{d_0} \right)^2, & d_i(q) < d_0 \\ 0, & d_i(q) \geq d_0 \end{cases} \tag{8}$$

where, q denotes the current position, U_{att} implements quadratic attraction toward the goal with a gain k_{att}, and $U_{\mathrm{rep,i}}$ applies obstacle repulsion within an influence distance d_0 (typically 5–20 m) from obstacle i with gain k_{rep}[51,92]. APF enables real-time reactive control with minimal computation (evaluatable at 50–100 Hz) but suffers from local minima traps where attractive and repulsive forces balance, halting progress toward unreachable goal configurations[92]. Hybrid approaches combining global A* or RRT* planning with local APF-based obstacle avoidance exploit their complementary strengths: Global methods ensure completeness, whereas local reactivity handles unforeseen dynamic obstacles[51,52,92].

Hybrid architectures and mission endurance

Though multirotors offer agility, their flight endurance (typically 20–40 min) is insufficient for persistent fire line monitoring[45,50,77]. This limitation has driven the development of hybrid aerial–ground vehicles (morphing drones) that can fly over obstacles and land to traverse terrain or monitor from a static position[50,54,96].

A platform's morphology fundamentally constrains its operational capability. Rotary-wing UAVs consume 150–400 W when hovering (given a payload of 2–8 kg), limiting endurance to 20–40 min on typical 200–400-Wh lithium-polymer batteries[45,50,77]. Fixed-wing aircraft achieve 60–120 min via aerodynamic lift at 50–150 W of cruise power but sacrifice stationary observation capability[50,76,77]. Ground vehicles offer extended endurance (2–6 h) at 100–400 W of locomotion power over varied terrain, but face mobility constraints from obstacles and limited operational range (1–3 km from the base station)[55,57,69,78].

Mission-adaptive UAS demonstrations, joint UAV–UGV exploration projects, and low-cost UGV navigation studies show that heterogeneous mobility is technically feasible, but they also reinforce the need to separate demonstrations of mobility from wildfire-grade autonomy, safety governance, and environmental survivability[109,110,113].

Hybrid platforms can provide substantial endurance benefits when they shift from propulsion-dominated flight to landed sensing or low-power ground mobility, but this benefit should be reported as a platform-specific mode-switching ratio rather than as a universal multiplier. For example, with a 710-Wh battery, a landed monitoring load of approximately 60 W for sensors, an onboard computer, communications, and low-power housekeeping yields about 11.8 h of stationary 'sense and wait' operation. The corresponding hover endurance cannot be inferred from a generic hover-power assumption; it must be computed from the measured hover power of the actual aircraft, which depends on the mass, rotor disk loading, payload, wind, propulsion efficiency, and duty cycle. Accordingly, the earlier fixed multiplier comparison has been replaced by a narrower engineering conclusion: Landed or ground-assisted monitoring can reduce propulsion energy costs and extend persistence when the mission allows the platform to stop, monitor, and relaunch rather than hover continuously[50,54,96].

The 'bimodal morphocopter' concept demonstrates quadcopter-to-wheeled-vehicle transformation and has been explicitly proposed for wildfire suppression-oriented hybrid mobility[36]. Its aerial mode enables rapid transit over obstacles, whereas the ground mode can support slower rolling surveillance or stationary monitoring with lower propulsion demand than hovering. Reported scenario-level analyses therefore suggest endurance advantages when the mission profile includes long dwell periods and intermittent repositioning rather than continuous flight; however, the magnitude of the advantage is platform-, terrain-, and duty-cycle-dependent and should be reported in terms of measured power budgets rather than a fixed hover-power ratio[50,54,96].

However, morphing introduces mechanical complexity: transformation mechanisms (articulated joints, deployable wheels), dual-mode propulsion (rotors plus wheel motors), and a reinforced structure to withstand landing impacts increase the component count by 30%–50% and, accordingly, the failure modes[50,54]. Field reliability data from deployments of the prototypes indicate a mean time between failures of 8–12 h of operation for morphing platforms versus 15–25 h for fixed-configuration vehicles, with transformation mechanism faults (jammed actuators, worn gear teeth) accounting for 40%–60% of the recorded failures[50,54]. Nonetheless, mission-level performance metrics demonstrate a net positive impact: Successful completion rates of 70%–85% for morphing systems versus 50%–65% for hovering-only platforms in long-duration (> 60 min) surveillance tasks, with energy constraints forcing a premature return to base representing the dominant failure mode for conventional multirotor platforms [50,51].

However, endurance gains should not be interpreted as purely mechanical. A hybrid platform that lands to monitor a hotspot may reduce propulsion power but still incur sustained computation, sensing, communication-retry, and thermal management loads. For long-duration wildfire missions, the useful endurance metric is therefore computation-aware mission endurance: The time during which the platform can maintain an actionable fire-state estimate while preserving its battery reserves, thermal margin, and safe retreat capability.

Synthesis of demonstrated capabilities

The reviewed literature demonstrates substantial technological maturity in individual system components but reveals integration gaps that constrain end-to-end operational capability. Sensing modalities have progressed from research prototypes to field-deployed systems, with thermal imaging and distributed sensor networks achieving operational status in wildfire management contexts. Computer vision algorithms based on deep learning architectures have demonstrated sufficient detection accuracy for practical deployment, with inference-level latency compatible with frame-level onboard detection on embedded hardware platforms.

However, performance comparisons across the reported systems remain constrained by heterogeneous evaluation methodologies, dataset differences, and inconsistent reporting of the operational conditions. Detection accuracy metrics vary substantially depending on the test set's composition, environmental factors including

https://doi.org/10.48130/emst-0026-0009

smoke density and illumination, and the specific definition of successful detection. Latency measurements inconsistently account for end-to-end system delays versus the isolated algorithmic inference time. Communication range specifications depend critically on the terrain, antenna placement, and RF environment characteristics.

These limitations notwithstanding, consistent patterns emerge regarding capability–cost trade-offs. Multimodal sensor fusion combining thermal and visible imaging consistently outperforms single-modality approaches in terms of both detection accuracy and reduced false alarms. Edge processing architectures achieve lower end-to-end latency than cloud-dependent pipelines but constrain the model's complexity and limit coordination across distributed platforms. Aerial platforms provide broader spatial coverage and deployment flexibility compared with fixed installations, at the cost of limited endurance and payload capacity.

The primary technical frontier separating demonstrated component capability from operational system deployment lies in robust integration across the sensing, perception, geolocation, communication, and autonomous coordination layers. Few reported systems characterize performance throughout the full pipeline from ignition to a verified alert under realistic field conditions including smoke occlusion, communication dropouts, and the need for persistent multiday operation.

This synthesis therefore leads directly to the central discussion question: Whether the reviewed systems provide only component connectivity or whether they define the arbitration, feedback, and uncertainty logic required for operational decision-making.

Comparative analysis and discussion

The cited literature reveals a critical divergence in wildfire-related robotics: Although mechanical platforms (the body) have evolved towards multimodal hybridity, electronic architectures (the brain) remain fragmented.

The following discussion therefore separates three levels that are often conflated in the literature: Components' performance, the connectivity between components, and closed-loop decision architectures.

Trade-offs in detection architectures

Our analysis of the literature indicates that no single AI architecture is consistently optimal across all metrics (see Table 1).

I. Latency vs. accuracy: For initial attack drones, lightweight CNNs (e.g., YOLOv8-Nano) offer the necessary low latency (< 30 ms) but can suffer from higher false negative rates when smoke density, motion blur, or thermal background clutter increase.

II. Robustness vs. efficiency: ViT systems can provide stronger robustness against smoke occlusion but require heavy computation nodes (e.g., Jetson AGX) that reduce flight endurance by 15%–20%.

The hybrid imperative

The integration of these technologies points towards a heterogeneous swarm architecture. As illustrated in the taxonomy (Fig. 1), future systems will likely combine fast RGB scouts with heavy thermal-equipped hybrid platforms that utilize 'sense and wait' strategies to balance energy consumption with processing power.

Figure 3 focuses on the cyber–physical connectivity layer that enables such heterogeneous teams. It distinguishes video streaming, telemetry, status exchange, and swarm synchronization among UAV nodes, UGV nodes, edge gateways, and cloud backends. The figure is discussed here to clarify a key argument of the review: Protocol connectivity is necessary for coordination but it is not equivalent to decision-level integration unless it is coupled with temporal fusion, confidence arbitration, and feedback-based replanning.

Table 6 synthesizes technologies' suitability across operational phases. Thermal imaging appears to be a highly versatile solution in the reviewed scenarios, maintaining high efficacy across active fire and night scenarios. Conversely, RGB cameras, though cost-effective, degrade severely under smoke or darkness. Hybrid aerial–ground platforms demonstrate operational flexibility critical for multiday suppression campaigns.

Representative studies supporting the operational suitability comparison in Table 6 include thermal remote sensing of active fires, UAV-IoT-based detection, UAS infrared pyrometrics, thermal sensor–drone collaboration, morphing wildfire platforms, and AI-enabled UAS wildfire management surveys[2,8,36,71,76,87,106,114].

However, heterogeneity also creates the architectural problem emphasized in this review: Platforms with different sensors, update rates, energy budgets, and failure modes cannot be coordinated reliably unless their observations are fused through explicit temporal alignment, confidence management, and feedback loops.

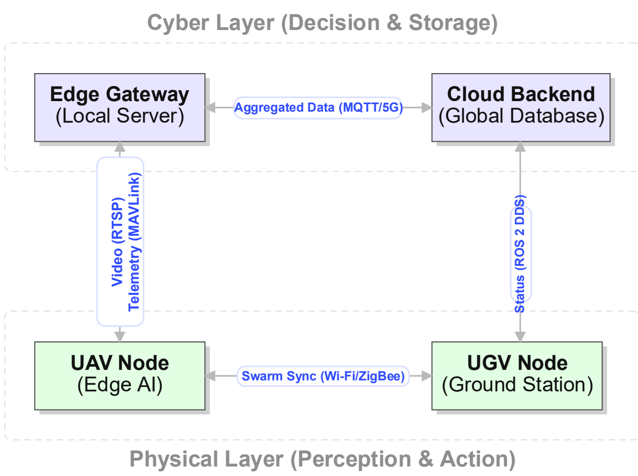

Fig. 3 Cyber–physical connectivity architecture.

Table 6. Operational suitability matrix.

Technology/scenario	Early ignition	Active crown fire	Smoldering	Night operation
RGB camera	High (smoke-free)	Low (smoke occlusion)	Low (ash blending)	None (requires light)
Thermal (LWIR)	Medium (range limits)	High (penetration)	High (hotspot detection)	High (passive infrared)
Satellite (LEO)	Low (spatial resolution)	High (macro-view)	Medium (revisit time)	Medium (swath dependent)
Hybrid UAV–UGV	High (rapid response)	Medium (turbulence)	High (Endurance)	High (sensor fusion)

System-level integration beyond connectivity

The preceding architecture and connectivity diagrams are useful for identifying which components exchange information, but they should not be mistaken for a complete system-level decision architecture. In the response to a wildfire, integration is not equivalent to connecting sensors, robots, gateways, and cloud services. True integration requires the system to maintain a time-stamped and uncertainty-aware belief state, decide which sources to trust under degraded conditions, arbitrate conflicting evidence, act under safety constraints, and then measure whether the intervention has changed the fire's state.

Figure 4 visualizes this distinction by showing the closed-loop path from multisource observations to confidence-managed fusion, mission/risk governance, human authorization, action execution, and post-action updates of the belief state.

Multirate temporal fusion

A deployable architecture should fuse observations at the rate at which they are trustworthy, not at the rate at which a diagram implies they are connected. A ground node reporting a temperature spike every few seconds should update a local ignition probability hypothesis, while a UAV's thermal frame should refine the fire's perimeter and hotspot geometry when the platform is close enough for reliable geolocation. A satellite hotspot, although valuable for regional awareness, may arrive too late for immediate targeting of suppression and should be treated as delayed evidence that corrects or validates the broader incident map. Multirate filtering therefore requires explicit handling of the acquisition time, communication latency, and observation age.

The common state should include, at minimum, the probability of a fire, the location of the fire front or hotspots, intensity estimates, the direction of smoke plumes direction, spatial uncertainty, temporal freshness, and expected spread direction. When a sensor stream is missing, the belief state should not collapse into an artificial 'no fire' conclusion; instead, uncertainty should grow, and the planner should decide whether to re-task a UAV, request confirmation from neighboring ground nodes, or wait for satellite/cloud updates depending on the risk and resource availability.

Confidence management and conflict arbitration

Conflict resolution is a core architectural requirement because wildfire sensors fail in different ways. A thermal camera may report a high-temperature anomaly from sun-heated rock or metal; RGB vision may miss a flame behind smoke; LiDAR may generate ghost returns from ash or particulate scattering; and GNSS may drift under canopy or in canyon-like terrain. A decision system that simply averages these outputs can become overconfident in precisely the conditions where caution is required.

A practical confidence manager should assign each observation a dynamic reliability score based on source health, environmental suitability, temporal freshness, spatial uncertainty, prior false alarm behavior, and cross-sensor agreement. For example, if GNSS uncertainty increases while LIO/visual-inertial odometry (VIO) remains stable, the planner should downweight GNSS for targeting while preserving it as a weak global prior. If thermal evidence is positive but RGB evidence is negative, the system should not automatically suppress; it should check temporal persistence, local background-normalized temperature gradients, gas/particulate evidence, and UAV-based re-observation from a different angle. If a ground node

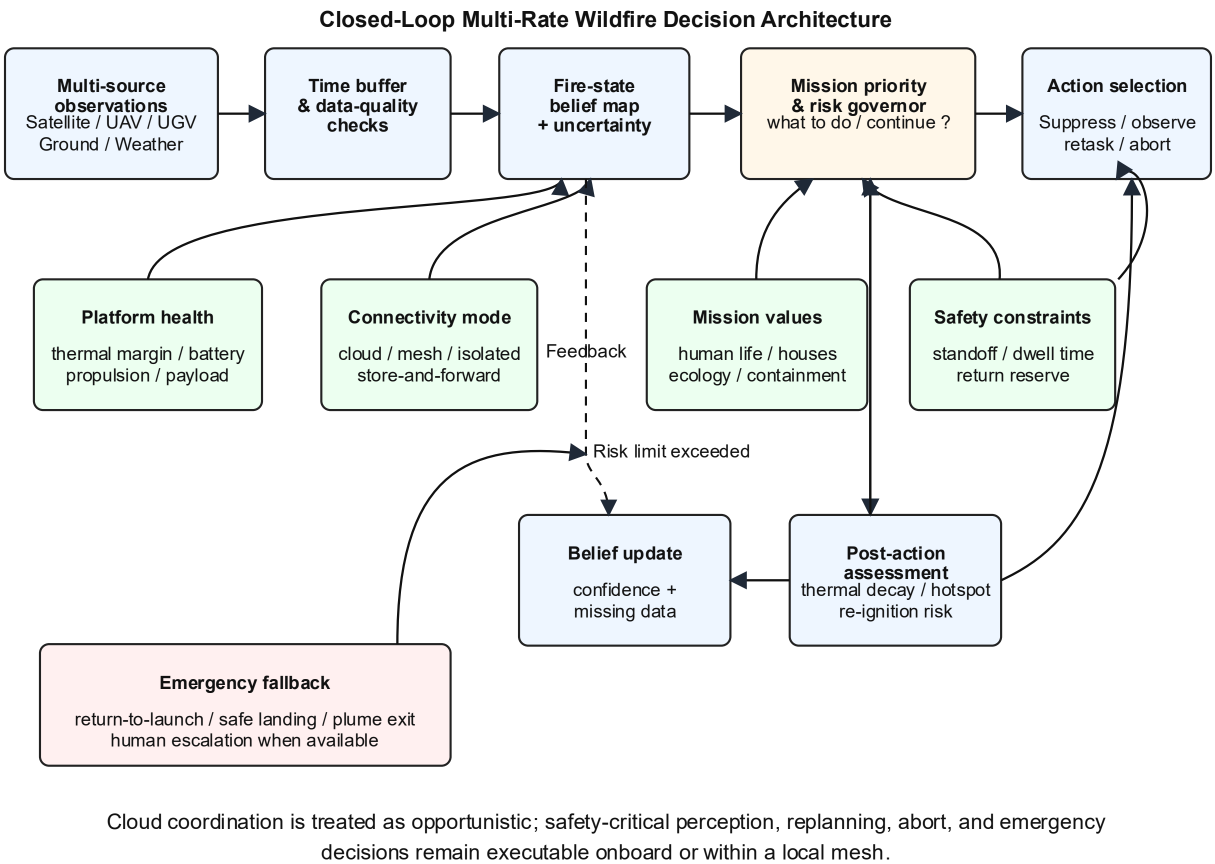

Fig. 4 Closed-loop supervised wildfire decision architecture.

https://doi.org/10.48130/emst-0026-0009

alarms but aerial verification is negative, the system should correlate neighboring nodes and decide whether the event is a localized false positive, an occluded smoldering source, or a stale alert. Table 7 summarizes the corresponding integration failure modes and countermeasures.

Closed-loop suppression feedback and replanning

For a suppression-oriented system, the perception–planning–action process is insufficient unless the action's outcomes are measured and fed back into the decision state. After a UAV releases water, retardant, or another suppressant, the system should automatically reassess the target area using thermal decay, reduction in the hotspot area, weakening of the smoke plume, change in the flame front's geometry, and the probability of reignition. A successful intervention should reduce the thermal intensity and spatial growth over a defined assessment window; an unsuccessful intervention should trigger re-attacks, an altered approach geometry, additional platform allocation, or escalation to human operators.

This feedback loop also improves perception and planning. If a region remains thermally active after suppression, the belief state should preserve high risk even if visible flames temporarily disappear. If the target cools but downwind smoke and neighboring ground nodes continue to intensify, the planner should infer that suppression may have shifted the active front rather than eliminated the incident. Therefore, post-action assessment should be treated as a mandatory state update, not as an optional reporting step.

Architectural implication

The central implication is that future wildfire systems should be evaluated not only by the detectors' accuracy or communication connectivity, but by the quality of their closed-loop decisions. Relevant metrics include time-stamped belief consistency, the accuracy of conflict resolution, rejection of stale data, calibration of uncertainty, post-suppression thermal reduction, re-tasking latency, and mission success under communication dropout. These metrics directly address the gap between a connected prototype and an operational decision-making architecture.

Energy–computation–-communication–thermal coupling

The architecture must be read as a coupled degradation problem. Increasing onboard computation power can improve perception and replanning, but it increases heat generation and battery draw. Offloading computation can reduce onboard heat, but it creates dependence on links that are likely to fail during large fires. Moving closer to the flame may improve the accuracy of suppression, but it increases radiative/convective heat exposure and shortens the safe dwell time. These constraints interact: Link loss increases onboard computation demand, higher computation increases the thermal load, thermal derating reduces the frame rate and planners' update frequency, and reduced update frequency can increase the navigation and suppression risk.

Consequently, future evaluations should report not only the accuracy, latency, and endurance, but also the operating envelope in which those metrics remain valid. Minimum reporting should include the communication state, onboard workload, platform mode, standoff zone, dwell time, thermal margin, battery reserve, and the fallback behavior used when any of these variables violates a safety threshold. Table 8 summarizes the degradation triggers and the required architectural responses.

From automation to mission-level autonomy

The previous sections show that wildfire systems increasingly automate sensing, localization, path planning, and control. However, this form of automation mainly concerns how a task is executed. Operational autonomy also requires a higher-level decision layer that decides what should be attempted, which objective should be prioritized, and when the mission must be modified or aborted. The ALFUS terminology similarly treats autonomy as a multidimensional property involving the mission's complexity, environmental difficulty, and requirements for human interaction rather than a single navigation capability[12].

For wildfire suppression, this distinction becomes critical when resources are constrained. If a single UAV detects several ignition points but has the payload and battery to treat only one, the decision cannot be reduced to selecting the nearest target. A mission priority policy should combine the expected risk reduction, the probability of suppression success, and the opportunity cost while penalizing platform exposure. In conceptual form, priority can be expressed as:

$$\begin{aligned}\text{Priority score} &= \mathrm{E}[\text{risk reduction}] \times \mathrm{P}(\text{success}) \\ &\quad - \text{platform risk} - \text{opportunity cost}\end{aligned} \tag{9}$$

The terms in this expression should be interpreted operationally rather than as a universal scalar formula. Risk reduction may represent protection of human life, evacuation corridors, houses, critical

Table 7. Integration failure modes and countermeasures.

Integration challenge	Example failure	Operational consequence	Architectural countermeasure	Ref.
Multirate observations	Satellite evidence is delayed when UAV/ground detections are near real-time	Stale or duplicated incident state	Timestamped buffers, observation age-based weighting, multirate filters	[7,65,70,102,104,105]
Missing data	UAV loses link or thermal camera is saturated	Planner acts on incomplete evidence	Uncertainty expansion, re-observation requests, fall back on local autonomy	[7,9–11,89]
Sensor conflict	Thermal positive, RGB negative, or LiDAR–GNSS disagreement	False alarm, missed fire, or wrong target coordinates	Dynamic reliability scoring and arbitration rules	[7,78,89,105]
GNSS drift	Canopy or terrain multipath shifts the fire's coordinates	The suppression drop misses the target	LIO/VIO correction, uncertainty ellipses, target confirmation pass	[7,45,69,78]
Communication dropout	Edge and cloud maps diverge	Inconsistent common operating picture	Local decision authority with delayed synchronization	[7,8,65,89]
No suppression feedback	Suppressant drop fails or reignition begins	Open-loop mission termination despite an active risk	Post-action thermal/RGB reassessment and feedback-driven replanning	[106,114–116]
Overconfident AI	High mAP model fails under smoke, glare, or hot background clutter	Unsafe autonomous action	Uncertainty-aware inference and hard-negative validation	[3,6,7,88]

infrastructure, or high-value ecological areas; the probability of success depends on the distance, payload, wind, terrain, the fire's intensity, and sensors' confidence; and platform risk includes thermal exposure, turbulence, communication status, battery reserves, and navigation uncertainty. Thus, a geometrically short RRT* or A* path is not necessarily an acceptable suppression path if it violates the standoff distance, dwell time, battery reserve, or thermal health constraints[9−11]. In safety-critical deployments, minimum return energy, maximum thermal exposure, controllable flight, and minimum localization confidence should be treated as hard constraints that can override any mission benefit score.

True autonomy also requires governance of failure states. Replanning is sufficient only when the platform remains controllable and sensors' uncertainty is bounded. Recoverable failures, such as temporary visual degradation or a single stale data stream, can be handled through re-observation, sensor fallback, or local replanning. Nonrecoverable or safety-critical failures, such as rotor degradation, battery overheating, thermal overload, total navigation loss, or uncontrollable flight dynamics, require mandatory emergency behaviors rather than continued mission optimization. The autonomy state of a machine should therefore include nominal operation, degraded operation, isolated operation, conservative mission continuation, abort/retreat, return to launch, safe landing, plume exit, and human escalation where communication is available. Table 9 summarizes the decision functions and fallback behaviors needed for mission-level autonomy.

Human–robot teaming and supervised autonomy in wildfire response

Autonomy in wildfire robotics should be framed as supervised human–robot teaming rather than as a replacement for incident commanders, firefighters, safety officers, or fleet operators. Foundational human–automation work distinguishes information acquisition, information analysis, decision selection, and action implementation, meaning that a wildfire system may automate sensing and local execution while preserving human authority over the mission's intent, approval, and high-consequence trade-offs[15−18]. This distinction is important because wildfire priorities involve tactical and ethical judgments, including whether to protect structures, evacuation routes, crews, infrastructure, ecological assets, or containment lines under limited time and resources.

For multi-UAV/UGV wildfire fleets, the operator should not be required to micromanage each platform. Human–swarm and mixed-initiative research instead supports supervisory interfaces in which humans issue region-, task-, or objective-level commands such as

Table 8. Degradation modes and architectural responses.

Degradation trigger	System-level risk	Required architectural response	Ref.
Cellular/5G/cloud loss	Global coordination and model offloading unavailable	Switch to local mesh or isolated onboard replanning; store-and-forward synchronization	[7,8,65,89]
Edge gateway loss	Swarm desynchronization and inconsistent fire maps	Exchange compact belief-state summaries among nearby UAV/UGV nodes	[53,89,105,116]
GNSS drift plus link loss	Incorrect fire coordinates and unsafe approach	Downweight GNSS; rely on LIO/VIO where available; increase uncertainty and retreat if localization confidence falls	[7,45,78]
Rising GPU/SoC temperature	Inference latency increases through throttling	Reduce the model's size, frame rate, and noncritical analytics; prioritize safety supervisor	[10,91,96]
External flame radiation/plume heat	Rapid reduction in the thermal margin	Enforce the standoff zone, dwell time, post-drop exit path, and thermal abort threshold	[9,11]
Battery or enclosure overheating	Power loss or hardware damage	Abort suppression, preserve reserve energy, and transfer the task to another platform if available	[9−11]

Table 9. Autonomy decision functions and fallback behavior.

Decision layer	Question answered	Wildfire example	Governance or emergency response	Ref.
Autonomy definition and authority boundary	What level of decision authority is allowed?	Differentiate teleoperation, supervised autonomy, and independent mission decisions.	Specify the human interaction threshold, the permitted action set, and when escalation is mandatory.	[12,14]
Execution automation	How is a command executed?	Follow waypoints, stabilize hover, or perform a drop command.	Low-level controllers execute tasks only within verified health and safety limits; an actuator anomaly triggers abort logic.	[58,117,118]
Perception and localization autonomy	What is observed and where is it?	Detect fire/smoke, estimate hotspots' location, or navigate under GNSS loss.	Use confidence thresholds, re-observation, sensor fallback, and uncertainty expansion when evidence is weak or conflicting.	[4,72,88,117]
Mission-priority arbitration	What should be done first?	Select one ignition point among several when the battery or suppressant payload is insufficient.	Rank tasks by expected risk reduction, feasibility of suppression, opportunity cost, and the availability of other assets.	[13,106,115,116]
Risk governor and platform survival	Should the task continue?	Reject a short path that crosses a high-temperature plume or leaves no return reserve.	Apply hard constraints on standoff, dwell time, thermal margin, battery reserves, turbulence, and link state.	[9−11]
Degraded and isolated autonomy	What if cloud or mesh support is lost?	Continue local monitoring or suppress only if the onboard belief and return reserve remain valid.	Maintain an onboard/local belief map, perform conservative replanning, store data for later synchronization, or retreat.	[8,10,60,119]
Emergency and fail-safe autonomy	What if severe failure occurs?	Rotor faults, smoke-blinded vision, battery overheating, or total navigation loss.	Execute a return to launch, controlled safe landing, plume exit, mission abort, or human handoff instead of ordinary replanning.	[14,117,118,120]

https://doi.org/10.48130/emst-0026-0009

'map the eastern fire edge,' 'verify this hotspot,' or 'prioritize structures north of the fireline,' wheres robots perform local navigation, perception, and coordination. To reduce cognitive overload, the interface should aggregate raw detection instances into incident-level summaries, rank alerts by risk and uncertainty, indicate platforms' health and autonomy state, and use exception-based supervision so that human attention is requested primarily when confidence drops, sensors' evidence conflicts, communication degrades, or safety boundaries are approached[17,19,20,23,24].

Explainability is essential for calibrating trust and safe intervention. The operator should be able to inspect why a scene was classified as flame, smoke, a hotspot, or a false alarm; why one ignition point was prioritized over another; why a route was rejected as thermally unsafe; why an autonomy mode changed; and/or why the system requested approval or aborted. Explanations should therefore combine sensor evidence, confidence, uncertainty, risk drivers, platform constraints, and alternative actions, rather than presenting only a neural network confidence score[21,22].

Human intervention should also be granular rather than binary. Direct teleoperation may be necessary for degraded sensing or close-proximity rescue operations, waypoint control may be appropriate for local inspection, task-level commands can trigger mapping or verification missions, objective-level commands can encode incident priorities, approval-based autonomy can gate release of the suppressant or high-risk approaches, and exception-based supervision can allow routine fleet actions to continue without continuous human input. Authority arbitration must be explicit: Hard safety constraints such as firefighter exclusion zones, no-fly areas, thermal standoff limits, loss of controllable flight, or insufficient return energy should override both autonomous optimization and unsafe human commands; within those constraints, human mission intent should guide priority selection and AI-based tactical optimization. Any override or refusal should be logged and explained for accountability[15,20,21,23,24].

Identified gaps and limitations in the current literature

Despite technological maturity in individual components, six critical gaps impede operational deployment.

A seventh widespread gap is the lack of quantified degradation envelopes for communication-denied and thermally constrained operation. Many studies report detectors' performance, flight time, or cloud-connected workflows without specifying whether the same system can maintain safe decision-making after losing links or during short-term exposure to near-fire radiant and convective heat.

Consequently, the remaining gaps are not only empirical but architectural: the field lacks common methods for proving that connected wildfire technologies can maintain a coherent belief state and make safe decisions when evidence is delayed, contradictory, or incomplete.

Lack of standardized benchmarks. Wildfire detection datasets vary in size (500–10,000 images), environmental conditions (day/night, smoke density), and annotation quality (bounding boxes vs. pixel masks). No consensus benchmark analogous to ImageNet or COCO exists, preventing objective cross-system comparisons. Proposals for standardized test protocols should specify: (i) the environmental parameter ranges (smoke's optical depth, illumination, wind speed), (ii) the measurement methodology of detection latency, and (iii) quantification of false positives under adverse conditions.

Methodological bias and evidence-level ambiguity. A second limitation is that the reported performance is often detached from the realism of the validation environment. Laboratory flames, web-scraped images, and controlled outdoor burns are useful for early algorithm development, but they cannot reproduce the coupled radiative, aerodynamic, optical, and operational stressors of an uncontrolled wildfire. Many studies also underreport whether the test set is geographically independent, whether hard negatives are included, whether sequential video frames leak across the training and testing partitions, and whether platform-related effects such as vibration, motion blur, auto-exposure changes, and communication loss are included. Therefore, published mAP, F1-score, or latency values should be interpreted according to the level of evidence rather than compared as if all studies validated the same operational problem. Table 10 summarizes the reliability levels used to judge the strength of evidence across reviewed studies.

Insufficient end-to-end validation. Most papers evaluate isolated components (e.g., sensors' performance or algorithmic accuracy) without characterizing full-system integration. For example, thermal camera specifications are reported independently of the accuracy of coordinate transformation, yet geolocation errors (50–200 cm at high altitude) may exceed the requirements for targeting fire suppression. Future work should adopt system-level metrics, such as time from ignition to alert dispatch, false alarm rates including sensor fusion, and the mission success rate under realistic smoke/wind conditions.

Limited reporting of operational reliability. Field deployments encounter environmental stressors that are absent in laboratory testing: Thermal extremes (> 50 °C near active flames), vibration during UAV flight, contamination of the optics by dust, and electromagnetic interference from RF communication. Only a minority of cited field-oriented studies report extended field trials (> 10 flight hours or > 30 d of deployment); most present proof-of-concept validations. Operational maturity requires chracterization of the failure

Table 10. Methodological reliability framework.

Level	Validation setting	Typical strength	Main limitation	Operational confidence	Ref.
1	Synthetic or web-scraped images	Low-cost large-scale pretraining and rapid benchmarking	No real sensor physics; severe domain gaps; weak negative coverage	Very low	[3,7,41,94,121]
2	Laboratory flame/smoke experiments	Repeatable conditions and controlled labels	Limited radiative flux, smoke dynamics, wind, clutter, and platform motion	Low	[3,7,78]
3	Controlled outdoor or prescribed burns	More realistic illumination, wind, and smoke than laboratory scenes	Scale and chaos remain constrained; safety protocols simplify the scene	Medium	[7,68,88,106]
4	Short-duration UAV/UGV field trials near fire lines	Captures platform vibration, latency, geolocation, and partial smoke effects	Often short, site-specific, and not statistically representative	Medium–high	[7,53,68,89]
5	Multiseason operational wildfire deployment	Includes uncontrolled smoke, heat, wind, terrain, communications, and maintenance effects	Expensive, hazardous, and difficult to standardize	High	[7–11,68]

modes, including sensors' degradation over time, false alarm rates under seasonal variations, lens contamination, thermal throttling, communication dropouts, and maintenance intervals.

Absence of safety-critical certification frameworks. Autonomous UAV operations near active fires often require mission profiles outside the baseline operating limits of small UASs, including operations beyond the visual line of sight, operations near people or infrastructure, degraded visibility, multivehicle coordination, and emergency contingency behavior. Rather than prescribing a universal flight-hour threshold, validation should be tied to the applicable operational risk category and civil aviation pathway. In the United States, FAA Part 107 operations that cannot comply with the listed limitations require an operational waiver demonstrating that equivalent safety is maintained[122]. In Europe, EASA's Specific Operations Risk Assessment (SORA) provides a risk-based method for classifying the ground and air risk of drone operations in a specific category, calculating the Specific Assurance and Integrity Level (SAIL), and identifying operational safety objectives and mitigations[123]. ISO 21384-3:2023 further specifies the requirements for safe commercial UAS operations, including safety-critical command-and-control link services[14]. ASTM F3322 should be cited only in the context of deployable parachute recovery systems intended to reduce impact energy when a small UAS cannot sustain normal stable flight, not as a telemetry logging standard[124].

Inadequate multimodal sensor fusion. Although thermal–RGB fusion is common (cited in 38 papers), higher-order integration, combining satellite-based hotspots, UAVs' thermal imagery, ground sensor networks, and meteorological data, remains exploratory. Fusion architectures should weight the inputs by reliability: Satellite-based detection instances are delayed by 15–30 min but cover vast areas, UAVs provide rapid local verification within a 5-km radius, and ground sensors offer sub-minute alerts but high false positives. Bayesian fusion frameworks (cited in nine papers) show promise but require validation at operational scale. The proposed architectural interpretation therefore reframes fusion as a governed decision process: Asynchronous observations are first temporally aligned, then weighted by the sources' reliability and environmental suitability, and finally incorporated into a closed-loop belief state that is updated after suppression actions.

Underdeveloped human–robot teaming and supervised autonomy. Existing wildfire robotics studies often emphasize autonomous detection, navigation, and planning, but give less attention to the operational mode through which human responders supervise, trust, override, or delegate tasks to multirobot systems. Future architectures should specify how operator workload is managed, explanation interfaces, the granularity of intervention, and authority arbitration, especially when autonomous recommendations conflict with incident command priorities or hard safety constraints[15–24].

Recommendations for future research directions

On the basis of the identified gaps, we propose six strategic research priorities.

1. Establish community benchmarks. Develop open-access wildfire detection datasets with (i) 10,000+ annotated images spanning day/night, clear/smoky conditions, (ii) the corresponding ground truth of fire locations (GPS coordinates, perimeter polygons), (iii) metadata documenting environmental parameters (temperature, humidity, wind, smoke optical depth, solar irradiance, camera altitude, and sensor type), (iv) mandatory hard-negative subsets including nonfire hot surfaces and smokelike distractors, and (v) standardized evaluation scripts for computing the mAP, IoU, latency, false positive rate, and stratified performance across evidence levels. We could host annual challenges analogous to COCO or RoboCup to incentivize improvements in performance and reproducibility.

2. Advance system-level integration and multiplatform coordination. Future work should transition from component-level validation to end-to-end system-based evaluation. Critical research priorities include characterizing achievable ignition-to-alert latency under realistic network delays and edge processing constraints, evaluating multivehicle coordination and the scalability of cooperative trajectory optimization to fleets of 20–50 UAVs operating under limited bandwidth conditions[61], and developing sensor fusion architectures that explicitly quantify false alarms, missed detections, and uncertainty trade-offs under realistic environmental stressors. Simulated environments incorporating realistic fire spread models, atmospheric turbulence, and communication constraints would enable rapid iteration before costly field trials. In addition, future studies should report the actual integration logic used to manage asynchronous data streams, missing observations, and sensor conflicts. Benchmark scenarios should include delayed satellite products, intermittent UAV links, conflicting thermal and RGB observations, and staged communication outages so that the fusion logic, uncertainty handling, and fallback behaviors are evaluated as system-wide properties rather than as isolated algorithms.

This priority should also include communication-denied trials in which cloud access, edge gateways, GNSS quality, or mesh connectivity are deliberately degraded. Success should be measured by whether the platform maintains a conservative local belief state, replans safely, preserves its thermal and battery margins, and synchronizes delayed data after the link returns.

3. Develop safety-critical certification pathways. Collaboration among research institutions, regulatory agencies, firefighting organizations, and platform manufacturers should establish risk-based validation pathways for autonomous wildfire responses. Technical milestones should be derived from the applicable operational risk category rather than from universal thresholds such as a fixed number of flight hours or a single failure rate target. A standards-aligned pathway should include operation-specific risk assessments and waivers or authorization logic for operations outside baseline limits for small UASs, including beyond visual line of sight (BVLOS) and multivehicle operation. It should also document command-and-control links' performance, containment,behavior with lost links, emergency procedures, operators' responsibilities, event logging, detect-and-avoid evidence, airspace deconfliction procedures, and fail-safe recovery or safe termination mechanisms. Validation should progress from simulation and hardware-in-the-loop testing to controlled field trials, supervised BVLOS demonstrations, and limited operational deployment under incident command oversight[118,122–124].

4. Explore transfer learning and few-shot adaptation. Current models trained on specific wildfire datasets may generalize poorly to novel fire types (grassland vs. forest), geographic regions (Mediterranean vs. boreal), or seasonal conditions (wet vs. dry seasons). Transferring learning from large-scale pre-trained models (CLIP, DINOv2) combined with few-shot fine-tuning on local data could enable rapid deployment-based adaptation. Federated learning across multiple wildfire agencies would allow collaborative model improvement without sharing sensitive geolocation data.

5. Quantify degradation-aware engineering envelopes. Future work should publish distance–time–computation and link loss

https://doi.org/10.48130/emst-0026-0009

envelopes for representative UAV/UGV platforms. The required variables include the standoff zone, dwell time, external radiant/convective exposure, onboard workload, SoC temperature, battery temperature, cooling strategy, reserve energy, and fallback behavior. These envelopes should become part of benchmark reporting because a model that is accurate in a lab but cannot run safely near a fire front is not operationally deployable.

6. Formalize mission-level and human-supervised autonomy. Future systems should explicitly encode mission priority arbitration, risk-governed path acceptance, platform survival constraints, emergency fallback behaviors, and supervised human–robot teaming. Evaluations should report not only the detection accuracy or path length but also whether the system selects the correct task under resource scarcity, rejects thermally unsafe trajectories, preserves return reserves, explains its recommendation to the operator, supports objective-level rather than joystick-level supervision, and executes a safe landing or return to launch under severe failure conditions[12−24].

Emerging dialog-driven and large language model-guided fire response concepts may assist operators in summarizing the state, proposing plans, or explaining autonomy decisions, but they should be treated as advisory decision-support layers until validated under safety-critical, communication-limited, and thermally degraded wildfire conditions[125].

Broader implications beyond wildfire applications

The system architectures and methodologies synthesized in this review extend to adjacent domains requiring autonomous environmental monitoring and intervention, as described below.

Agricultural precision intervention. Thermal–RGB sensor fusion for fire detection translates directly to crop stress monitoring via NDVI analysis and thermal anomaly detection. Autonomous UAV coordination strategies enable targeted pesticide applications in orchards that are inaccessible to ground machinery. Cost projections of, for example, a 20-vehicle fleet ($\sim\$300,000$ capital + $\$50,000$ annual operations) are competitive with manned agricultural aviation ($\$150$–250 per hour of flight time).

Search and rescue in hazardous environments. Hotspot localization algorithms can generalize to detecting the thermal signature of trapped individuals in collapsed structures or avalanche debris. Multiagent coordination protocols developed for fire suppression can apply to distributed search patterns, maximizing the probability of detection within time-critical windows (e.g., the golden hour for trauma patients).

Industrial inspection and leak detection. Thermal imaging for fire detection can be repurposed for identifying gas leaks in refineries and chemical plants (methane leaks exhibit characteristic thermal signatures). Autonomous platforms navigating a wildfire's terrain can adapt to confined industrial spaces, such as pipe galleries and cooling tower interiors, that are inaccessible to human inspectors. The economic case showes that replacing scaffolding-based inspection would cost $\$10,000$–50,000 per event, with UAV-UGV systems costing $\$2,000$ per inspection after capital amortization.

Conclusions

Wildfire detection and suppression increasingly rely on integrated electronic systems spanning wireless sensor networks, thermal imaging, deep learning inference, autonomous platforms, and multisource data fusion. This structured review synthesizes the cited peer-reviewed and technical literature using an engineering taxonomy covering sensing hardware, computer vision algorithms, geospatial technologies, decision support architectures, and autonomous coordination strategies. The reference list is intentionally limited to works that are cited in the manuscript to support specific claims.

Key findings reveal technological maturity in isolated components—thermal cameras detect fires at multikilometer ranges, YOLO-based models achieve 85%–90% detection accuracy with an inference latency below 50 ms on edge devices, and UAV platforms provide rapid wide-area coverage—yet system-level integration remains fragmented. Recurring engineering trade-offs characterize the current solutions, such as accuracy vs. inference speed, range vs. resolution (long-distance sensing limits the pixel density), and endurance vs. payload capacity (extended flight time constrains sensors' capabilities).

Critical gaps impeding operational deployment include the absence of standardized benchmarks, which prevents objective system comparison, and insufficient end-to-end validation under realistic environmental conditions. Additional barriers include the limited reporting of long-term reliability, calibration drift, maintenance requirements, and the lack of certification frameworks for safety-critical autonomous operations. Bridging these gaps requires coordinated efforts among fire management agencies, robotics researchers, remote sensing specialists, and regulatory bodies.

The analysis also highlights that edge–cloud coordination must be replaced by an edge-first, degradation-aware view of autonomy. Communication links should be considered to be intermittent by design, and thermal limits should be modeled as a combined effect of onboard heat generation, external radiant heat, convective plume exposure, cooling degradation, and the mission's dwell time rather than as ambient temperature-based throttling alone[9−11].

The review also reframes autonomy as supervised human–robot teaming rather than human replacement. In operational wildfire responses, humans should retain mission intent, approval authority, override capability, and accountability for high-consequence decisions, whereas autonomous platforms reduce exposure and workload through persistent sensing, local planning, explainable recommendations, and exception-based reporting[15−24].

Future research should prioritize standardized evaluation protocols that enable reproducible cross-system comparisons under explicit environmental stressors, together with end-to-end integration demonstrating complete detection-to-suppression pipelines. Safety validation should be anchored to risk-based operational frameworks for UASs, with documented failure modes, contingency procedures, command-and-control link behavior, and operational evidence accumulated across simulation, controlled field trials, and supervised deployments. Future work should also improve transfer learning approaches that enable rapid adaptation to novel fire types and geographic regions. The methodologies and architectures synthesized in this review extend beyond wildfire applications to precision agriculture, search and rescue operations, and industrial inspection; indeed, they could extend to any domain requiring autonomous environmental monitoring and time-critical intervention in hazardous conditions. Future validation should further quantify whether suppression-oriented autonomy can maintain safe perception, localization, replanning, and retreat capability under communication loss, thermal exposure, and degraded sensing rather than reporting isolated components' accuracy alone.

In practical terms, future studies should report whether a UAV/UGV can continue tasks such as safe perception, localization, replanning, and retreat under complete communication loss while maintaining thermal and energy margins. This requirement connects physical failure mechanisms, closed-loop architecture, and realistic engineering quantification into a single operational-readiness criterion.

By providing an engineering-focused capability map, this review equips researchers with reference architectures, practitioners with technology selection criteria, and funding agencies with research priorities. The transition from isolated technical demonstrations to operationally mature wildfire response systems demands interdisciplinary collaboration spanning electrical engineering, computer science, atmospheric science, and firefighting operations—a challenge this synthesis aims to facilitate through structured knowledge organization and gap identification. The autonomy framing therefore treats wildfire robots not merely as automated perception–planning–control pipelines, but as mission-governed systems that must arbitrate priorities, respect survival constraints, and transition to fail-safe modes when the operating envelope collapses.

Ethical statements

Not applicable.

Author contributions

The authors confirm their contributions to the paper as follows: conception and design of the review, literature collection, analysis and interpretation of the cited studies, manuscript preparation, and critical revision: Acar O, Baydemir P. Both authors reviewed and approved the final version of the manuscript.

Data availability

Data sharing is not applicable to this article, as no datasets were generated or analyzed during the current study. The literature sources analyzed for the manuscript are listed in the References.

Acknowledgments

The authors thank Sempa Pompa A.Ş. for providing financial support to create this paper.

Conflict of interest

The authors declare that they have no conflict of interest.

Dates

Received 13 January 2026; Revised 30 June 2026; Accepted 7 July 2026; Published online 4 August 2026

References

[1] Bugarić M, Krstinić D, Šerić L, Stipaničev D. 2025. Current trends in wildfire detection, monitoring and surveillance. *Fire* 8(9):356

[2] Boroujeni SPH, Razi A, Khoshdel S, Afghah F, Coen JL, et al. 2024. A comprehensive survey of research towards AI-enabled unmanned aerial systems in pre-, active-, and post-wildfire management. *Information Fusion* 108:102369

[3] Bouguettaya A, Zarzour H, Taberkit AM, Kechida A. 2022. A review on early wildfire detection from unmanned aerial vehicles using deep learning-based computer vision algorithms. *Signal Processing* 190:108309

[4] Ramos L, Casas E, Bendek E, Romero C, Rivas-Echeverría F. 2024. Computer vision for wildfire detection: a critical brief review. *Multimedia Tools and Applications* 83(35):83427−83470

[5] Tasnim S, Singh A, Hasan T, Ahmed I. 2024. UAV-based efficient and reliable wildfire detection on the edge. In *2024 IEEE 10th World Forum on Internet of Things (WF-IoT). November 10–13, 2024, Ottawa, ON, Canada*. IEEE. pp. 511−516 doi: 10.1109/WF-IoT62078.2024.10811358

[6] Al-Smadi Y, Alauthman M, Al-Qerem A, Aldweesh A, Quaddoura R, et al. 2023. Early wildfire smoke detection using different YOLO models. *Machines* 11(2):246

[7] Patel J, Bhusnoor M, Patel D, Mehta A, Sainkar S, et al. 2023. Unmanned aerial vehicle-based forest fire detection systems: a comprehensive review. *SSRN Electronic Journal*

[8] Ramadan MNA, Basmaji T, Gad A, Hamdan H, Akgün BT, et al. 2024. Towards early forest fire detection and prevention using AI-powered drones and the IoT. *Internet of Things* 27:101248

[9] Häusermann D, Bodry S, Wiesemüller F, Miriyev A, Siegrist S, et al. 2023. FireDrone: multi-environment thermally agnostic aerial robot. *Advanced Intelligent Systems* 5(9):2300101

[10] Coloe BT, Atkins EM. 2016. Cyber-physical optimization of small unmanned aircraft systems with thermal consideration. In *2016 American Control Conference (ACC). July 6–8, 2016, Boston, MA, USA*. IEEE. pp. 1373−1380 doi: 10.1109/ACC.2016.7525109

[11] Hilton JE, Leonard J, Blanchi R, Newnham G, Opie K, et al. 2017. Dynamic modelling of radiant heat from wildfires. In *International Congress on Modelling and Simulation, December 3–8, 2017, Hobart, Tasmania, Australia*. Modelling and Simulation Society of Australia and New Zealand (MSSANZ), Inc. pp. 1104−1110 doi: 10.36334/modsim.2017.h10.hilton

[12] Huang HM, Messina E, Albus J, ALFUS Working Group. 2004. Autonomy levels for unmanned systems (ALFUS) framework, volume I: terminology. *NIST Special Publication 1011*. National Institute of Standards and Technology, Gaithersburg, Maryland, USA. https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication1011-I-2.0.pdf

[13] Gerkey BP, Matarić MJ. 2004. A formal analysis and taxonomy of task allocation in multi-robot systems. *The International Journal of Robotics Research* 23(9):939−954

[14] International Organization for Standardization (ISO). 2023. Unmanned aircraft systems – Part 3: Operational procedures. *ISO 21384-3:2003*, ISO, Geneva, Switzerland. www.iso.org/files/live/sites/isoorg/files/news/magazine/ISOupdate/EN/2023/ISOupdate-August-2023.pdf

[15] Parasuraman R, Sheridan TB, Wickens CD. 2000. A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans* 30(3):286−297

[16] Endsley MR, Kaber DB. 1999. Level of automation effects on performance, situation awareness and workload in a dynamic control task. *Ergonomics* 42(3):462−492

[17] Kaber DB, Endsley MR. 2004. The effects of level of automation and adaptive automation on human performance, situation awareness and workload in a dynamic control task. *Theoretical Issues in Ergonomics Science* 5(2):113−153

[18] Goodrich MA, Schultz AC. 2008. Human–robot interaction: a survey. *Foundations and Trends® in Human–Computer Interaction* 1(3):203−275

[19] Kolling A, Walker P, Chakraborty N, Sycara K, Lewis M. 2016. Human interaction with robot swarms: a survey. *IEEE Transactions on Human-Machine Systems* 46(1):9−26

[20] Jiang S, Arkin RC. 2016. Mixed-initiative human-robot interaction: definition, taxonomy, and survey. In *2015 IEEE International*

Conference on Systems, Man, and Cybernetics. October 9–12, 2015, Hong Kong, China. IEEE. pp. 954–961 doi: 10.1109/SMC.2015.174
[21] Methnani L, Chiou M, Dignum V, Theodorou A. 2024. Who's in charge here? A survey on trustworthy AI in variable autonomy robotic systems. *ACM Computing Surveys* 56(7):1–32
[22] Kong X, Xing Y, Tsourdos A, Wang Z, Guo W, et al. 2024. Explainable interface for human-autonomy teaming: a survey. *arXiv* 2405.02583
[23] Agrawal A, Cleland-Huang J, Steghofer JP. 2020. Model-driven requirements for humans-on-the-loop multi-UAV missions. In *2020 IEEE Tenth International Model-Driven Requirements Engineering (MoDRE). August 31, 2020. Zurich, Switzerland*. IEEE. pp. 1–10 doi: 10.1109/modre51215.2020.00007
[24] Frering L, Koefler A, Huber M, Pfister S, Feischl R, et al. 2024. Multi-robot support system for fighting wildfires in challenging environments: system design and field test report. In *2023 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR), November 13–15, 2023, Naraha, Fukushima, Japan*. IEEE. pp. 32–38 doi: 10.1109/SSRR59696.2023.10499951
[25] Xing J, Wang C, Liu Y, Chao Z, Guo J, et al. 2023. UAV multispectral imagery predicts dead fuel moisture content. *Forests* 14(9):1724
[26] Ghali R, Akhloufi MA, Mseddi WS. 2022. Deep learning and transformer approaches for UAV-based wildfire detection and segmentation. *Sensors* 22(5):1977
[27] Casas E, Ramos L, Bendek E, Rivas-Echeverria F. 2024. YOLOv5 vs. YOLOv8: performance benchmarking in wildfire and smoke detection scenarios. *Journal of Image and Graphics* 12(2):127–136
[28] Hristov G, Raychev J, Kinaneva D, Zahariev P. 2018. Emerging methods for early detection of forest fires using unmanned aerial vehicles and lorawan sensor networks. In *2018 28th EAEEIE Annual Conference (EAEEIE). September 26–28, 2018, Hafnarfjordur, Iceland*. IEEE. pp. 1–9 doi: 10.1109/eaeeie.2018.8534245
[29] Al-Fugara A, Mabdeh AN, Ahmadlou M, Pourghasemi HR, Al-Adamat R, et al. 2021. Wildland fire susceptibility mapping using support vector regression and adaptive neuro-fuzzy inference system-based whale optimization algorithm and simulated annealing. *ISPRS International Journal of Geo-Information* 10(6):382
[30] Lenz C, Quenzel J, Periyasamy AS, Razlaw J, Rochow A, et al. 2022. Autonomous wall-building and firefighting: team NimbRo's UGV solution for MBZIRC 2020. *Field Robotics* 2:55–74
[31] Wang Y, Wang Y, Xu C, Wang X, Zhang Y. 2024. Computer vision-driven forest wildfire and smoke recognition via IoT drone cameras. *Wireless Networks* 30(9):7603–7616
[32] Munawar HS, Gharineiat Z, Akram J, Imran Khan S. 2022. A framework for burnt area mapping and evacuation problem using aerial imagery analysis. *Fire* 5(4):122
[33] Hartley RJL, Davidson SJ, Watt MS, Massam PD, Aguilar-Arguello S, et al. 2022. A mixed methods approach for fuel characterisation in gorse (*Ulex europaeus* L.) scrub from high-density UAV laser scanning point clouds and semantic segmentation of UAV imagery. *Remote Sensing* 14(19):4775
[34] Di Paola D, Milella A, Cicirelli G, Distante A. 2010. An autonomous mobile robotic system for surveillance of indoor environments. *International Journal of Advanced Robotic Systems* 7:8
[35] Hussain SH, Hussain A, Shah R, Abro SA. 2019. Mini rover-object detecting ground vehicle (ugv). *University of Sindh Journal of Information and Communication Technology* 3(2):89–94
[36] Acar O, Honkavaara E, Sefaci S. 2025. Bimodal MorphoCopter for wildfire suppression: transforming quadcopters to all-terrain-vehicles (ATV). *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLVIII-2/W11-2025:1–6
[37] Reilly S, Clark ML, Bentley LP, Matley C, Piazza E, et al. 2021. The potential of multispectral imagery and 3D point clouds from unoccupied aerial systems (UAS) for monitoring forest structure and the impacts of wildfire in Mediterranean-climate forests. *Remote Sensing* 13(19):3810
[38] Pasini D, Jiang C, Jolly MP. 2023. UAV and UGV autonomous cooperation for wildfire hotspot surveillance. In *2022 IEEE MIT Undergraduate Research Technology Conference (URTC). September 30–October 2, 2022, Cambridge, MA, USA*. IEEE. pp. 1–5 doi: 10.1109/URTC56832.2022.10002208
[39] Dilfanian E, Zhang Y, Dong H, Qiao L, Taherzadeh A, et al. 2025. Autonomous vision-guided high-precision firefighting using unmanned aerial vehicles. In *IECON 2024 - 50th Annual Conference of the IEEE Industrial Electronics Society. November 3–6, 2024, Chicago, IL, USA*. IEEE. pp. 1–6 doi: 10.1109/IECON55916.2024.10905606
[40] Shahid M, Chen SF, Hsu YL, Chen YY, Chen YL, et al. 2023. Forest fire segmentation *via* temporal transformer from aerial images. *Forests* 14(3):563
[41] Wang M, Yue P, Jiang L, Yu D, Tuo T, et al. 2024. An open flame and smoke detection dataset for deep learning in remote sensing based fire detection. *Geo-spatial Information Science* 28:511–526
[42] Andrada ME, Ferreira JF, Portugal D, Couceiro MS. 2022. Integration of an artificial perception system for identification of live flammable material in forestry robotics. In *2022 IEEE/SICE International Symposium on System Integration (SII). January 9–12, 2022, Narvik, Norway*. IEEE. pp. 103–108 doi: 10.1109/SII52469.2022.9708734
[43] Raveendran R, Ariram S, Tikanmaki A, Roning J. 2020. Development of task-oriented ROS-based autonomous UGV with 3D object detection. In *2020 IEEE International Conference on Real-time Computing and Robotics (RCAR). September 28–29, 2020. Asahikawa, Japan*. IEEE. pp. 427–432 doi: 10.1109/RCAR49640.2020.9303034
[44] Stow D, Riggan P, Schag G, Brewer W, Tissell R, et al. 2019. Assessing uncertainty and demonstrating potential for estimating fire rate of spread at landscape scales based on time sequential airborne thermal infrared imaging. *International Journal of Remote Sensing* 40(13):4876–4897
[45] Luo Z, Zhang Y, Mu L, Huang J, Xin J, et al. 2021. A UAV path planning algorithm based on an improved D* lite algorithm for forest firefighting. In *2020 Chinese Automation Congress (CAC). November 6–8, 2020, Shanghai, China*. IEEE. pp. 4233–4237 doi: 10.1109/CAC51589.2020.9327111
[46] Vasić MK, Drak A, Bugarin N, Kružić S, Musić J, et al. 2020. Deep semantic image segmentation for UAV-UGV cooperative path planning: a car park use case. In *2020 International Conference on Software, Telecommunications and Computer Networks (SoftCOM). September 17–19, 2020, Split, Croatia*. IEEE. pp. 1–6 doi: 10.23919/SoftCOM50211.2020.9238313
[47] Ando H, Ambe Y, Ishii A, Konyo M, Tadakuma K, et al. 2018. Aerial hose type robot by water jet for fire fighting. *IEEE Robotics and Automation Letters* 3(2):1128–1135
[48] Lakov D. 1996. Fuzzy sets and fuzzy logic, theory and applications. *Fuzzy Sets and Systems* 84(1):114
[49] Sandhu AS, Singh RP, Prabha C. 2025. Edge AI and IoT-integrated wildfire real-time monitoring system. In *2025 6th International Conference on Electronics and Sustainable Communication Systems (ICESC). September 10–12, 2025, Coimbatore, India*. IEEE. pp. 639–644 doi: 10.1109/ICESC65114.2025.11212548
[50] Verma U. 2022. Recent trends and challenges in analysis of UAV aerial images for post-disaster scene understanding. In *IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium. July 17–22, 2022, Kuala Lumpur, Malaysia*. IEEE. pp. 4647–4649 doi: 10.1109/IGARSS46834.2022.9883265
[51] Ahmad A, Tahar KN, Udin WS, Hashim KA, Darwin N, et al. 2014. Digital aerial imagery of unmanned aerial vehicle for various applications. In *2013 IEEE International Conference on Control System, Computing and Engineering. November 29–December 1, 2013, Penang, Malaysia*. IEEE. pp. 535–540 doi: 10.1109/ICCSCE.2013.6720023
[52] Akhloufi MA, Couturier A, Castro NA. 2021. Unmanned aerial vehicles for wildland fires: sensing, perception, cooperation and assistance. *Drones* 5(1):15

[53] Pham HX, La HM, Feil-Seifer D, Deans MC. 2020. A distributed control framework of multiple unmanned aerial vehicles for dynamic wildfire tracking. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 50(4):1537–1548

[54] Mseddi WS, Ghali R, Jmal M, Attia R. 2021. Fire detection and segmentation using YOLOv5 and U-NET. In *2021 29th European Signal Processing Conference (EUSIPCO). August 23–27, 2021, Dublin, Ireland*. IEEE. pp. 741–745 doi: 10.23919/EUSIPCO54536.2021.9616026

[55] Chen J, Luo X, Hua C, Mu D, Sun F. 2024. Modeling and robust adaptive practical predefined time and precision tracking control of unmanned fire fighting robot. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 54(7):4273–4283

[56] Moyer W, Patil N, Toffoli E, Hadi N, Guan A, et al. 2022. Drought monitoring with drones a hundred fields at a time, In *Technical Memorandum NASA/TM-20220014853*, National Aeronautics and Space Administration (NASA), USA. https://ntrs.nasa.gov/citations/20220014853

[57] Badiger VC, Bhalerao SB, Mankar AA, Nimbalkar A. 2020. Wireless sensor network-assisted forest fire detection and control firefighting robot. *SAMRIDDHI: A Journal of Physical Sciences, Engineering and Technology* 12(S2):57–61

[58] Qin H, Cui JQ, Li J, Bi Y, Lan M, et al. 2016. Design and implementation of an unmanned aerial vehicle for autonomous firefighting missions. In *2016 12th IEEE International Conference on Control and Automation (ICCA). June 1–3, 2016, Kathmandu, Nepal*. IEEE. pp. 62–67 doi: 10.1109/ICCA.2016.7505253

[59] Jiao Z, Zhang Y, Mu L, Xin J, Jiao S, et al. 2020. A YOLOv3-based learning strategy for real-time UAV-based forest fire detection. In *2020 Chinese Control and Decision Conference (CCDC). August 22–24, 2020. Hefei, China*. IEEE. pp. 4963–4967 doi: 10.1109/ccdc49329.2020.9163816

[60] Erdelj M, Natalizio E, Chowdhury KR, Akyildiz IF. 2017. Help from the sky: leveraging UAVs for disaster management. *IEEE Pervasive Computing* 16(1):24–32

[61] Yao P, Wei X. 2022. Multi-UAV information fusion and cooperative trajectory optimization in target search. *IEEE Systems Journal* 16(3):4325–4333

[62] Chen Y, Liu Z, Xu Z, Lin J, Guan X, et al. 2025. UAVs-UGV cooperative boom sprayer system based on swarm control. *Computers and Electronics in Agriculture* 235:110339

[63] Chintalapati B, Precht A, Hanra S, Laufer R, Liwicki M, et al. 2025. Opportunities and challenges of on-board AI-based image recognition for small satellite Earth observation missions. *Advances in Space Research* 75(9):6734–6751

[64] Fulé PZ, Crouse JE, Roccaforte JP, Kalies EL. 2012. Do thinning and/or burning treatments in western USA ponderosa or Jeffrey pine-dominated forests help restore natural fire behavior? *Forest Ecology and Management* 269:68–81

[65] Valero MM, Rios O, Mata C, Pastor E, Planas E. 2017. An integrated approach for tactical monitoring and data-driven spread forecasting of wildfires. *Fire Safety Journal* 91:835–844

[66] Khakzad N. 2019. Modeling wildfire spread in wildland-industrial interfaces using dynamic Bayesian network. *Reliability Engineering & System Safety* 189:165–176

[67] Mohd Daud SMS, Mohd Yusof MYP, Heo CC, Khoo LS, Chainchel Singh MK, et al. 2022. Applications of drone in disaster management: a scoping review. *Science & Justice* 62(1):30–42

[68] Karma S, Zorba E, Pallis GC, Statheropoulos G, Balta I, et al. 2015. Use of unmanned vehicles in search and rescue operations in forest fires: advantages and limitations observed in a field trial. *International Journal of Disaster Risk Reduction* 13:307–312

[69] Ahn Y, Choi H, Kim BS. 2023. Development of early fire detection model for buildings using computer vision-based CCTV. *Journal of Building Engineering* 65:105647

[70] Schroeder W, Giglio L. 2017. Advances in the remote sensing of active fires. *Report*. GOFCGOLD Fire Implementation Team, College Park, MD, USA. www.naturalhazards.com.au/crc-collection/downloads/2.2.3_advances_in_the_remote_sensing_of_active_fires_report_v1.0.pdf

[71] Riggan PJ, Tissell RG. 2008. Chapter 6 Airborne remote sensing of wildland fires. *Developments in Environmental Science* 8:139–168

[72] Dimitrov K, Damyanov I. 2024. Thermal infrared imaging of near-freezing objects with unmanned aerial vehicles and the influence of observation height in real conditions. In *2024 32nd National Conference with International Participation (TELECOM). November 21–22, 2024, Sofia, Bulgaria*. IEEE. pp. 1–4 doi: 10.1109/TELECOM63374.2024.10812235

[73] Wörner J, Eimler J, Pein-Hackelbusch M. 2025. Long-term drift behavior in metal oxide gas sensor arrays: a one-year dataset from an electronic nose. *Scientific Data* 12:1628

[74] Yeo JD, Woo D. 2022. Real-time non-uniformity correction without TEC for microbolometer array. *Electronics* 11(19):3083

[75] Labed S, Touatİ H, Her İda A, Kerbab S, Saİrİ A . 2023. An AI-based image recognition system for early detection of forest and field fires. *European Journal of Forest Engineering* 9(2):48–56

[76] Wooster MJ, Roberts G, Smith AMS, Johnston J, Freeborn P, et al. 2013. Thermal remote sensing of active vegetation fires and biomass burning events. In *Thermal Infrared Remote Sensing. Remote Sensing and Digital Image Processing*, Vol. 17. eds. Kuenzer C, Dech S. Dordrecht: Springer. pp. 347–390 doi: 10.1007/978-94-007-6639-6_18

[77] Cansler CA, Kane VR, Hessburg PF, Kane JT, Jeronimo SMA, et al. 2022. Previous wildfires and management treatments moderate subsequent fire severity. *Forest Ecology and Management* 504:119764

[78] Ladstädter R, Schraml S, Studnicka N. 2024. Multi-sensor visibility study based on a hot smoke experiment in a tunnel. In *Urban Operations Expert Talks 2023 (#UOET23), June 27–29, 2023, Vienna, Austria*. pp. 1–21 doi: 10.5281/zenodo.10762118

[79] Buchelt A, Adrowitzer A, Kieseberg P, Gollob C, Nothdurft A, et al. 2024. Exploring artificial intelligence for applications of drones in forest ecology and management. *Forest Ecology and Management* 551:121530

[80] Martinez-Rozas S, Rey R, Alejo D, Acedo D, Cobano JA, et al. 2021. Skyeye Team at MBZIRC 2020: a team of aerial and ground robots for GPS-denied autonomous fire extinguishing in an urban building scenario. *arXiv* 2104.01834

[81] Khoshdel S, Luo Q, Afghah F. 2024. PyroTrack: belief-based deep reinforcement learning path planning for aerial wildfire monitoring in partially observable environments. *arXiv* 2403.11095

[82] Kong L, Li J, Guo S, Zhou X, Wu D. 2024. Computer vision based early fire-detection and firefighting mobile robots oriented for onsite construction. *Journal of Civil Engineering and Management* 30(8):720–737

[83] Kondylatos S, Prapas I, Ronco M, Papoutsis I, Camps-Valls G, et al. 2022. Wildfire danger prediction and understanding with deep learning. *Geophysical Research Letters* 49(17):e2022GL099368

[84] Condat R, Vasseur P, Allibert G. 2024. Focusing on object extremities for tree instance segmentation in forest environments. *IEEE Robotics and Automation Letters* 9(6):5480–5487

[85] Barzani AR, Pahlavani P, Ghorbanzadeh O. 2023. Ensembling of decision trees, knn, and logistic regression with soft-voting method for wildfire susceptibility mapping. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* X-4/W1-2022:647–652

[86] Spadavecchia C, Belcore E, Grasso N, Piras M. 2023. A fully automatic forest parameters extraction at single-tree level: a comparison of mls and tls applications. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLVIII-1/W1-2023:457–463

[87] McFadden AJ, Stow DA, Riggan PJ, Tissell R, O'Leary J, et al. 2024. Estimating fire radiative energy density with repeat-pass aerial thermal-infrared imaging of actively progressing wildfires. *Fire* 7(6):179

[88] Hopkins B, ONeill L, Marinaccio M, Habibpour M, Rowell E, et al. 2024. FLAME 3 dataset: unleashing the power of radiometric thermal UAV imagery for wildfire management. *arXiv* 2412.02831
[89] Lewicki T, Liu K. 2022. Multimodal wildfire surveillance with UAV. In *2021 IEEE Global Communications Conference (GLOBECOM). December 7–11, 2021, Madrid, Spain*. IEEE. pp. 1–6 doi: 10.1109/GLOBECOM46510.2021.9685547
[90] Briley AA, Afghah F. 2024. Hardware acceleration for real-time wildfire detection onboard drone networks. In *IEEE INFOCOM 2024 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS). May 20–20, 2024, Vancouver, BC, Canada*. IEEE. pp. 1–6 doi: 10.1109/INFOCOMWKSHPS61880.2024.10620877
[91] Hristov G, Kinaneva D, Georgiev G, Zahariev P, Kyuchukov P. 2020. An overview of the use of unmanned aerial vehicles for precision agriculture. In *2020 International Conference on Biomedical Innovations and Applications (BIA). September 24–27, 2020, Varna, Bulgaria*. IEEE. pp. 137–140 doi: 10.1109/BIA50171.2020.9244519
[92] Zasetska A, Urblík Ľ, Krupáš M, Kajáti E, Papcun P. 2025. Automated landing of a drone on a moving vehicle using computer vision. In *2025 IEEE 23rd World Symposium on Applied Machine Intelligence and Informatics (SAMI). January 23–25, 2025, Stará Lesná, Slovakia*. IEEE. pp. 385–388 doi: 10.1109/SAMI63904.2025.10883170
[93] Girshick R. 2016. Fast R-CNN. In *2015 IEEE International Conference on Computer Vision (ICCV). December 7–13, 2015, Santiago, Chile*. IEEE. pp. 1440–1448 doi: 10.1109/ICCV.2015.169
[94] Ghosh A, Wild T, Du J, Tan J, Grudnitsky A, et al. 2025. A unified future: integrated sensing and communication (ISAC) in 6G. *IEEE Journal of Selected Topics in Electromagnetics, Antennas and Propagation* 1(1):365–374
[95] Quan TM, Hildebrand DGC, Jeong WK. 2021. FusionNet: a deep fully residual convolutional neural network for image segmentation in connectomics. *Frontiers in Computer Science* 3:613981
[96] Tang Z, Liu X, Chen H, Hupy J, Yang B. 2020. Deep learning based wildfire event object detection from 4K aerial images acquired by UAS. *AI* 1(2):166–179
[97] Caron N, Noura HN, Nakache L, Guyeux C, Aynes B. 2025. AI for wildfire management: from prediction to detection, simulation, and impact analysis—bridging lab metrics and real-world validation. *AI* 6(10):253
[98] Yuan W, Qiao L, Tang L. 2024. Forest wildfire detection from images captured by drones using window transformer without shift. *Forests* 15(8):1337
[99] Xu Y, Wang H, Bi Y, Nie G, Wang X. 2025. MCDet: target-aware fusion for RGB-T fire detection. *Forests* 16(7):1088
[100] Yan Z, Wang L, Qin K, Zhou F, Ouyang J, et al. 2023. Unsupervised domain adaptation for forest fire recognition using transferable knowledge from public datasets. *Forests* 14(1):52
[101] Di Martino F, Pedrycz W, Sessa S. 2018. Spatiotemporal extended fuzzy C-means clustering algorithm for hotspots detection and prediction. *Fuzzy Sets and Systems* 340:109–126
[102] Chmaj G, Selvaraj H. 2015. Distributed processing applications for UAV/drones: a survey. In *Progress in Systems Engineering. Advances in Intelligent Systems and Computing*, Vol. 366, eds. Selvaraj H, Zydek D, Chmaj G. Cham: Springer. pp. 449–454 doi:10.1007/978-3-319-08422-0_66
[103] Ramasamy S, Mondal MS, Humann JD, Dotterweich JM, Reddinger JF, et al. 2024. Computationally efficient multi-agent optimization framework for online routing of UAV-UGV system. In *2024 IEEE 20th International Conference on Automation Science and Engineering (CASE). August 28–September 1, 2024, Bari, Italy*. IEEE. pp. 204–211 doi: 10.1109/CASE59546.2024.10711338
[104] Ravichandran A. 2024. Making satellite data mainstream: gaps, challenges, and opportunities for remote sensing experts [Industry Profiles and Activities]. *IEEE Geoscience and Remote Sensing Magazine* 12(4):373–376
[105] Kanand T, Kemper G, König R, Kemper H. 2020. Wildfire detection and disaster monitoring system using uas and sensor fusion technologies. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLIII-B3-2020:1671–1675
[106] Momeni M, Mirzapour Al-e-Hashem SMJ. 2024. Collaboration of thermal sensors and drones in fighting wildfires; Mathematical model and heuristic approach. *Internet of Things* 26:101168
[107] Guan M, Sagar K, Xu N. 2023. GPS-denied UAV-UGV relative positioning system via vision-UWB fusion. In *2023 IEEE 18th Conference on Industrial Electronics and Applications (ICIEA). August 18–22, 2023, Ningbo, China*. IEEE. pp. 227–232 doi: 10.1109/ICIEA58696.2023.10241577
[108] Dhrafani D, Liu Y, Jong A, Shin U, He Y, et al. 2025. FIReStereo: forest InfraRed stereo dataset for UAS depth perception in visually degraded environments. *IEEE Robotics and Automation Letters* 10(4):3302–3309
[109] Alamouri A, Gerke M, Batzdorfer S, Becker M, Bestmann U, et al. 2019. The joint research project ankommen–exploration using automated uav and ugv. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLII-2/W13:165–172
[110] Teskeredzic E, Akagic A. 2020. Low cost UGV platform for autonomous 2D navigation and map-building based on a single sensory input. In *2020 7th International Conference on Control, Decision and Information Technologies (CoDIT). June 29–July 2, 2020. Prague, Czech Republic*. IEEE. pp. 988–993 doi: 10.1109/CoDIT49905.2020.9263975
[111] Cai X, Tang W. 2025. Dynamic UAV path planning under wildfire scenarios using PPO algorithm. In *2025 4th International Conference on Robotics, Artificial Intelligence and Intelligent Control (RAIIC). July 4–6, 2025, Chengdu, China*. IEEE. pp. 196–199 doi: 10.1109/RAIIC65850.2025.11169796
[112] Alzorgan H, Razi A, Moshayedi AJ. 2023. Invited paper: actuator trajectory planning for UAVs with overhead manipulator using reinforcement learning. In *2023 IEEE 34th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC). September 5–8, 2023, Toronto, ON, Canada*. IEEE. pp. 1–6 doi: 10.1109/PIMRC56721.2023.10293891
[113] Dunagan S, Fladeland M, Ippolito C, Knudson M, Young Z. 2015. Mission adaptive uas capabilities for earth science and resource assessment. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XL-7/W3:1163–1170
[114] O'Neill L, Fulé PZ, Watts A, Moran C, Hopkins B, et al. 2024. Pixels to pyrometrics: UAS-derived infrared imagery to evaluate and monitor prescribed fire behaviour and effects. *International Journal of Wildland Fire* 33(11):WF24067
[115] Zhang C, Hu B, Yan F. 2021. Multiple UAVs Forest Fire Fighting using a hierarchical task planning method. In *2021 2nd International Conference on Electronics, Communications and Information Technology (CECIT). December 27–29, 2021. Sanya, China*. IEEE. pp. 760–765 doi: 10.1109/CECIT53797.2021.00138
[116] Ghamry KA, Kamel MA, Zhang Y. 2017. Multiple UAVs in forest fire fighting mission using particle swarm optimization. In *2017 International Conference on Unmanned Aircraft Systems (ICUAS). June 13–16, 2017, Miami, FL, USA*. IEEE. 1404–1409 doi: 10.1109/ICUAS.2017.7991527
[117] Martínez-Rozas S, Rey R, Alejo D, Acedo D, Cobano JA, et al. 2022. An aerial/ground robot team for autonomous firefighting in urban GNSS-denied scenarios. *Field Robotics* 2:241–273
[118] Real F, Castaño ÁR, Torres-González A, Capitán J, Sánchez-Cuevas PJ, et al. 2021. Autonomous fire-fighting with heterogeneous team of unmanned aerial vehicles. *Field Robotics* 1:1–28
[119] Liu K, Chauhan S, Devaraj R, Shahi S, Sreekumar U. 2019. Enabling autonomous unmanned aerial systems via edge computing. In *2019 IEEE International Conference on Service-Oriented System Engineering (SOSE). April 4–9, 2019. San Francisco East Bay, CA, USA*. IEEE. pp. 374–3745 doi: 10.1109/SOSE.2019.00063
[120] Walter V, Spurný V, Petrlík M, Báča T, Žaitlík D, et al. 2022. Extinguishing real fires by fully autonomous multirotor UAVs in the MBZIRC 2020 competition. *Field Robotics* 2:406–436

[121] Al-Kaff A, Moreno FM, José LJS, García F, Martín D, et al. 2017. VBII-UAV: vision-based infrastructure inspection-UAV. In *Recent Advances in Information Systems and Technologies*. eds. Rocha Á, Correia A, Adeli H, Reis L, Costanzo S. Cham: Springer International Publishing. pp. 221−231 doi: 10.1007/978-3-319-56538-5_24

[122] Federal Aviation Administration. 2026. *Part 107 Waivers*. FAA. Available: at www.faa.gov/uas/commercial_operators/part_107_waivers. Accessed May 7, 2026

[123] European Union Aviation Safety Agency. 2026. *Specific Operations Risk Assessment (SORA)*. EASA. Available at: www.easa.europa.eu/en/domains/drones-air-mobility/operating-drone/specific-category-civil-drones/specific-operations-risk-assessment-sora. Accessed: May 7, 2026

[124] ASTM International. 2024. *ASTM F3322-24a: Standard Specification for Small Unmanned Aircraft System (sUAS) Parachutes*. ASTM International. Available at: https://store.astm.org/standards/f3322. Accessed: May 7, 2026

[125] Majumdar S, Kirkley SE, Basu Mallik B. 2025. LLM-guided hybrid architecture for autonomous fire response: dialog-driven planning in space and disaster missions. In *2025 IEEE World AI IoT Congress (AIIoT). May 28–30, 2025, Seattle, WA, USA*. IEEE. pp. 1049−1054 doi: 10.1109/AIIoT65859.2025.11105239