NASA/TM–20260007959

A Multi-Tier Autonomous Aerial Architecture for Wildfire Detection, Characterization, and Communication in Infrastructure-Denied Environments

** All authors contributed equally to this memo. The names are listed alphabetically by last name.*

Alexei Barrow
Basis Flagstaff, Flagstaff, Arizona

Miriam Blinder
Lowell High School, San Franisco, California

Ian Gilliam
St. John’s College High School, Bowie, Maryland

Zubin Kalavade
Mountain View Highschool, Mountain View, California

Sathvik Majumdar
Metuchen High School, Metuchen, New Jersey

Esha Penmatsa
Green Level High School, Cary, North Carolina

August 2026

NASA STI Program Report Series

The NASA STI Program collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NTRS Registered and its public interface, the NASA Technical Reports Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA Programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at http://www.sti.nasa.gov

- Help desk contact information:

https://www.sti.nasa.gov/sti-contact-form/ and select the "General" help request type.

NASA/TM–20260007959

A Multi-Tier Autonomous Aerial Architecture for Wildfire Detection, Characterization, and Communication in Infrastructure-Denied Environments

** All authors contributed equally to this memo. The names are listed alphabetically by last name.*

Alexei Barrow
Basis Flagstaff, Flagstaff, Arizona

Miriam Blinder
Lowell High School, San Franisco, California

Ian Gilliam
St. John's College High School, Bowie, Maryland

Zubin Kalavade
Mountain View Highschool, Mountain View, California

Sathvik Majumdar
Metuchen High School, Metuchen, New Jersey

Esha Penmatsa
Green Level High School, Cary, North Carolina

National Aeronautics and
Space Administration

Ames Research Center
Mountain View, California

August 2026

Available from:

NASA STI Program / Mail Stop 050
NASA Langley Research Center
Hampton, VA 23681-2199

Acknowledgments

This technical memorandum was made possible by the NASA internship program. The authors of the report would like to thank Hokkwan Ng and Swati Saxena for their extensive guidance and advice throughout the research and writing process. In addition, the authors would like to thank everyone at Ames Research Center, primarily Shannon Zelinski and Divya Bhadoria, for making this experience possible.

Table of Contents

List of Figures

List of Tables

Abstract

Wildfire response depends on how fast an ignition can be confirmed and located, especially in remote regions where ground-based communication and monitoring may be limited. Geostationary sensors provide frequent observations but at kilometer-scale resolution, which is too coarse to resolve small fires in remote terrain. Ground camera networks require sightlines and infrastructure that back-country areas lack. To address these limitations, this work proposes a Multi-Tier Autonomous Wildfire Intelligence System that combines wide-area monitoring with targeted, high-resolution sensing. A solar-powered high-altitude long-endurance (HALE) platform operating at approximately 60,000 ft provides persistent wide-area thermal and optical surveillance, running onboard edge inference to screen candidate ignitions and reduce false positives and downlink bandwidth. When a candidate ignition is detected, low-altitude uncrewed aircraft systems (UAS) can be deployed to conduct localized observations, including high-resolution imaging and atmospheric measurements such as wind and plume observation. By combining persistent detection with local sensing, the proposed architecture is designed to provide first responders with timely, high-resolution information about fire location and behavior to aid in emergency decision making.

1. Introduction

Wildfire regimes across the western United States, the Mediterranean basin, and Australia have shifted toward larger, faster, and more destructive events, and an increasing share of them intersect the wildland– urban interface, the zone where housing meets wildland fuel. The January 2025 Palisades and Eaton fires in Los Angeles County burned roughly 37,500 acres between them yet destroyed more than 16,000 structures and killed 31 people [1]. This shows that fire behavior at the wildland–urban interface evolves on a timescale of minutes while the information available to responders' updates on a timescale of hours.

That mismatch arises because there exists no system that offers fine spatial detail and continuous temporal persistence over a fixed location.

- Geostationary imagers such as the Geostationary Operational Environmental Satellite R-Series (GOES-R) Advanced Baseline Imager (ABI) refresh the conterminous United States every five minutes and a mesoscale sector every minute, but the infrared channels used for fire detection have a footprint of roughly two kilometers at the sub-satellite point [2], [3]. A fire small enough to be contained by an initial-attack crew is frequently smaller than a single pixel.
- Polar-orbiting imagers resolve far more finely, the Visible Infrared Imaging Radiometer Suite (VIIRS) active fire product at 375 m [4], the Moderate Resolution Imaging Spectroradiometer (MODIS) at 1 km [5], but orbital mechanics limits any given cell to a small number of brief overpasses per day, and near-real-time distribution adds roughly three hours of latency [6].
- Terrestrial camera networks such as ALERT California now exceed 1,200 installations and provide continuous, machine-triaged coverage [7], [8]. They work well where they have sightlines, but they cannot see through a ridgeline, and an ignition at the bottom of a canyon stays invisible until its smoke column clears the terrain.
- Ground sensor networks placed in the fuel itself can detect the gases of a smoldering fire before any smoke is visible, which is the earliest warning available today, but each unit covers only its immediate surroundings, so protecting a landscape means installing very many of them in advance [9].

The Earth Fire Alliance and Muon Space FireSat program placed its first three operational satellites in orbit in July 2026, carrying a multispectral infrared payload specified to detect fires as small as five meters on a side; the present configuration delivers at least twice-daily revisit over fire-prone regions, with a stated target of hourly revisit by 2029 and roughly twenty-minute global revisit at a full constellation of more than fifty satellites [10], [11]. Even at the twenty-minute target, a purely spaceborne asset cannot loiter over one incident, and does not measure the wind field that drives spread and cannot carry the radio traffic of the crews working the fire.

1.1 Focus and Contributions

This report treats detection, characterization, and connectivity as one problem with one architecture, and it is organized around the question of how the wildfire information gap can be fixed. The report's contributions are as follows:

1. A phase-by-phase response concept of operations that follows a single ignition from standing watch through autonomous confirmation, cued characterization, fused prediction, sustained communications support, and reignition watch.
2. An onboard reduction and detection chain ordered by energy cost, which discards most pixels in fixed logic before the general-purpose processor is woken, and whose output is a georeferenced polygon and a small context image.
3. A cued-handoff and role-change protocol for the low-altitude tier, including the in-situ measurements that operational spread models currently must assume, the combustion-chemistry products that distinguish flaming from smoldering fire, and the transition from sensing to communications relay once the incident is established.

1.2 Scope and Status

This is a systems-architecture and concept-of-operations report. The architecture's own budgets, energy, downlink rate, geolocation error, and latency, are described here as relationships and constraints rather than hardware specific numbers, because they depend on differing airframe and payload selections this report deliberately leaves open.

2. Background and Related Work

2.1 Uncrewed Aircraft in Wildfire Response

Uncrewed aircraft have been studied both as sensors and as network infrastructure for disaster response [12], [13]. Wildfire-specific work has concentrated on airborne fire and smoke segmentation from visible and thermal imagery, supported by the public Fire Luminosity Airborne-based Machine Learning Evaluation (FLAME and FLAME 2) datasets [14], [15] and the Corsican Fire Database [16]. A parallel body of work has used small aircraft as chemical samplers: NASA Ames researchers have proposed a smoke-sampling payload spanning methane, carbon dioxide, carbon monoxide, nitrogen species, ozone, organic vapors, and particle size distribution [17], and field campaigns over prescribed burns have produced position-tagged concentration maps and vertical plume structure from repeated flights [18], [19].

2.2 High-Altitude Platforms

The Zephyr solar electric aircraft flew continuously for 67 days in the stratosphere in 2025, superseding its own 64-day record and operating near twenty kilometers [20]. Networking research has developed a correspondingly detailed vision of such high-altitude long-endurance (HALE) aircraft as a persistent aerial communications layer [21].

2.3 Fire-Spread Modelling and Machine Learning

Operational spread prediction largely descends from the Rothermel formulation [22], packaged into landscape tools of which FARSITE is the most widely used [23], with open-source real-time implementations now available that track the fire front directly and accept externally supplied wind, fuel moisture, and topography inputs [24], [25], [26]. Learned approaches have advanced quickly on the back of shared datasets for training and comparison [27], [28], and a recent evaluation trained several architectures on a large set of Hawaii wildfires and compared the best performers against FARSITE on the Maui fires, finding them competitive but concluding that learned models cannot yet replace established physical models [29].

One property is shared by every model in this literature, learned or physical: the wind, the fuel moisture, and the state of combustion are supplied from outside, usually from a weather station some distance away and a vegetation map compiled months earlier. How the fire is burning, at this spot, right now, is estimated. Closing that specific loop is what the low-altitude tier of this architecture exists to do.

2.4 Gap Statement

Prior work addresses persistent coverage, fine-grained sensing, in-situ chemistry, spread prediction, and emergency connectivity as separate problems, typically with separate platforms and separate programs. What is missing is an architecture in which one persistent stratospheric asset performs autonomous detection at the edge, uses that detection to cue a second tier that measures precisely the quantities the models need, feeds those measurements into prediction while they are still timely, and then converts that first tier into the communications infrastructure of the response.

3. Proposed Solution

The architecture comprises three tiers: Tier 1 is a solar-electric high-altitude long-endurance (HALE) aircraft station-keeping in the lower stratosphere with a multiband infrared and visible payload and an onboard processing stack; its function is persistent wide-area screening, autonomous confirmation, and geolocation. Tier 2 is a fleet of small uncrewed aircraft held at readiness and dispatched on cue; its function is to measure microclimate temperature and wind data. Tier 3 is the ground segment: incident command, the fusion and prediction engine, and the dissemination path to crews, aviation, and public alerting authorities.

Figure 1 shows our proposed 3-tier architecture, describing each tier and major actions that occur in each tier. On the right side, a flow diagram of actions represents broadly how our system operates.

System Architecture

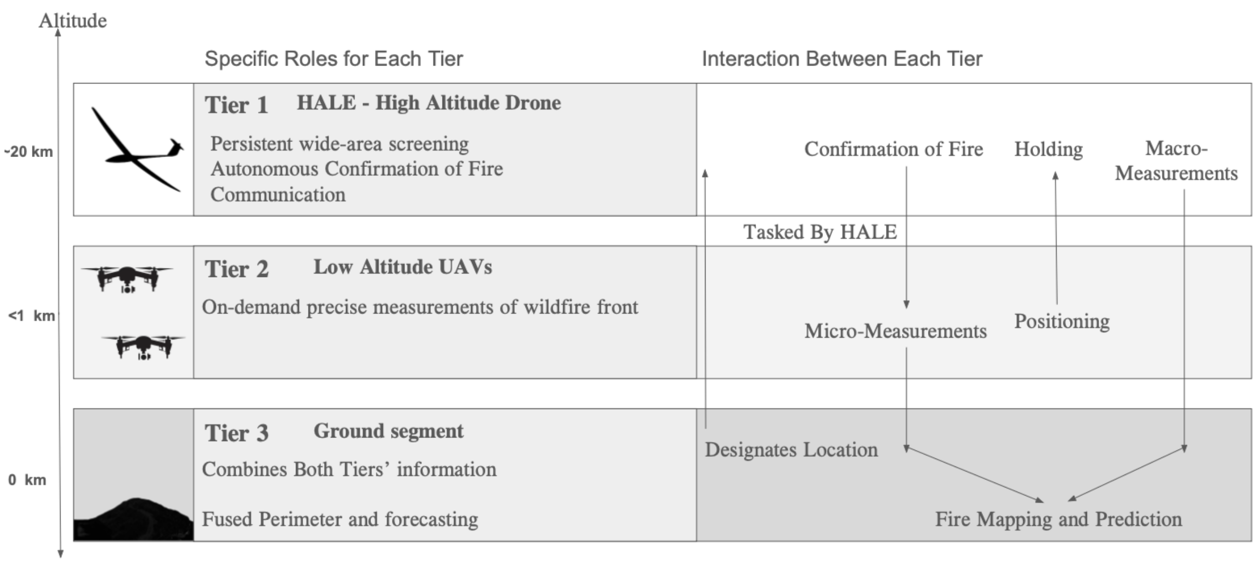

Figure 1. The three-tier architecture. The altitude ladder at left places each tier at its operating altitude. Arrows at right show the operational exchange between tiers. The HALE confirms the fire and passes macro-scale measurements down, tasks and positions the small aircraft, which return micro-scale measurements, with all products converging on fire mapping and prediction at the ground segment.

3.1 Design Philosophy

The design of the solution utilizes persistent monitoring over a specified area. This allows for constant monitoring through the fire season in fire-prone areas and significantly reduces the time between ignition and a response. The architecture is designed around components that work together to maximize coverage, speed, and accuracy.

3.2 Why Three Tiers

A three-tier approach allows for different types of information to be collected at different rates and scales, a necessity for wildfire monitoring and detection. While high altitude drones can monitor a large area, they cannot capture detailed conditions near a fire, which are typically collected by a low altitude drone. A ground station then, can combine all the data into useful information for responders. By dividing the responsibilities into three tiers, the system is enhanced and becomes more effective than a single platform.

3.3 Order of Operations

The response unfolds in six phases, and the rest of the paper is organized around them. In Phase 0 the high-altitude drone holds station over an assigned block and screens each frame against its own recent frames of the same ground; nothing is transmitted, and the ground segment maintains the terrain, fuel, weather, and route baseline the later phases draw on. In Phase 1, a thermal anomaly survives the onboard detection cascade and is confirmed on the aircraft, producing a georeferenced perimeter, a confidence value, and a context image. In Phase 2, that confirmed detection is promoted to a cue, and the low-altitude aircraft are dispatched with an approach vector. In Phase 3, those aircraft measure at the fire line what the stratosphere cannot resolve, such as, near-surface wind, fuel moisture, hot-spot geometry, and plume structure. In Phase 4, the ground segment fuses both scales into a single perimeter and a spread forecast. In Phase 5 the high-altitude drone changes role, becoming the communications relay for crews that have no other connectivity. Table 1 summarizes the phases, the tier acting in each, and the product that leaves the system.

Table 1 shows each stage of wildfire ignition and primary tiers associated with each. The fire-based event column shows the major event about the ignition of the wildfires. The acting tier informs on the Tier most heavily involved in the process. Finally, the system role describes what the proposed system is doing during each event.

Table 1. *Order of Operations*

Phase	Fire-Based Event	Primary tier Involved	System Role
0 — Surveillance	HALE deployment over an assigned block	Tier 1	A continuously refreshed observation baseline and a current terrain, fuel, weather, and route picture
1 — Detection and confirmation	A thermal anomaly survives the onboard cascade	Tier 1	Confirming hotspot as fire and compressing the perimeter vector to the low-altitude UAS.
2 — Cue, handoff, dispatch	Promotion of a confirmed detection	Tier 1 and Tier 2	A dispatch cue with approach vector; an updated shared airspace picture
3 — Characterization	Low-altitude aircraft collect information on fire	Primarily Tier 2 alongside support from Tier 1	Near-surface wind field, fuel moisture and structure, hot-spot geometry, combustion state classification, vertical plume structure

Phase	Fire-Based Event	Primary tier Involved	System Role
4 — Fusion, prediction, delivery	First fused observation set	Tier 3	Common operating picture: perimeter and rate of spread, forecast isochrones with uncertainty, smoke advisory, escape-route viability
5 — Sustained support	Incident established; characterization sufficient	Tier 1	Continuous connectivity for crews and aviation; sustains the perimeter

3.4 Representative System Requirements

The proposed architecture is designed around system level performance rather than hardware-specific specifications. Because the platform is presented as a concept of operations, Table 2 below represents engineering targets that guide the system and provide measurable goals for future implementation.

Table 2. *Suggested System Requirements*

Requirement	Purpose
Low detection latency	Ensuring that initial response is fast enough for effective suppression
Accurate fire geolocation	Certifies accurate dispatch of drones to the optimal location
Near zero false positive rate	Prevent unnecessary deployments that waste resources
Onboard data reduction	Reduce bandwidth and power requirements
Routine prediction refresh	Maintain constant awareness to ensure rapid response
Constant communications availability	Maintain responder connectivity and data transfer

Concept of Operations (Sections 4-9)

4. Phase 0 — Pre-ignition: The High-Altitude Drone Surveillance

The high-altitude drone holds station over an assigned block of high-value terrain and screens every frame against its own recent frames of the same ground. The general-purpose processor is asleep. Overnight the payload runs a reduced watch, keeping the primary hot-spot channel active while the remaining bands and the higher-power processing stages stay dormant to save energy in this diurnal energy cycle.

The low-altitude tier is parked and serviced during this phase, and the ground segment is maintaining the baseline it will need later: terrain and fuel maps, the current fire-weather picture, road and escape-route geometry, and the calibration for instrumentation (Figure 2). On the left, a scale represents the true scale

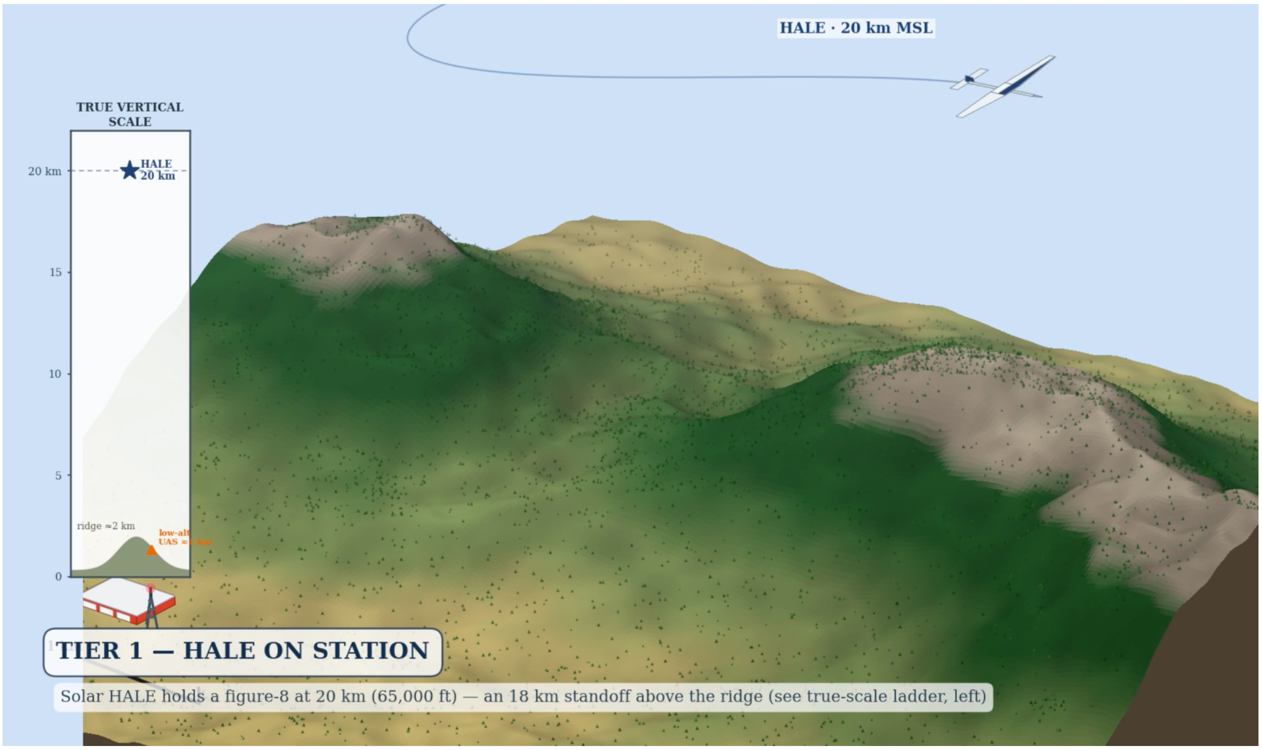

Figure 2. Tier 1 on station. *The solar-electric high-altitude drone holds a figure-eight at its operating altitude above the ridgeline it is assigned to watch; the true-scale ladder at left shows the altitude between the HALE and the terrain below it.*

4.1 Why the Lower Stratosphere (20 km)

The operating band is selected for four reasons. It sits above essentially all tropospheric weather, convective clouds, and precipitation, so the observation baseline is not interrupted by the weather the fire itself generates. It sits above certified commercial air traffic, which reduces airspace integration to only takeoff and landing. It lies near the altitude of minimum wind at mid-latitudes, which is the only reason a low thrust solar aircraft can hold station against the prevailing flow at all. And it places the platform above the thermal turbulence and boundary-layer gustiness that degrades lower-altitude uncrewed aircraft.

4.2 Station-Keeping and Search Patterns

Thin air imposes real maneuvering constraints, and they shape how the HALE responds. It is a lightweight, high-aspect-ratio solar-electric glider whose flexible structure cannot tolerate the wing-root loads of a steep bank. The bank angle is forced shallow, so the turn radius is large and effectively fixed.

Because the sensors are likely body-mounted, the field of regard is fixed relative to the airframe, and this produces two distinct blind spots. While the aircraft is banked, its own wing and fuselage mask part of the scene on the inboard side. Looking down at the shallow angles a wide orbit forces, ridgelines occlude the terrain behind them. A single circular orbit flown in one direction therefore leaves parts of a high-value block persistently under-observed, and if that ground is a populated canyon the gap is exactly where it must not be.

Two patterns follow from this, and the choice between them is a response decision. A figure-eight holds the aircraft over a single high-value point while reversing bank direction each lap, so terrain masked during one lobe is exposed during the next; it is the mode entered on promotion of a candidate detection and the mode held through characterization. A back-and-forth raster maximizes uniform coverage of a rectangular block and keeps re-look intervals even; it is the default during standing watch. Because pattern geometry is set by the turn radius at the airspeed the aircraft must hold against the prevailing stratospheric wind, pattern selection is a joint tasking and flight-condition decision, and the turn radius rather than the sensor sets the smallest area over which multi-angle overwatch can be maintained (see Figure 3).

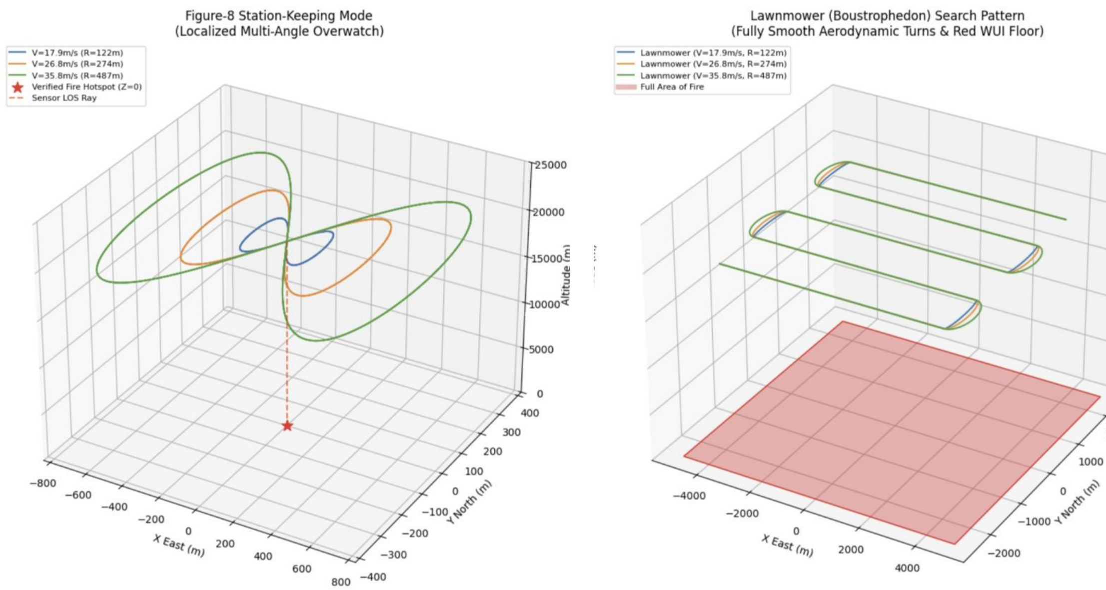

(a) Figure-eight station-keeping (b) Lawnmower raster

Figure 3. The two station-keeping patterns. *The figure-eight holds the aircraft over a single point and reverses bank direction each lap, exposing terrain that the airframe or a ridgeline masked on the previous lobe; the raster maximizes uniform coverage of a rectangular block with track spacing set so that successive swaths overlap given the fixed field of regard.*

4.3 Payload of HALE Aircraft

The band selection follows from the physics of what a fire radiates. A blackbody's peak emission wavelength scales inversely with temperature, so ambient terrain near three hundred kelvin peaks in the longwave window while a flame front at eight hundred to twelve hundred kelvin peaks in the mid-wave. More usefully, mid-wave radiance rises very steeply with temperature, so a fire occupying a small fraction of a pixel footprint produces a large mid-wave signal against a cool background, whereas the same subpixel fire appears in the longwave band only as a modest warming of the pixel average. Burning vegetation additionally releases neutral potassium, which emits a resonance doublet near 766.5 and 769.9 nanometers; this line emission is a chemical fingerprint of active biomass combustion [30], which is what promotes a hot pixel to a burning pixel.

Table 3 sets out the bands the payload carries and the job each one does. The mid-wave channel is the primary detector of small new fires and the longwave channel is the baseline it is compared against; the short-wave channel is what keeps a detection alive once a plume has thickened over the fire; the potassium doublet is what separates burning vegetation from a merely hot surface; and the visible and near-infrared bands supply plume geometry and the context image a commander actually looks at. No row is sufficient on its own, and that is the point: each band rejects a class of false alarm the others cannot, which is what makes the cascade of Section 5 possible.

Table 3. Payload bands carried by the HALE, and the discriminant each contribute to the detection cascade.

Band	Nominal Wavelength range	Role in the response
Mid-wave infrared	3–5 μm	Primary hot-spot discriminant. A combusting surface is orders of magnitude brighter here than cool background, which makes this the cheapest first rejection test and the sensor that finds sub-pixel ignitions.
Longwave infrared	8–12 μm	Contextual background temperature. Distinguishes a true thermal anomaly from uniformly sun-warmed terrain, rock, or bare soil, and supplies the local baseline for the contextual gate.
Short-wave infrared	1.4–1.8 μm	Penetrates smoke, so detection survives a thickening plume. Supplies the context image and supports discrimination of flaming from smoldering combustion.
Potassium line	766.5 / 769.9 nm [30]	Positive confirmation that vegetation is burning. The one band that a heavy overlying plume can defeat, which is why Section 5.6 provides an override.
Visible and near-infrared	0.45–0.9 μm	Plume geometry, fuel condition, and the human-interpretable context image. Daylight only.

4.4 Energy Constraints of HALE Aircraft

Everything the HALE does in a response is bounded by weight and electrical power, because it is powered solely by its battery. Survival across a single night requires that the energy banked by sunset exceed everything drawn before sunrise, and this must hold on to the worst night of the deployment window: the longest night, following the dimmest day, at the coldest ambient temperature. If it fails, the aircraft cannot hold altitude until morning and would make an uncontrolled descent. Therefore, this condition sets battery mass, and battery mass competes directly with payload mass, which is why the payload is capped at a few kilograms and why every watt of continuous draw needs to be efficient.

Ambient temperature in this band is extremely cold, detector noise falls with temperature, and passive radiative cooling to the cold sky improves infrared sensitivity for free. The second works against it: the processors and the battery management system have minimum operating temperatures and require active electrical heating, so onboard computation carries a parasitic thermal cost on top of its direct draw, and that cost peaks on the coldest nights when the energy margin is thinnest.

5. Phase 1 — Ignition: Onboard Detection and Confirmation

An ignition produces a small, very hot area against a cool background, which is a strong signal in the midwave infrared. The HALE screens for it continuously, then subjects each candidate to a sequence of independent tests described in the subsections that follow: cheap physical gates that reject water glint, cloud, and uniformly sun-warmed terrain; a comparison against the local background; a check for the potassium emission that accompanies burning vegetation; a learned classifier that rejects the false-alarm classes; and a persistence requirement that the candidate reappear in the same georeferenced location across consecutive observation cycles (Figure 6).

The following two figures (Figure 4 & Figure 5) represent two crucial steps in the progression of a wildfire. Figure 4 represents the early ignition of a wildfire. This is intended to highlight how detection is not instantaneous because if early ignitions were vectorized too early, it may lead to an increase in false positives. This means there is a delay between ignition and the detection and confirmation stages. Figure 5 shows the HALE categorizing the fire after some time from a designated size from ignition.

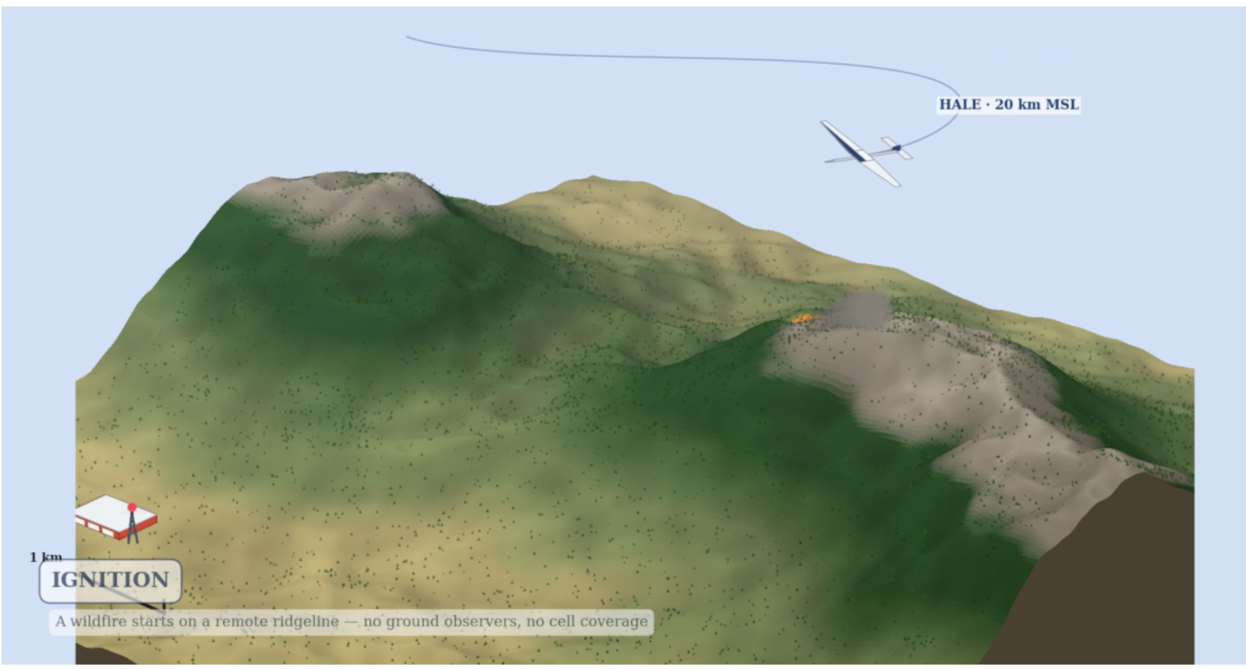

Figure 4. Ignition. *A wildfire starts on a remote ridgeline with no ground observers and no cellular coverage; inside the block the HALE is already observing against its own prior frames.*

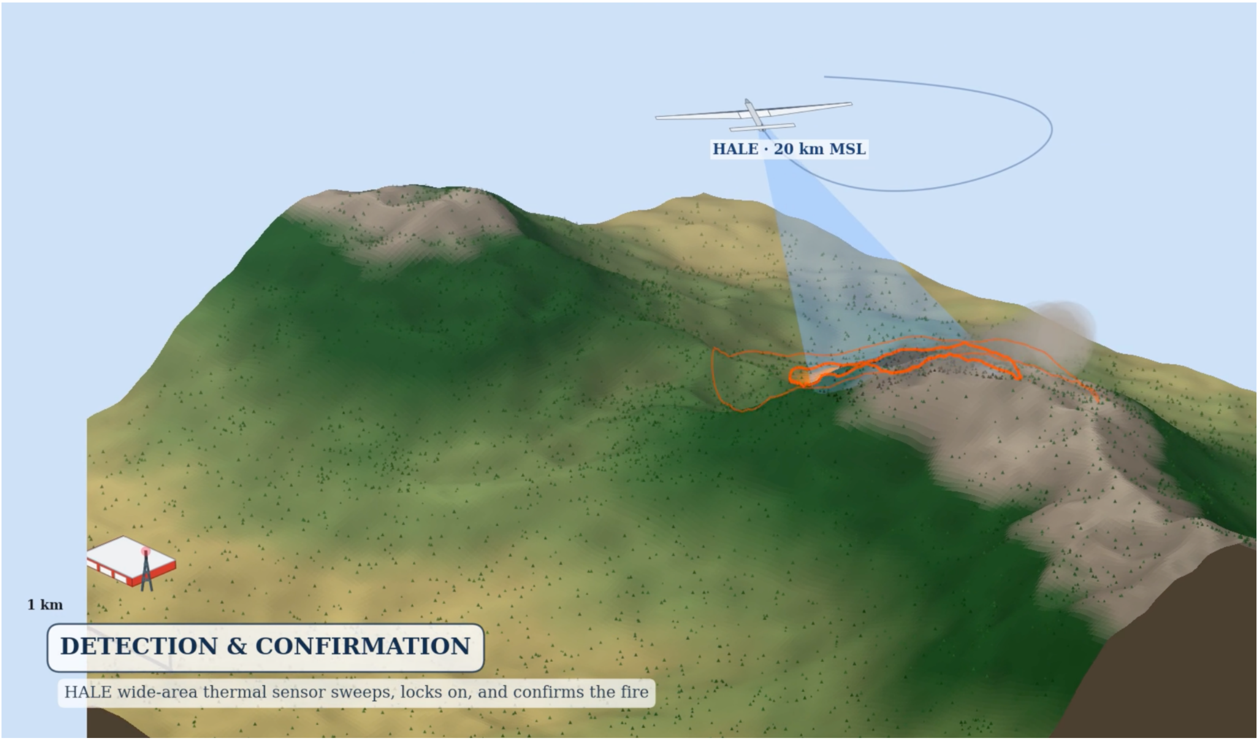

Figure 5. *Detection and confirmation. The wide-area thermal sweep locks on and confirms the fire, and the mapped perimeter becomes the first product to leave the aircraft.*

5.1 Energy Constraint

The governing constraint is the ratio of available power to available bandwidth. Several simultaneous sensor streams at useful frame rate and bit depth represent far more raw data than the HALE can encode and radiate within its power and spectrum budget, and the transmitter is itself one of the larger discretionary loads on the aircraft. The architecture therefore does not attempt to move imagery. It analyses and filters at the edge and reduces the scene to compact products, vector perimeter and a small, compressed context image, before anything reaches the transmitter, and the chain is ordered so that the most power intensive stages run least.

5.2 Band Registration and the Cascaded Gates

Raw sensor outputs are moved directly into fixed programmable logic rather than into a general-purpose processor, for two reasons: simple per-pixel operations map onto parallel logic at very high throughput per watt, and the imagery can be processed without waking the processor, which is the single largest discretionary power sink. The cameras are neither identical nor co-aligned, so a point on the ground falls on a slightly different pixel in each band and the bands must be registered into a common frame before any multi-band logic can run.

Co-registered pixels then pass through a cascade of cheap tests ordered so that the test rejecting the most pixels for the least energy runs first. A land mask is applied first, because water produces solar glint that is bright in the same bands a fire is. A cloud test follows, because reflective tropospheric weather blocks the infrared line of sight to the ground entirely and any measurement through it is meaningless. Mid-wave radiance is then compared against an absolute threshold; most of any surviving frame is cool background and fails immediately, so those pixels are discarded at essentially zero cost. Finally, the mid-wave to longwave difference is compared against a contextual threshold computed from the surrounding neighborhood which is the gate that rejects hot metal roofs, bare rock, and sunbaked urban surfaces while remaining robust to diurnal and terrain-driven baseline drift.

The value of this cascade is power management: the general-purpose processor is held in a low-power state and woken only when a spatially connected cluster of pixels survives every gate, which also suppresses single-pixel artefacts from detector noise and cosmic-ray hits. The cloud gate is where operating altitude pays a direct algorithmic dividend, since loitering above most cloud layers raises the pass rate of a test that would otherwise reject a large fraction of observations outright. Figure 6 explains the process of going through cascaded gates.

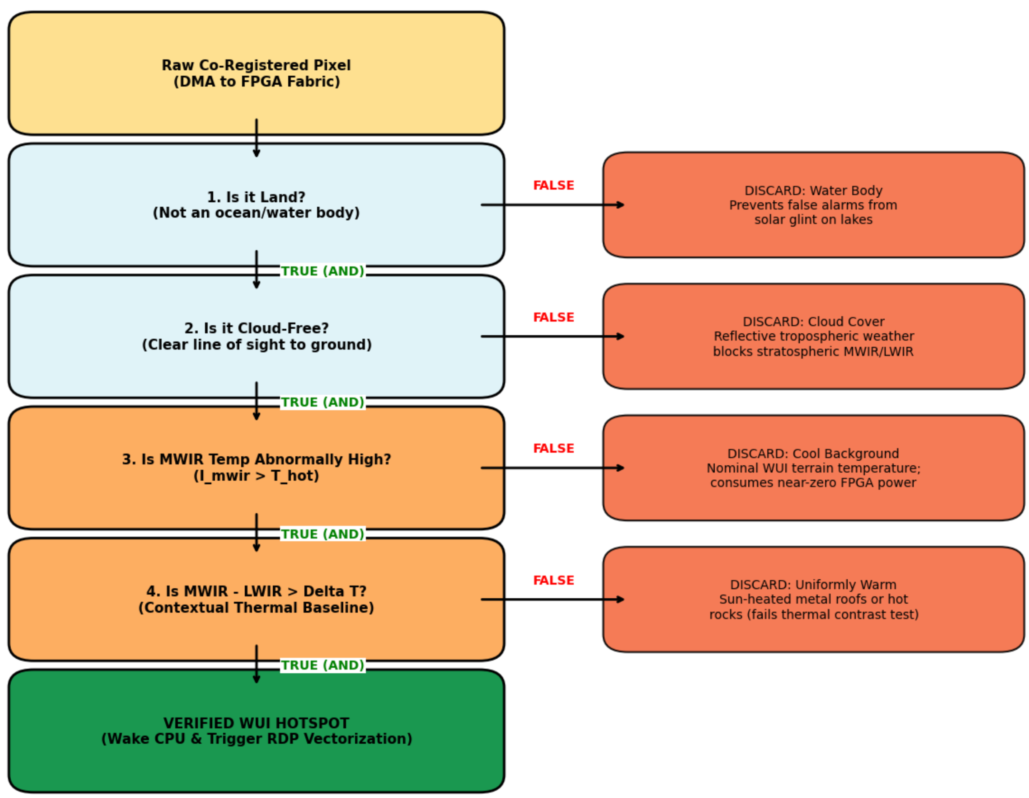

Figure 6. The onboard cascade. *Each gate rejects a distinct class of false alarm at negligible power cost, and the general-purpose processor is woken only when a candidate survives all of them. The ordering is by energy cost per rejected pixel. Each color represents a different stage in the Boolean logic. Yellow is the initial step where an anomaly is registered. Blue represents a visual gateway that does not use difference in infrared wavelengths. Orange is the gateway that uses infrared sensors. Green is the verified step. Finally, red is a rejected pathway.*

5.3 The Learned Screen and the Persistence Gate

Once the processor is awake, a compact convolutional classifier of the kind designed for power-constrained inference [31] evaluates the candidate cluster together with its surrounding context and assigns a confidence value. Placing the learned component here, behind two physically grounded gates, is a deliberate containment decision discussed in later sections.

A final temporal gate requires that a candidate reappear in the same georeferenced location across consecutive observation cycles before it is promoted because most remaining artefacts are transient while a fire is not.

5.4 Perimeter Vectorization and Geolocation

A confirmed cluster is isolated into a binary fire mask, separated into distinct fires where necessary, and traced into an ordered list of boundary points. Transmitting that raw per-pixel edge would be grossly inefficient for what is geometrically a simple closed curve, so the perimeter is simplified by Douglas– Peucker line reduction [32]: a straight segment is drawn between the endpoints of the boundary, the vertex of maximum perpendicular deviation is found, and if that deviation falls within a tolerance every intermediate vertex is discarded, while if it does not the far vertex is retained as a corner and the procedure recurses on the two sub-segments. The result is a polygon of critical vertices with a guaranteed bound on how far any point of the true perimeter can lie from the simplified shape, achieved with a fixed error bound, which is represented in Figure 7.

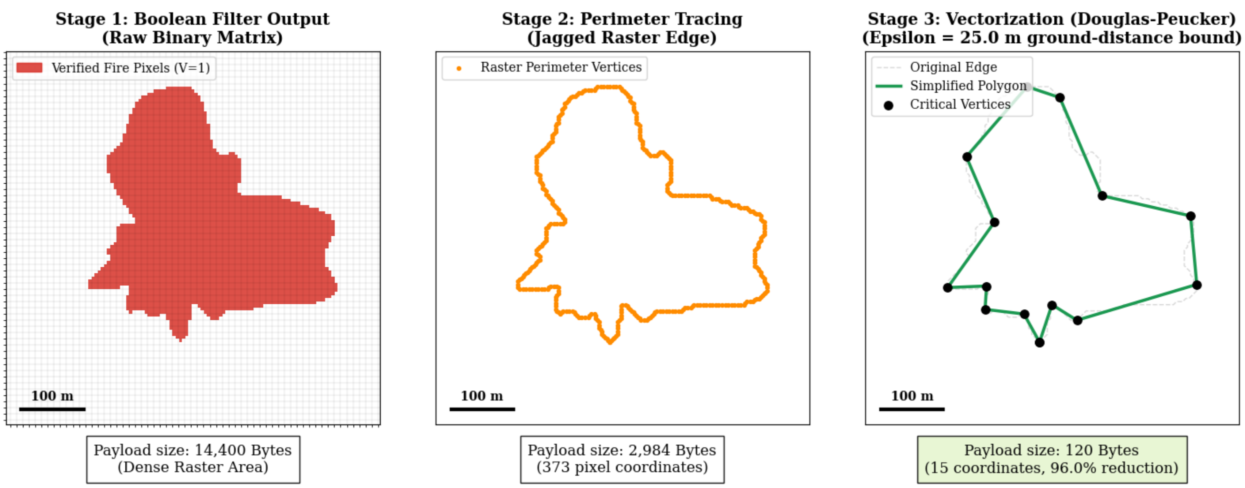

Figure 7. Perimeter vectorization. *The verified fire mask is traced to a raster boundary and then simplified to a polygon of critical vertices, collapsing a dense raster edge into a compact transmissible product under a fixed deviation bound.*

The vertices are still image coordinates and are not yet geolocated. Each is projected onto Earth map tiles by converting the pixel to a direction in the camera frame, rotating that direction into the local Earth frame using the HALE's measured attitude composed with the fixed camera mounting, and intersecting the resulting ray with a terrain model. A locally flat datum is inadequate, because terrain relief at stratospheric slant ranges displaces the intersection substantially on sloped ground, so the intersection must be taken against a terrain elevation model. And because the ray is anchored by measured attitude, pointing knowledge rather than camera resolution is the binding accuracy constraint: to first order a pointing error produces a ground error equal to the slant range multiplied by that angular error, so the long range that gives the HALE its coverage also amplifies every fraction of a degree of attitude uncertainty.

5.5 The Context Image

Commanders still require a visual of terrain and smoke for interpretation, so rather than downlinking full frames the system crops a small image centered on the anomaly from the short-wave infrared and visible bands. Compression uses the low complexity predictive coder developed by the space data systems community, which forms a prediction of each sample from already-processed neighbors both spatially and across bands, transmits only the residual, and supports a near-lossless mode with user-specified absolute error bounds [33], [34]. The fire is thus reported as a precise georeferenced polygon plus a small context image, never as full-motion video.

6. Phase 2 — Cue, Handoff, and Dispatch

Because the HALE loiters in the stratosphere it holds an unobstructed line of sight both to the ground cell where the fire has started and to any uncrewed asset within a wide radius, so it broadcasts the cue directly rather than routing it through terrestrial infrastructure that may be impacted by fire. The cue carries the polygon, the centroid, a confidence value, the context image (as shown in Figure 9), and a recommended approach vector derived from the observed plume orientation.

Three things happen in parallel during this phase: the low-altitude tier launches toward the coordinates, tasked by the HALE. The high-altitude drone transitions from area search to point overwatch, entering the figure-eight pattern over the hotspot, as shown in Figure 8.

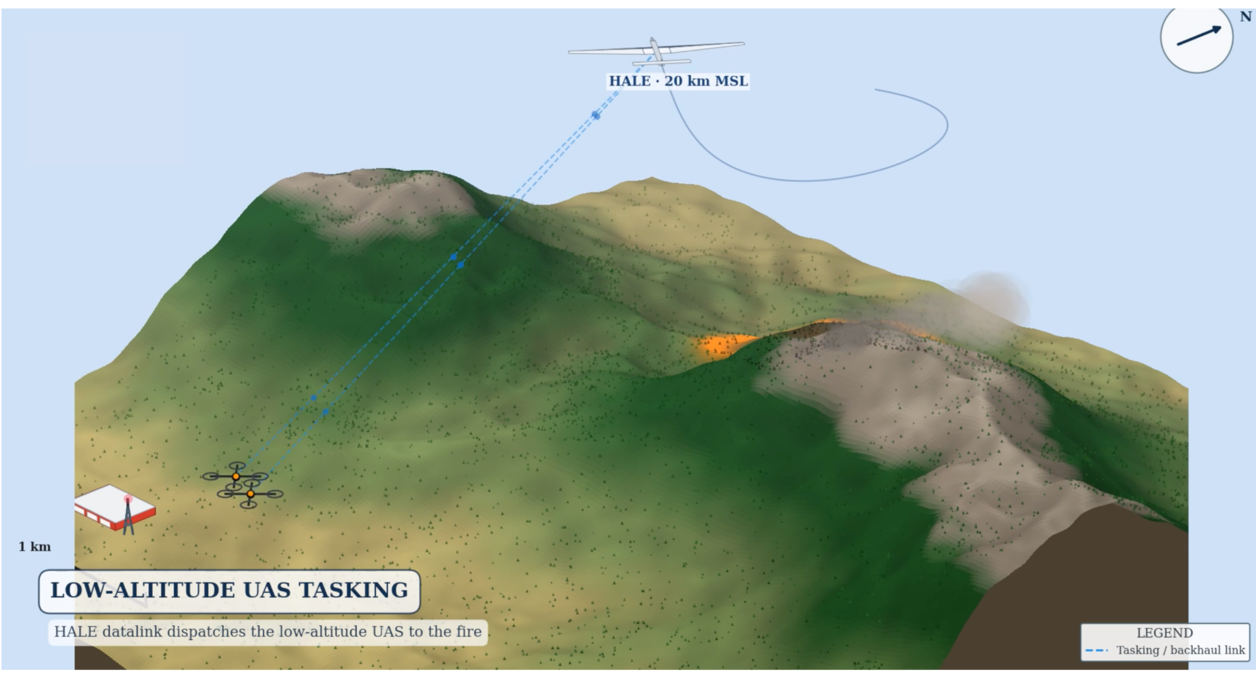

Figure 8. Tasking the low-altitude tier. *The HALE broadcasts the confirmed location directly to the small uncrewed aircraft, which launch toward the fire without routing through terrestrial infrastructure.*

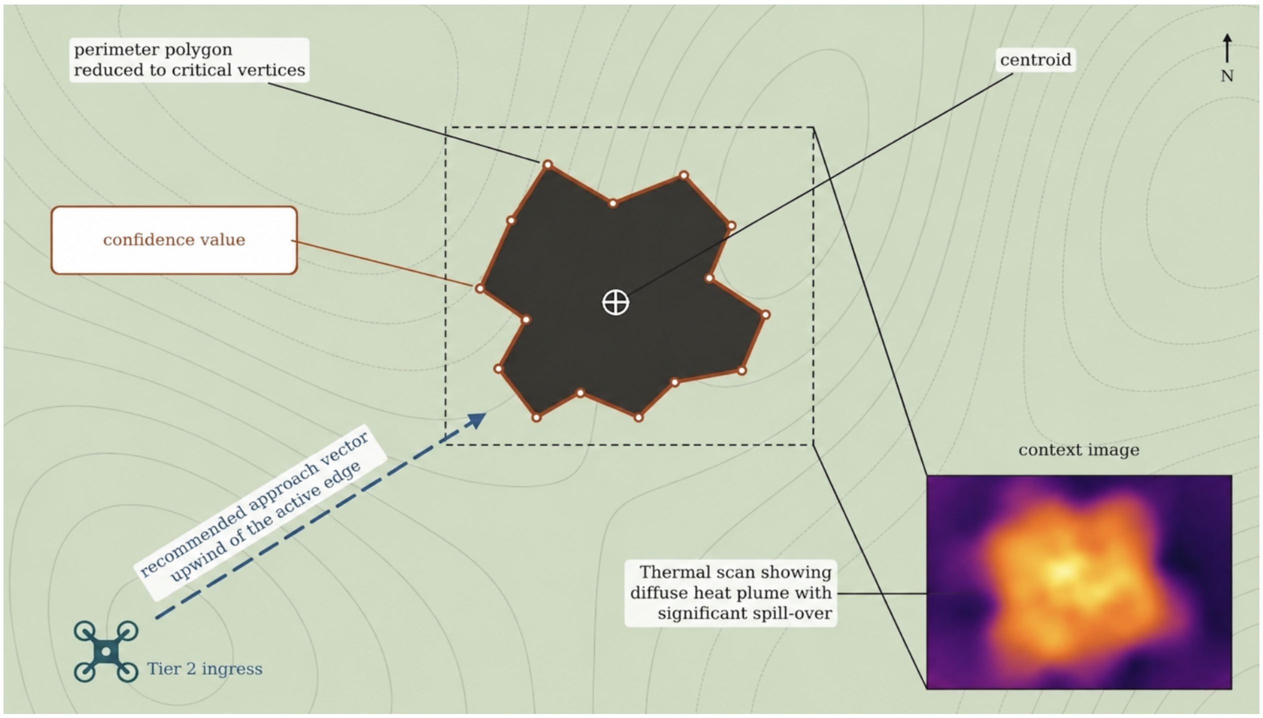

Figure 9. Product of Tier 1. *This figure shows the products of Tier 1, which include a perimeter polygon, a confidence value, a centroid, a context images, and a recommended approach vector.*

7. Phase 3 — Characterization at the Fire Line: The Low-Altitude Tier

The arriving aircraft fly beneath the smoke and cloud and measure the quantities that determine near-term fire behavior and that are physically inaccessible measurement from the HALE. The primary quantity among them is the near-surface wind field, because spread rate is dominated by wind and slope in every operational model from Rothermel onward [22], [23], and the wind that matters is the wind closest to the fuel bed.

It begins with a clean-air measurement upwind of the fire to establish a baseline. And it repeats those transects at intervals, because the change between passes, a hot spot growing, a plume rising, a wind backing, carbon monoxide climbing over a smoldering interior, is the signal that can feed preexisting prediction models as visualized in Figure 10.

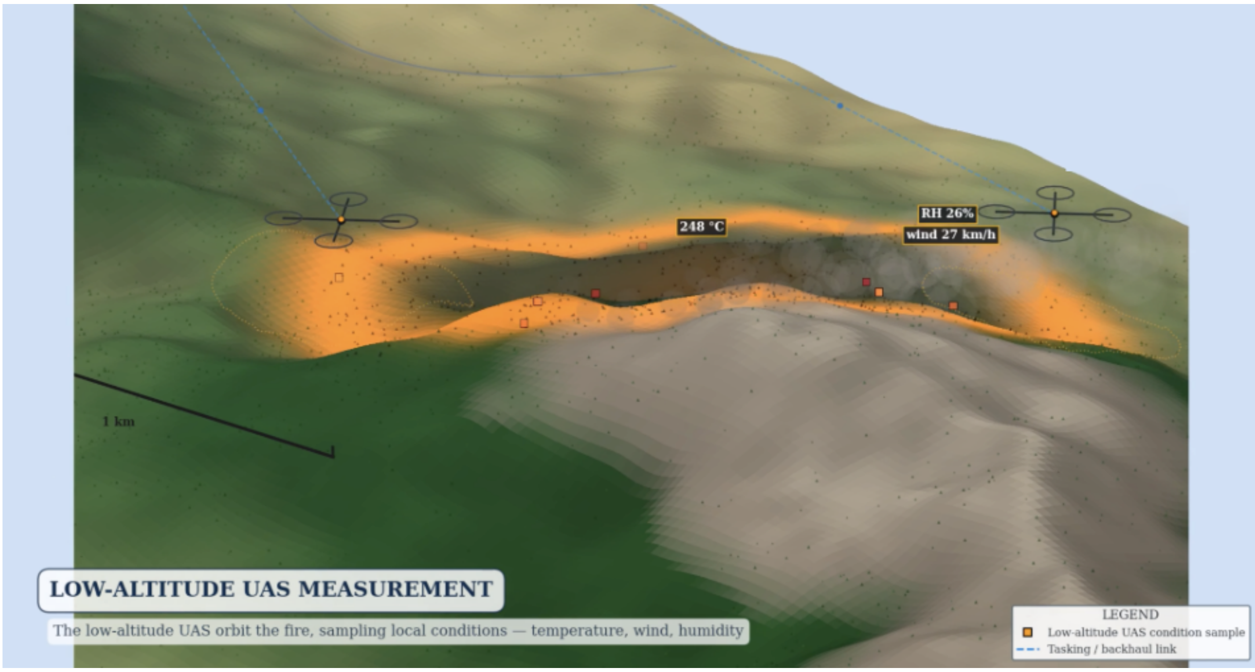

Figure 10. Close-in measurement. *Quadcopters orbit the fire beneath the smoke layer, sampling the local conditions: temperature, wind, and humidity at the fire line.*

7.1 What the Low-Altitude Tier Measures, and Why

The low-altitude tier exists because the quantities that determine what a fire does next are not observable from the stratosphere. Its measurements fall into four groups, summarized in Table 4: the micrometeorology at the fuel bed, the state of the fuel itself, the geometry and intensity of the active edge, and the chemistry of the combustion.

Vertical structure deserves particular emphasis, because it is invisible from above and decisive for response. A drone can begin below a plume and climb through it, recording how temperature, humidity, wind, smoke, and gases change with height. A plume typically has a hot rising core, cooler edges, layers where smoke concentrates, and different wind directions at different altitudes, and those differences determine where smoke travels and whether embers are lofted toward structures. This is the localized sounding that NASA FireSense has tested uncrewed aircraft to provide when weather balloons cannot be launched because firefighting aircraft are working overhead [35].

Read across its rows, Table 4 summarizes the measurements taken by the low-altitude UAS.

Table 4. *Low-altitude measurements, their role in the response, and the sensing approach used to obtain them.*

Measurement group	Quantities	Why the response needs it	Sensing approach
Micrometeorology	Wind vector, temperature, relative humidity, pressure at and above the fuel bed	Spread rate is dominated by wind and slope [27], [28]; the relevant wind is meters above the fuel, not the synoptic wind	Onboard meteorological package with a validated wind-estimation method; repeated vertical profiles

Fuel state	Fuel type and structure class, fuel load, fuel moisture content	Determines ignition probability, spread rate, flame character, and how hard the fire will be to hold	Visible imagery with a learned classifier; multispectral moisture indices; laser scanning for height, canopy depth, and density; radar, which works through cloud
Active edge	Surface temperature, hotspot geometry, relative intensity, sub-meter edge detail	Refines the wide-area perimeter locally and identifies newly active and apparently contained sections	Radiometric longwave imager with per-pixel temperature values and georeferenced export; visible imagery
Combustion products	Carbon dioxide, carbon monoxide, methane, organic vapors, nitrogen species, fine and coarse particulate matter	Distinguishes flaming from smoldering, identifies reignition potential, and drives smoke and exposure advisories [17], [36], [37]	Electrochemical, infrared-absorption, laser-scattering, photoionization, and laser-absorption instruments; upwind background reference each sortie

7.2 Measurement Integrity

Direct sampling introduces error: instruments respond with a lag, so an aircraft in forward flight records a concentration displaced from where the sampled air entered the inlet, and the processing chain must shift readings back toward the probable sampling location. Low-cost optical particulate instruments require wildfire-specific correction, having been found to substantially overestimate smoke concentrations before calibration [38].

Every observation carries four elements: the value, the position and altitude, an accurate timestamp, and a statement of uncertainty and calibration state. And every observation carries a confidence classification, reliable, usable with caution, or unreliable, that propagates all the way into the displayed product, so that a saturated or drifting reading cannot silently become a contour on a map.

8. Phase 4 — Data Diffusion: Fusion, Prediction, and Delivery

Everything from both tiers arrives tagged with position and time, which is what allows observations made by different aircraft to be combined into one prediction. The ground segment associates them, maintains the perimeter as a time-indexed polygon with per-vertex uncertainty, and lets a small number of locally precise low-altitude corrections adjust the wide-area estimate without discarding it. Prediction then proceeds from the fused perimeter and the measured local wind and moisture fields through established spread models, producing forecast perimeters at successive horizons with explicit uncertainty.

Delivery is differentiated by user, because bandwidth and attention differ by user. Incident command receives the wide picture: perimeter, rate of spread, forecast isochrones, smoke and air-quality advisories, and a prioritized prompt for resource allocation. Crews on the line receive something deliberately smaller on handheld devices, perimeter, hot spots, wind vectors, team positions, and escape-route status, because a video feed is both impossible on their bandwidth and a distraction in their situation as shown in Figure 11 and Figure 12.

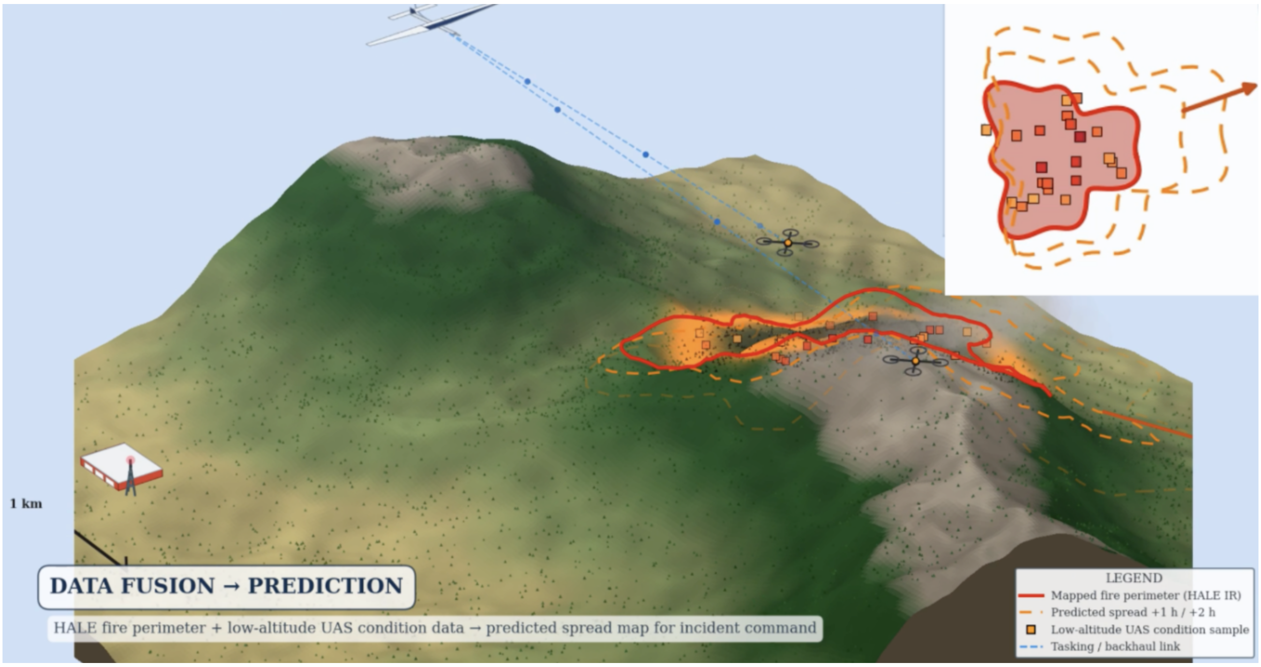

Figure 11. Fusion and prediction. *The mapped perimeter and the local condition data are combined into the predicted spread map delivered to incident command, with forecast perimeters shown at successive horizons.*

8.1 Data Synchronization and Preprocessing

Data Synchronization

The data must be synchronized and preprocessed before it can be used in decision making. The following four steps outline this process:

Frame Selection

Not every captured frame contributes useful information. Frames affected by excessive motion blur, or other forms of obstruction, are automatically rejected before entering the fusion pipeline. Consecutive frames with minimal scene change might also be discarded to reduce compute load and the bandwidth requirements while preserving all significant information.

Sensor Registration

After quality filtering, imagery from different sensors is spatially registered so that corresponding pixels represent the same physical location on the ground. Registration compensates for differences in sensor position, field of view, and aircraft altitude before multi-sensor analysis is performed.

Edge Compression

The synchronized observations are then compressed into compact products, including vectorized fire perimeters, environmental measurements, and small context image. Rather than sending full images and data, only the processed products are transmitted to the ground segment, significantly reducing bandwidth while preserving the information required for the ground system's prediction.

8.2 Association and Perimeter Maintenance

Fusion is keyed on the only two quantities they share unambiguously: position and coordinated universal time. Every polygon vertex from the HALE and every sample from the low-altitude tier carries both, so association reduces to a spatio-temporal nearest-neighbor problem within a tolerance set by the respective geolocation uncertainties: dominated, for the HALE and for the low-altitude tier by satellite navigation accuracy and HALE attitude.

The fused perimeter is maintained as a time-indexed polygon with per-vertex uncertainty. An uncertainty-weighted update allows a small number of precise local observations to correct the global estimate without discarding it, which is the property that makes the two tiers complementary.

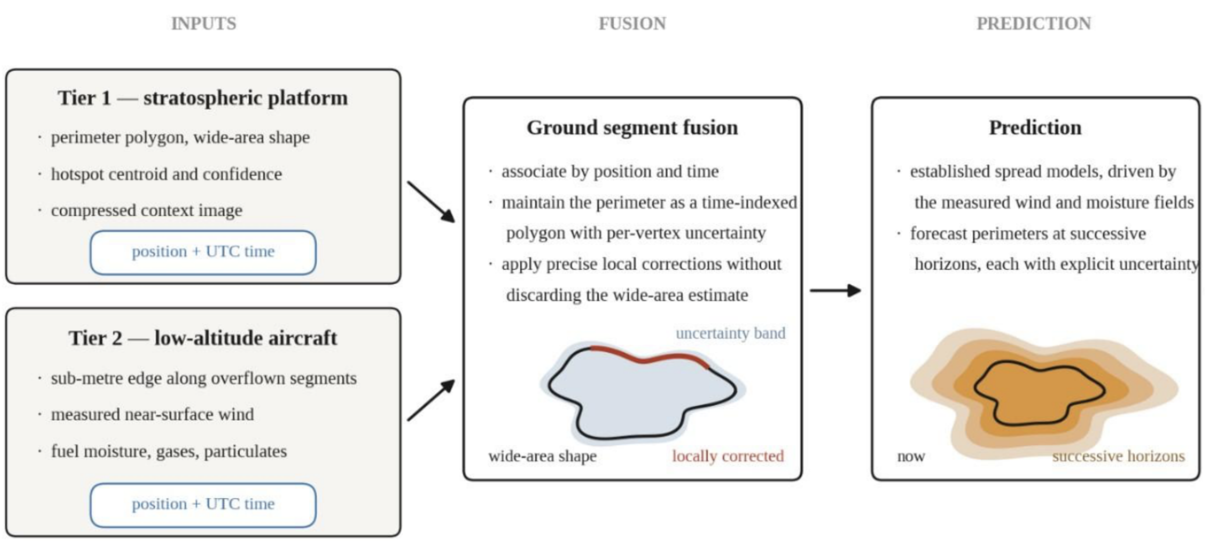

Figure 12. Data Combination. *Observations from both tiers become one estimate of one fire, are turned into forecasts with explicit uncertainty and are delivered at the fidelity each user can act on.*

8.3 Prediction

Prediction proceeds from the fused perimeter and the measured local wind and moisture fields through established semi-empirical spread models [22], [23], producing forecast perimeters at successive horizons. Three properties of the arrangement matter for response. The perimeter that starts the simulation is measured since the HALE's vectorized polygon can be ingested directly as the initial front by open-source real-time implementations built for this purpose [24], [25], [26]. The wind and fuel-moisture fields those implementations require are supplied by the low-altitude tier rather than by a distant station, which is the substitution that makes a short-horizon forecast worth issuing. And one of those implementations supports mitigation experiments, so commands can ask not only where the fire will go but how it would respond to a line cut at a particular place, which converts a forecast into an answer about a decision as shown in Figure 13.

From two scales to a forecast

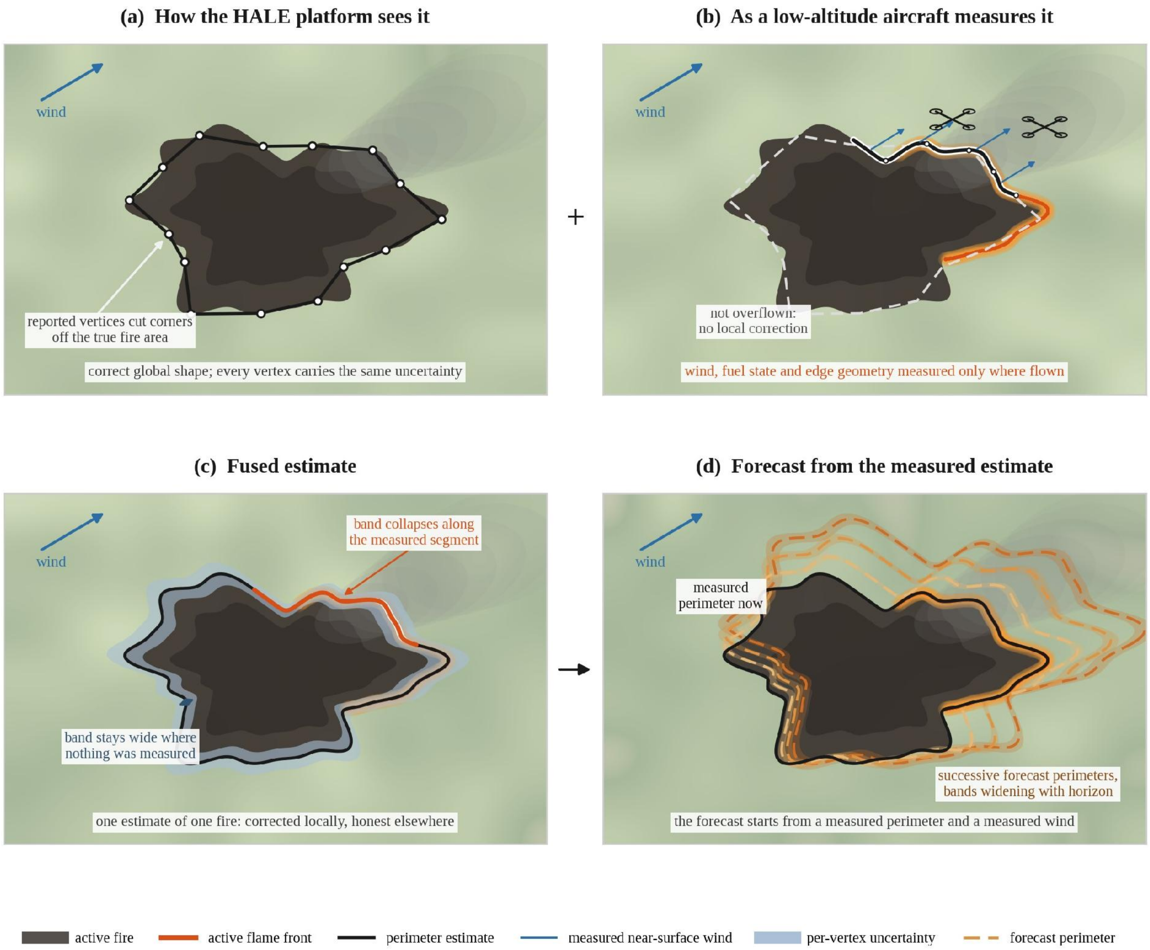

Figure 13. From two scales to a forecast (schematic). The high-altitude drone reports the global shape of the perimeter at high cadence but with coarse vertex spacing; a low-altitude aircraft measures sub-meter edge geometry, near-surface wind, and fuel state along the segment it overflew. An uncertainty-weighted update reconciles the two into a single estimate whose uncertainty band collapses where measurement exists and stays wide where it does not, and that measured estimate is the initial condition for the spread forecast, whose successive perimeters carry widening uncertainty with horizon. Terrain, smoke, and perimeter geometry are illustrative.

Figure 14 highlights how information gathered through the differing tiers is cleaned and combined to be put into existing wildfire spread models to create more accurate predictions and representations using highly precise data this architecture gathered.

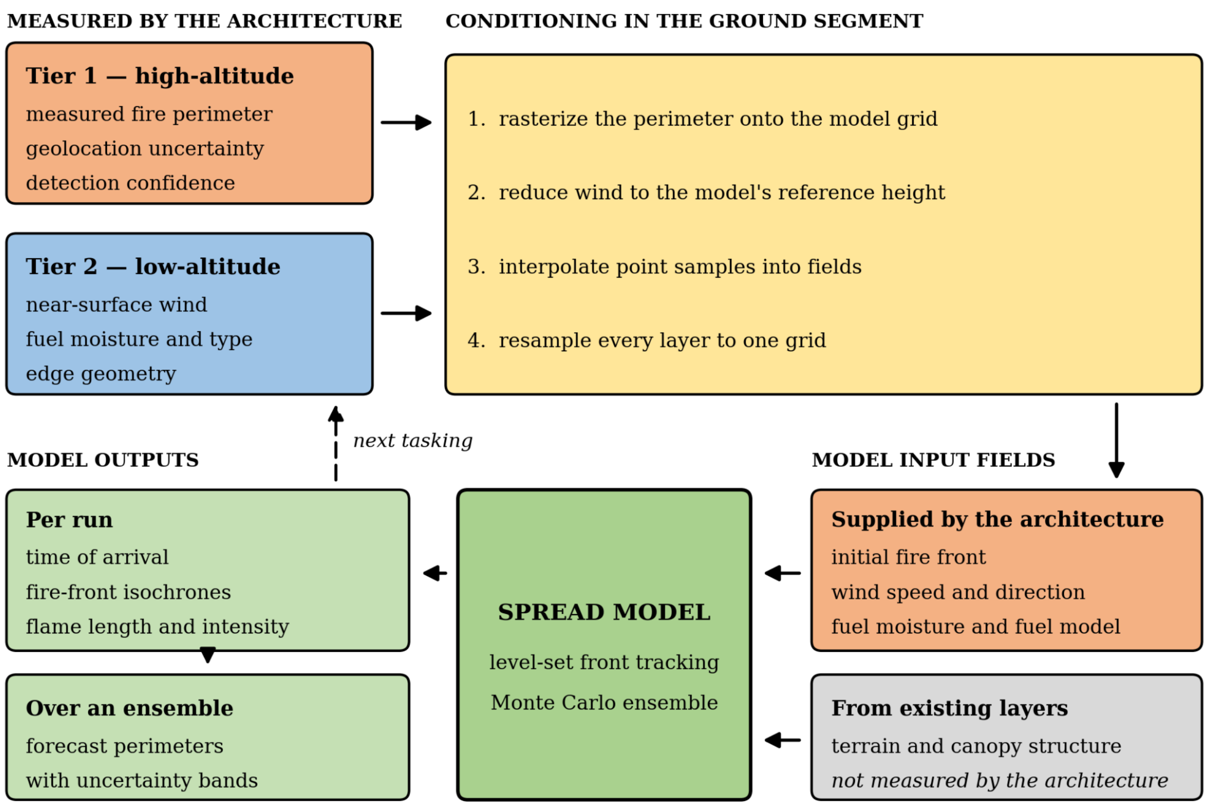

Figure 14. The prediction interface. Observations from both tiers are conditioned by the ground segment into the input fields a spread model accepts, run forward, and returns as forecast products. Terrain and canopy structure are inherited from existing landscape layers rather than measured. Field names follow the ELMFIRE input specification [24].

Where higher fidelity is warranted and time permits, coupled fire–atmosphere modelling captures the two-way interaction between the fire and the air above it [39], and the architecture is aimed at the assimilation pattern NASA is developing for its Wildfire Digital Twin, in which a running physical model is repeatedly corrected by fresh observations and produces an ensemble forecast [40].

8.4 Products

The common operating picture supplies current conditions, forecasts, and recommendations. Current conditions include the verified perimeter, active hot spots, relative intensity, flaming and smoldering sectors, the measured wind field, plume position, ground-level particulate exposure, and where the aircraft have flown and how old each layer is. Forecasts include perimeters at successive horizons with uncertainty, per-cell probability of spread, expected smoke transport and air-quality impact, and reignition-risk sectors. Escape-route viability is the product most directly tied to the safety of first responders.

9. Phase 5 — Sustained Support and the Role Change to Relay

Once the incident is established and the ground-level picture has been gathered the role of the HALE transforms to become a communication relay because more crews and more aircraft are arriving in terrain that frequently has no cellular service and no line of sight to a repeater. Because the HALE loiters in the lower stratosphere, its radio horizon extends far beyond that of any terrestrial mast or ridgetop repeater, and unlike a satellite it holds that geometry over one incident continuously rather than passing over it. It is also close enough that the link budget is met with low transmit power and negligible propagation delay, which is what allows it to work with the handheld radios crews already carry instead of requiring dedicated satellite terminals. The low-altitude aircraft, already on station over the fire, are repositioned as local relays: holding station above the incident, they link ground units that are shadowed by terrain and retain their own upward line of sight to the HALE.

Figure 15 represents how the communication architecture would exist in a communication-limited zone. Tier 1 acts as the primary airborne relay backbone, while Tier 2 serves as the local tactical radio relay.

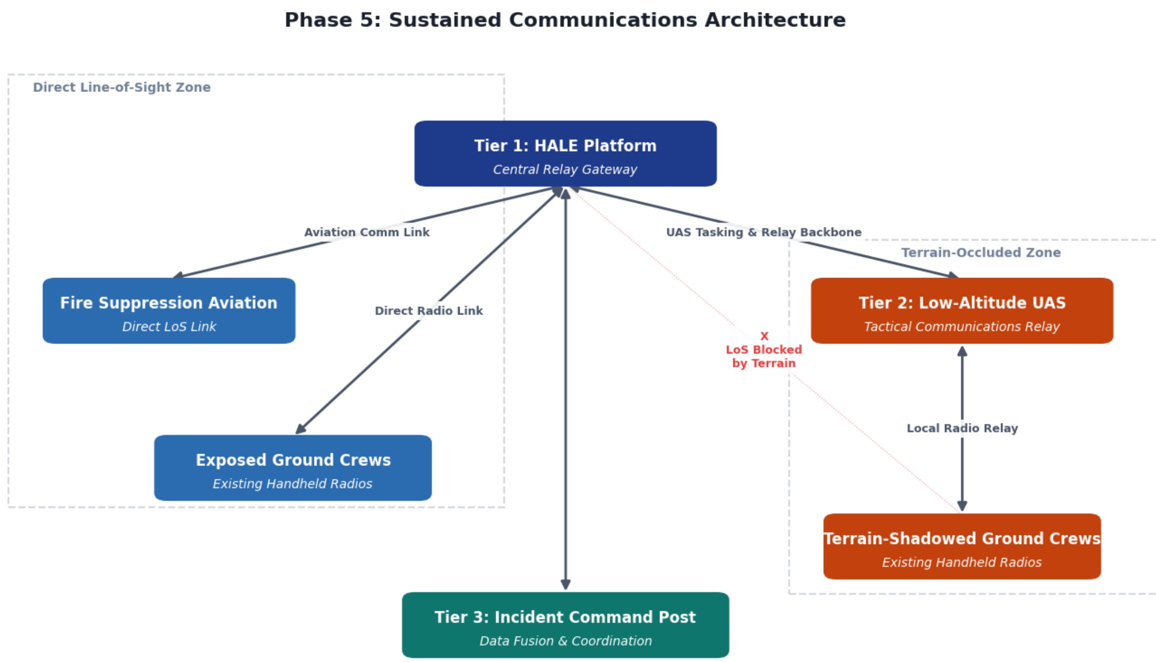

Figure 15. System Communication Architecture: *This demonstrates a proposed system that would allow Tier 1 to maintain communication among existing air traffic and both occluded and non-occluded ground crews.*

9.1 Whole System Integration

The proposed architecture functions as a single integrated system through continuous communication and data exchange between the three tiers. The high-altitude platform provides persistent monitoring and autonomous fire detection, while low-altitude drones collect detailed environmental measurements that improve fire suppression through better characterization and spread prediction. The ground station fuses these observations into a picture for responders and can update future drone tasks. This feedback loop enables faster detection, more accurate predictions, and improved awareness of the fire.

Information Flow between Tiers

The high-altitude platform constantly monitors its assigned region and performs edge node processing to identify and confirm potential ignitions. Once the fire is confirmed, required data, such as the fire perimeter, the confidence estimate and the image are transmitted to the ground segment. The ground station uses this data to appropriately dispatch the low altitude tier, which collects more specialized and localized measurements. These measurements are returned to the ground station where they are combined with the persistent observations from the HALE to maintain a continuously updated picture for responders.

Adaptive Mission Planning

Rather than a fixed plan created after initial dispatch, the mission is continuously updated based on incoming data from both the low and high-altitude tiers. As the fire evolves, updated predictions and measurements can trigger new flight paths, additional sampling locations or the deployment of an additional UAS. This adaptive tasking allows for the resources to remain focused on the regions of the fire that hold the most importance to situational awareness.

Data Fusion and Continuous Feedback

Every observation transmitted throughout the system includes its position, timestamp, and uncertainty, allowing measurements collected by different sensors to be accurately appropriated despite differences in frame rate, altitude, and collection time. The fusion process combines information from both aerial tiers into a single estimate of fire perimeter and surrounding environmental conditions. This data is then combined in a prediction engine, which generates revised predictions. These predictions subsequently influence future UAS tasking creating a continuous feedback loop in which each new observation improves future measurements and predictions. With the data and predictions provided by the feedback loop the system can be used to assist responders to the fire.

10. Challenges and Limitations

A solar-electric HALE is bounded seasonally and geographically because of its fragile launch and its dependency on battery power. Realistic deployment therefore implies seasonal repositioning, and a fleet-sizing analysis of the kind NASA has performed for high-altitude communications support [41].

Airspace integration. Low-altitude operations over an active fire must be deconflicted against crewed aviation working low and fast in degraded visibility. This is currently being refined by NASA's portable airspace management work. The HALEs take-off and landing require climbing and descending through commercial airspace, which would require pre-planning of take-off and landing roots to ensure no conflict between commercial air traffic and the HALE.

Privacy and mission scoping. A persistent fine-resolution imaging HALE loitering over inhabited terrain raises surveillance concerns because the architecture transmits polygons and small images centered on thermal anomalies rather than continuous imagery; wide-area imagery of non-fire areas is never transmitted and need never be retained.

Relationship to dedicated fire constellations. As FireSat and comparable programs scale, the marginal value of the high-altitude drone's detection function will fall. Its value as a persistent cueing and communications-anchoring node will not, because those functions depend on loiter and proximity. A plausible medium-term architecture uses a dedicated constellation for global detection and reserves high-altitude drones for surge deployment over high-risk regions during peak fire weather, where persistence and relay dominate.

Cost. A comparison of acquisition and operating cost against traditional firefighting aircraft and existing detection networks remains open and should accompany any proposal to field the system.

11. Conclusion

Wildfire response is limited by the age of the information available to the people making decisions, and that age is dominated by how long the sensing system waits before looking at the right spot. This report has specified how a three-tier architecture responds to a fire from the moment of ignition: a solar-electric high-altitude drone that simply stays overhead and confirms an ignition against several independent tests before anyone on the ground is informed; an onboard chain that converts a firehose of raw radiance into a compact georeferenced polygon without waking its expensive processor; a cued low-altitude tier that measures the wind field, the fuel state, and the combustion chemistry that operational models have to this point assumed; a ground segment that turns those measurements into fire prediction forecasts; and a role change that turns the Tier 1 into the communications infrastructure of the response once the fire is established.

12. References

[1] County of Los Angeles. Key facts about the Eaton and Palisades fires. https://file.lacounty.gov/SDSInter/lac/1199046_Eaton_Palisades_Fires_Anniversary_Fact_Sheet.pdf, 2026.

[2] T. J. Schmit, P. Griffith, M. M. Gunshor, J. M. Daniels, S. J. Goodman, and W. J. Lebair. A closer look at the ABI on the GOES-R series. Bulletin of the American Meteorological Society, 98(4):681–698, 2017. doi: 10.1175/BAMS-D-15-00230.1.

[3] NOAA/NESDIS/STAR GOES-R Algorithm Working Group. GOES-R Advanced Baseline Imager (ABI) algorithm theoretical basis document for fire / hot spot characterization. NOAA Center for Satellite Applications and Research, version 3.0, 2020.

[4] W. Schroeder, P. Oliva, L. Giglio, and I. A. Csiszar. The new VIIRS 375 m active fire detection data product: algorithm description and initial assessment. Remote Sensing of Environment, 143:85–96, 2014. doi: 10.1016/j.rse.2013.12.008.

[5] L. Giglio, W. Schroeder, and C. O. Justice. The Collection 6 MODIS active fire detection algorithm and fire products. Remote Sensing of Environment, 178:31–41, 2016. doi: 10.1016/j.rse.2016.02.054.

[6] NASA Earthdata. Fire Information for Resource Management System (FIRMS) frequently asked questions. https://www.earthdata.nasa.gov/data/tools/firms/faq, 2024. Near-real-time active fire data are distributed within approximately three hours of observation.

[7] ALERTCalifornia, University of California San Diego. Research and technology: statewide camera network and AI-assisted wildfire detection. https://alertcalifornia.org/technology/, 2026.

[8] P. Barmpoutis, P. Papaioannou, K. Dimitropoulos, and N. Grammalidis. A review on early forest fire detection systems using optical remote sensing. Sensors, 20(22):6442, 2020. doi: 10.3390/s20226442.

[9] Dryad Networks. Silvanet wildfire sensor. https://www.dryad.net, 2026.

[10] Earth Fire Alliance. Earth Fire Alliance's first three operational FireSats reach orbit. https://earthfirealliance.org/news-article/, July 2026.

[11] Muon Space. Muon Space deploys first three operational FireSat satellites for Earth Fire Alliance. https://www.muonspace.com/press/, July 2026.

[12] M. Erdelj, E. Natalizio, K. R. Chowdhury, and I. F. Akyildiz. Help from the sky: leveraging UAVs for disaster management. IEEE Pervasive Computing, 16(1):24–32, 2017. doi: 10.1109/MPRV.2017.11.

[13] M. Mozaffari, W. Saad, M. Bennis, Y.-H. Nam, and M. Debbah. A tutorial on UAVs for wireless networks: applications, challenges, and open problems. IEEE Communications Surveys & Tutorials, 21(3):2334–2360, 2019. doi: 10.1109/COMST.2019.2902862.

[14] A. Shamsoshoara, F. Afghah, A. Razi, L. Zheng, P. Z. Fulé, and E. Blasch. Aerial imagery pile burn detection using deep learning: the FLAME dataset. Computer Networks, 193:108001, 2021. doi: 10.1016/j.comnet.2021.108001.

[15] B. Hopkins, L. O'Neill, F. Afghah, A. Razi, E. Rowell, A. Watts, P. Fulé, and J. Coen. FLAME 2: fire detection and modeling — aerial multi-spectral image dataset. IEEE DataPort, 2022.

[16] T. Toulouse, L. Rossi, A. Campana, T. Celik, and M. A. Akhloufi. Computer vision for wildfire research: an evolving image dataset for processing and analysis. Fire Safety Journal, 92:188–194, 2017. doi: 10.1016/j.firesaf.2017.06.012.

[17] H. Miller and L. T. Iraci. UAS payload design for in-situ smoke plume sensing. NTRS 20240009739, NASA, 2024.

[18] UAV-based wildland fire air toxics data collection and analysis. National Library of Medicine, PMC10098707, 2023.

[19] Wildland fire emission sampling at Fishlake National Forest, Utah, using an unmanned aircraft system. Atmospheric Environment, 247, 2021.

[20] AALTO HAPS. Zephyr sets world record for longest continuous flight, flying 67 days in the stratosphere. https://www.aaltohaps.com/, May 2025.

[21] G. Karabulut Kurt, M. G. Khoshkholgh, S. Alfattani, A. Ibrahim, T. S. J. Darwish, M. S. Alam, H. Yanikomeroglu, and A. Yongacoglu. A vision and framework for the high altitude platform station (HAPS) networks of the future. IEEE Communications Surveys & Tutorials, 23(2):729–779, 2021. doi: 10.1109/COMST.2021.3066905.

[22] R. C. Rothermel. A mathematical model for predicting fire spread in wildland fuels. Research Paper INT-115, USDA Forest Service, 1972.

[23] M. A. Finney. FARSITE: fire area simulator — model development and evaluation. Research Paper RMRSRP-4, USDA Forest Service, 1998.

[24] C. Lautenberger. ELMFIRE: Eulerian level set model of fire spread. https://elmfire.io, accessed 2026.

[25] The MITRE Corporation. SimFire: wildfire simulation tool. https://github.com/mitrefireline/simfire, accessed 2026.

[26] Université de Corse / CNRS. ForeFire: open-source wildland fire front simulator. https://forefire.univ-corse.fr, accessed 2026.

[27] F. Huot, R. L. Hu, N. Goyal, T. Sankar, M. Ihme, and Y.-F. Chen. Next Day Wildfire Spread: a machine learning dataset to predict wildfire spreading from remote-sensing data. IEEE Transactions on Geoscience and Remote Sensing, 60:1–13, 2022.

[28] S. Gerard, Y. Zhao, and J. Sullivan. WildfireSpreadTS: a dataset of multi-modal time series for wildfire spread prediction. Advances in Neural Information Processing Systems, 36, 2023.

[29] J. Kim et al. Assessment of deep learning models integrated with weather and environmental variables for wildfire spread prediction and a case study of the 2023 Maui fires. Natural Hazards, 2025. doi: 10.1007/s11069025-07807-x.

[30] A. Vodacek, R. L. Kremens, A. J. Fordham, S. C. VanGorden, D. Luisi, and J. R. Schott. Remote optical detection of biomass burning using a potassium emission signature. International Journal of Remote Sensing, 23(13):2721–2726, 2002.

[31] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. MobileNets: efficient convolutional neural networks for mobile vision applications. arXiv:1704.04861, 2017.

[32] D. H. Douglas and T. K. Peucker. Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. Cartographica, 10(2):112–122, 1973. doi: 10.3138/FM57-6770-U75U-7727.

[33] M. Hernández-Cabronero, A. B. Kiely, M. Klimesh, I. Blanes, J. Ligo, E. Magli, and J. Serra-Sagristà. The CCSDS 123.0-B-2 low-complexity lossless and near-lossless multispectral and hyperspectral image compression standard: a comprehensive review. IEEE Geoscience and Remote Sensing Magazine, 9(4):102–119, 2021. doi: 10.1109/MGRS.2020.3048443.

[34] Consultative Committee for Space Data Systems. Low-complexity lossless and near-lossless multispectral and hyperspectral image compression. CCSDS 123.0-B-2, Blue Book, 2019.

[35] Earth Observer Staff. The FireSense project. NASA Science, March 20, 2025. https://science.nasa.gov/scienceresearch/earth-science/the-firesense-project.

[36] L. Xu et al. Parameterizations of U.S. wildfire and prescribed fire emission ratios and emission factors based on FIREX-AQ aircraft measurements. Atmospheric Chemistry and Physics, 24, 2024.

[37] NASA. Airborne observations of carbon dioxide and methane emission ratios from the Yosemite Rim wildfire, California. NTRS 20190000686.

[38] A. L. Holder et al. Field evaluation of low-cost particulate matter sensors for measuring wildfire smoke. Sensors, 20(17):4796, 2020. doi: 10.3390/s20174796.

[39] J. Mandel, J. D. Beezley, and A. K. Kochanski. Coupled atmosphere-wildland fire modeling with WRF 3.3 and SFIRE 2011. Geoscientific Model Development, 4:591–610, 2011.

[40] NASA Science Editorial Team. NASA “Wildfire Digital Twin” pioneers new AI models and streaming data techniques for forecasting fire and smoke. NASA Science, May 21, 2024.

[41] A. J. Burns, M. Johnson, J. Jung, and M. Fladeland. Estimating high altitude platform system requirements for wildland firefighting communications support. NASA/TM–20230018267, NASA Ames Research Center, December 2023.