

PROBLEM FRAMING

The Integration Gap: Why Wildfire Tech Keeps Failing the Last Mile

We can see almost every fire now. We still can't reliably turn seeing into doing. The distance between the two has quietly become the part of the problem most worth solving.

The 2025 Los Angeles County wildfires killed at least 30 people, and property and economic losses may exceed \$130 billion. They were not invisible. By the time they ignited, California had spent two years running an artificial-intelligence system that scans for smoke across more than 1,100 cameras statewide; satellites tracked the plumes from orbit; aircraft and ground sensors filled in the rest. We have become good at one thing in particular: knowing that a fire exists, and roughly where. We remain strikingly bad at converting that knowledge into coordinated action quickly enough to change the outcome.

That distance — between seeing a fire and acting on it — is the subject of this brief. It is also, increasingly, the field's own diagnosis.

The part we've largely solved

Detection has matured fast. The U.S. Government Accountability Office's (GAO) 2025 assessment describes a detection landscape improving across every modality — satellites, aircraft, drones, fixed cameras, and in-field sensors — with the underlying algorithms steadily getting better at spotting smoke and cutting false alarms. Two recent peer-reviewed surveys reach the same conclusion: the research front has moved quickly, and the core scientific questions of early detection — is something burning, and where — are increasingly answerable in software and in real time. None of this is finished. Satellites were never built for the job, and drones are still being qualified for it. But the trajectory is clear, and it is steep. Detection has advanced faster than the systems meant to act on it.

The last mile is still done by hand

Watch what happens after a credible detection. According to the GAO, the precise location of a suspected fire often still has to be pinned down by trained personnel; data from the patchwork of public and private systems, in the agency's words, "may not be easily integrated and compatible for analyses and quick response"; and the analysis itself is performed by a mix of algorithms and people, with responders dispatched as the situation seems to warrant — agency by agency. Each jurisdiction tends to run its own tools, its own feeds, its own picture. The alert does not flow into a single, machine-readable operating picture that everyone trusts. It lands in inboxes, radios, and phone calls — a manual coordination process improvised in real time, under pressure, across organizational seams. That is the last mile, and it is where modern detection quietly hands the problem back to the twentieth century.

Why the gap — not the sensor — is the binding constraint

Here the evidence turns counterintuitive in a useful way. A 2024 study in *PLoS One* examined six years of wildfire data from western Canada and found that detection-and-reporting delay accounts for only about three percent of total suppression cost — and that shaving a full hour off that delay reduces cost by roughly a quarter of one percent. Read carelessly, that says timing doesn't matter. Read carefully, it says something sharper: the leverage is not in detecting faster, but in what happens to a detection once you have it. If an hour of earlier warning barely moves the outcome, the constraint has migrated downstream — out of the sensor and into the coordination.

The field seems to know this. The request GAO hears most often from practitioners is not for better cameras; it is for a single tool that fuses federal, state, local, and private data into one real-time view. And this is not a fringe position. GAO alone has produced a dozen reports and more than twenty recommendations on wildfire management over the past decade. The connective tissue, not the next sensor, is where the findings keep pointing.

What closing the gap looks like

Operationally, closing the gap is unglamorous and specific. A confirmed detection should flow automatically into a common operating picture that every responding entity sees the same way — same fire, same location, same confidence, same moment. From there, software should do the work software is good at: fuse the feeds, geolocate and characterize the fire, check it against airspace and other active incidents, and propose a coordinated response — which asset, sent where, sequenced how. A human then decides. The machines integrate and compute; people keep authority over what actually happens on the ground. That division of labor is the entire point. Closing the integration gap does not require inventing another way to see fires. It requires building the layer between seeing and acting — and that layer can be built now, from technology that already exists.

Sources

- U.S. Government Accountability Office, *Science & Tech Spotlight: Wildfire Detection Technologies* (GAO-25-108161), May 2025.
- U.S. Government Accountability Office, *Wildfire Management: Technologies for Forecasting, Detection, Mitigation, and Response* (GAO-25-108589), June 2025.
- A. Mohapatra & T. Trinh, "Early Wildfire Detection Technologies in Practice—A Review," *Sustainability* 14(19), 2022.
- R. Honary, J. Shelton & P. Kavehpour, "A Review of Technologies for the Early Detection of Wildfires," *ASME Open Journal of Engineering* 4, 2025.
- M. S. Huang & B. Wichmann, "Machine learning estimates on the impacts of detection times on wildfire suppression costs," *PLoS One* 19(11), 2024.

About Inari Watch

Inari Watch is a veteran-founded company building an integrated system of systems for regional wildfire defense. Our architecture unifies a distributed sensing mesh, an aerial fleet, and ground control through a layer of *bounded agentic coordination* — *AI at the edge* that perceives and proposes within strict policy limits while humans retain authority over every consequential action.