Edge Computing for Automated Hazard Alerting from Orbit

Sarvesh Garimella
MyRadar
Portland, OR USA
vesh@myradar.com

Douglas Franz
MyRadar
Orlando, FL USA
douglas.franz@myradar.com

Abstract— **This paper details recent advances in the use of small satellites for hazard alerting, specifically focusing on wildfire detection from Low Earth Orbit (LEO). The HORIS platform leverages hyperspectral imaging, onboard edge AI, and thermal/visible sensors to monitor the dynamic urban-wildland interface. The onboard processing system supports real-time detection, scene classification, and contextual calibration, enabling autonomous alert generation. We review scene classifier accuracy, reflectance calibration using cloud targets, aerosol model assignment, and onboard wildfire identification. Testing with a UAV platform validates spatial accuracy and sensor performance. These capabilities support future hazard monitoring constellations capable of near real-time global alerting.**

Keywords— CubeSat, wildfire detection, hyperspectral imaging, edge computing, AI, remote sensing, hazard alerting

I. Introduction

Traditional satellite-based remote sensing platforms for hazard monitoring—such as GOES, MODIS, VIIRS, and AVHRR—have made significant contributions to Earth observation. However, these missions have inherent tradeoffs in temporal and spatial resolution, as well as high cost and latency. In contrast, recent advances in CubeSats have opened the door for low-cost, frequent-return platforms capable of targeted environmental hazard monitoring from orbit. The Hyperspectral Orbital Remote Imaging Spectrometer (HORIS) mission, developed by MyRadar, is one such platform. HORIS leverages miniaturized hyperspectral sensors and onboard artificial intelligence (AI) to autonomously detect scenes of interest and generate alerts in near real-time. The mission specifically targets fire hazard detection at the wildland-urban interface (WUI), where conventional systems often fall short. Hyperspectral imaging enables precise surface material identification through the capture of dense spectral data across hundreds of contiguous bands. Filchev [1] and Jablonski et al. [2] describe hardware advances that have made compact hyperspectral imagers viable. Dillon [5] and NOAA's Hazard Mapping System offer critical context for fire hazard potential and current mapping limitations. Meanwhile, transfer learning and synthetic training data have gained traction for environmental remote sensing AI applications.

II. Methods

The HORIS CubeSat is a 1U-class spacecraft platform that integrates three primary sensors, including a thermal infrared (TIR) camera for radiance thresholding and hotspot detection, a visible spectrum camera for land surface imaging and albedo tracking, and a Red-NIR hyperspectral sensor for vegetation indices, aerosol environment, and burned area assessment.

The system is designed with modularity in mind, allowing sensor payload updates as technology evolves. Its orbit and design prioritize observation of dynamic hazard-prone zones. The wildfire detection methodology on HORIS enables fire detection through multiple onboard processes:

- Radiance Thresholding: Detects thermal anomalies corresponding to active fires
- Hyperspectral Analysis: Identifies combustion products and vegetation stress
- Scene Context Classification: Determines the type of land cover and aerosol model needed for correction
- Onboard Alert Logic: Triggers data capture and downlink sequences when specific thresholds are met

Real-time onboard scene filtering helps prioritize only relevant data, reducing downlink volume and enabling rapid alerts. To enable onboard AI and Edge Computing HORIS uses a three-tiered computing stack. A low-power always-on AI processor (~0.07W idle) performs basic scene-of-interest classification using a 6-layer ResNet-like CNN trained on ~23,000 labeled images. A main image processor performs bulk data capture and deeper analysis (~8.65W peak). A flight control processor manages telemetry, tasking, and system health. Aerosol scene classification (e.g., biomass burning, continental, maritime) supports radiative transfer corrections and spectral retrieval. Onboard lookup tables (LUTs) aid in selecting the correct atmospheric correction model. Onboard calibration and validation are performed using deep convective clouds (DCCs), which provide consistent VNIR reflectance across orbits. DCCs serve as ideal pseudo-invariant references, minimizing atmospheric and surface contamination. HORIS uses equatorial cloud fields below 30° latitude for this purpose. Additionally, airborne validation using fixed-wing test flights emulates orbital observation conditions. Contextual images and pixel masks aid in AI refinement and cross-validation of geolocation accuracy (~100m RMSE from WGS-84 baseline) [6].

III. Results

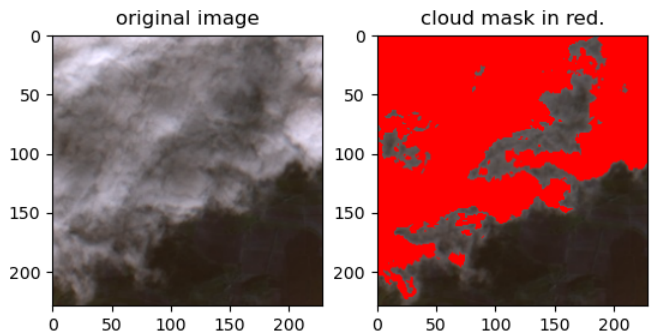

Fig. 1. Example calibration scene image (left) and cloud mask calibration pixels used for onboard LUT retrieval validations.

HORIS utilizes an onboard cloud masking scheme used for sensor calibration (Figure 1). This example compares a raw satellite image (left) with the corresponding segmentation mask (right), autonomously generated by the onboard AI to support reflectance retrieval calibration. Fig. 1 presents an example calibration scene image alongside the cloud mask calibration pixels used for onboard lookup table (LUT) retrieval validations, highlighting how specific cloud features are isolated to validate and adjust atmospheric correction parameters. These segmentation masks can distinguish between surface types such as vegetation, bare soil, water, clouds, and smoke—information that is crucial for accurate atmospheric correction and radiative transfer modeling. By identifying contextual features in each scene, the onboard AI dynamically adjusts the retrieval algorithms to apply surface- and atmosphere-specific corrections. This figure underscores how HORIS uses contextual awareness to self-calibrate in orbit, improving retrieval accuracy without manual intervention. Moreover, these labeled masks are also logged as training data, creating a feedback loop for continuous AI improvement and enabling large-scale generation of synthetic data with known context.

Figure 2 illustrates HORIS's onboard capability to autonomously detect potential biomass burning events using an power-efficient radiance thresholding method. This technique identifies thermally anomalous pixels—typically corresponding to active fires—by flagging regions where radiance values exceed calibrated thresholds. The figure shows a progression from raw thermal imagery to a binary hotspot mask, demonstrating how the AI filters and interprets sensor data in real time. Fire alerting can be performed within minutes of observation, demonstrating latency suitable for emergency alerting and underscoring the system's readiness for real-world deployment. This rapid response capability enables timely identification of dynamic environmental features such as wildfires, allowing HORIS to issue alerts without requiring data downlink or ground-based processing. Such low-latency detection is critical for time-sensitive applications like disaster response and air quality monitoring, and it exemplifies how edge AI systems can serve as autonomous watch towers from orbit.

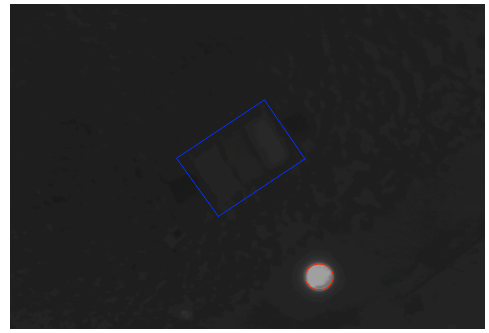

Fig. 2. Thermal infrared image collected during UAV arial characterization of sensor performance. Bright reflectance targets are indicated in blue and the red circle is the identfied hotspot from ground-truth warming fire target.

IV. Conclusions and Future Work

The HORIS platform demonstrates a compact, efficient model for autonomous hazard alerting using CubeSat-class systems (Figure 3). With its AI-enabled onboard computing, spectral sensors, and validated calibration, it is well-positioned to serve as a scalable node in future global wildfire monitoring constellations. Its lessons apply more broadly to remote sensing of dynamic, time-critical environmental hazards.

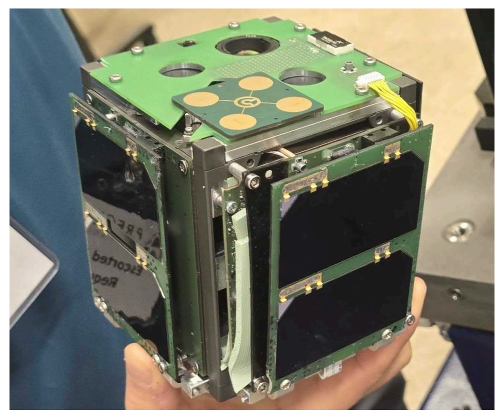

Fig. 3. HORIS pathfinder flight model prior to loading into deployer for vibrational testing.

Acknowledgment

We gratefully acknowledge funding from NOAA SBIR Award # NA23OAR0210337-T1-01.

References

[1] Filchev, L. (2014). Satellite Hyperspectral Earth Observation Missions – A Review. Aerospace Research in Bulgaria, 26.

[2] Jablonski, M., et al. (2016). Best practices in passive remote sensing VNIR hyperspectral system hardware calibrations. Proc. SPIE, 986004. https://doi.org/10.1117/12.2223941

[3] Gallagher, B., & St. Germain, K. (2018). The NOAA Satellite Observing System Architecture (NSOSA) Study.

[4] Anthes, R., & Schreiner, W. (2019). Six new satellites watch the atmosphere over Earth's equator. EOS, AGU, 100. https://eos.org

[5] Dillon, G. (2015). Wildfire Hazard Potential for the Conterminous US (270-m GRID), version 2014. U.S. Forest Service. https://doi.org/10.2737/RDS-2015-0047

[6] Garimella, S., & Franz, D. (2025). Satellite edge AI for automated training data collection, transfer learning, and synthetic training data generation. Proc. SPIE (in press).