

Cross-Sensor Portability of On-Board Wildfire Detection: Towards Cooperative Satellite Constellations

A. Ferrari¹, L. Manunza¹, C. Clementini¹, M. Maranesi¹, M. Frezza¹, P. Soresina¹, T. Scalia², A. D. Ciminelli²

¹GMATICS S.r.l., Rome, Italy ²Agenzia Spaziale Italiana (ASI), Rome, Italy

1 Introduction & Motivation

Wildfires evolve over minutes, and the acquisition-to-alert interval directly conditions emergency response.

- Conventional workflows downlink raw imagery before fire detection can begin, introducing delays incompatible with emergency response.
- Moving detection on board changes this paradigm, transmitting compact alerts and fire masks instead of full images, cutting downlink volume and RF power demand.
- Sub-hourly revisit is today achieved only at kilometre-scale sampling from geostationary satellites, too coarse to resolve small ignitions. A central design choice is therefore to **detect active fire without a dedicated thermal channel**, relying on SWIR/NIR bands common to multispectral and hyperspectral payloads: a model restricted to these channels can run on many general-purpose satellites, shortening effective revisit at high resolution.
- Because such payloads are general-purpose rather than fire-dedicated, **cross-sensor transferability becomes the enabling condition for cooperative multi-mission constellations**, in which a coordinating platform re-tasks heterogeneous assets according to priorities inferred from on-board detections. Transferability is quantified here on Sentinel-2 and Landsat-8.



2 Multi/Hyperspectral approach

We rely only on bands common to multispectral and hyperspectral payloads. Every general-purpose optical satellite becomes a candidate fire sensor, letting us leverage:


- A wide range of satellites: many already operational satellites share the near-infrared bands needed for this approach
- Short revisit time: relying on many satellites reduces the revisit time over the same area
- Cross-sensor portability: we test if the features of wildfires can be extracted so that any sensor with the same bands can provide detection.

3 Datasets

Two publicly available datasets are the starting points for our experiments: **SegTHRawS** [1] and **ActiveFire** [2]. The former provides Sentinel-2 while the latter Landsat-8 images. Both provide labels obtained from voting of a common set of algorithms.



4 Data Processing



5 Lightweight Segmentation Models

Residual U-Net (ResUNet)	Attention U-Net (ResUNet)
32 filters	32 filters
1,97 Million parameters	2,17 Million parameters
~7,9 MB (< 50 MB)	~8,7 MB (< 50 MB)
~1.5 s inference latency on CubeSat-compatible hardware	~1.5 s inference latency on CubeSat-compatible hardware
256x256 pixel input patch with 3 channels (SWIR2/SWIR1/NIR)	256x256 pixel input patch with 3 channels (SWIR2/SWIR1/NIR)

6 Experiments & Evaluation


Three categories of experiments are carried out:

- In-domain baseline:** Each model is trained and tested on the same dataset.
- Zero-shot transfer:** Each model is trained on a dataset but then tested on the other one.
- Dual sensor training:** Each model is trained on a combination of both datasets, then tested on one of them at a time.

Evaluation is performed considering the metrics **F1 score** and **IoU**

7 Datasets

The datasets we use for our experiments, in their final versions:



8 Results




Fig. 3 — In-domain baseline results, both models trained and tested on the same dataset.




Fig. 4 — Zero-shot transfer results, both trained models are tested in different datasets




Fig. 5 — Comparison among baseline, zero-shot and multi-source performances on S2 and L8h datasets

9 Conclusions & Next Steps

What we found:

- Multi-domain training demonstrated:** one model trained on S2 + L8h holds comparable accuracy on both sensors (0.905 on the combined test set), with no per-sensor model.
- Radiometric harmonisation:** the S2→L8 offset is additive and stable across biomes; one per-band correction recovers about a third of the zero-shot gap.
- Transfer is asymmetric:** L8→S2 is easier than the reverse; fire/background contrast and SWIR2 saturation, not the architecture, are the limit.
- On-board feasibility:** 1.97 M parameters (7.9 MB), ~1.5 s per scene at <30 W, within CubeSat-class budgets; attention gates gave no systematic gain.

Future improvements:

Data Augmentation + Model Quantization + PRISMA integration (hyperspectral data)



- [1] Castro Traba et al., 2026 — SegTHRawS
- [2] Pereira et al., 2021 — ActiveFire