Advancing Canadian wildfire technology through onboard processing and on the ground collaboration

Tim Heydrich[1], Andrew J. Macdonald[1], Tanya Harrison[1], Becca Bonham-Carter[1], Luis Chavier[1], Alexis Pascual[1], Shannon Paul[1], Saleh Abdelrahman[1], Scott Brass[2], Cordy Tymstra[3], Rob McAlpine[4], Eric Choi[1], Yolanda Brown[1]

[1] Mission Control, 162 Elm Street West, Ottawa, K1R 6N5, ON Canada – (tim, macdonald, tanya, becca, luis, alexis, shannon, saleh, eric, yolanda)@missioncontrolspace.com
[2] Eagle Flight Network, Manyhorses Drive, Tsuu T'ina Nation 145, T3Z 1A1, AB Canada - sbrass@eagleflight.network
[3] Whitebark & Sage Wildfire Science and Management, Edmonton, AB T6W 1E9 Canada – whitebarkandsage.wildfire@gmail.com
[4] Western University, 1151 Richmond Street, London ON N6A 3K7 – Rmcalpi4@uwo.ca

Keywords: Wildfire Detection, Onboard processing, Machine Learning, Satellites, Multidisciplinary Collaboration.

Abstract

Wildfire regimes are changing at a time when new satellite platforms can run machine learning analysis to detect and refine wildfire data products onboard, allowing explorations of new concepts of operations. Combining wildfire operations and science expertise with indigenous community knowledge and a technology demonstration satellite designed to deploy and test machine learning algorithms on orbit we characterise Fire Band Analysis Networks (FireBAN) models for wildfire detection from optical, multispectral, and hyperspectral data. Testing on a custom training dataset from Sentinel-2 and AVIRIS data prior to launch of the onboard models shows a test accuracy of 95.1% and mean intersection over union (mIoU) of 64.6% on Sentinel-2 and accuracy of 94.6% and mIoU of 72.6% for the best performing model architecture, which shows band-dependency effects under an ablation study. On-orbit performance demonstrates the ability to detect series of wildfire smoke plumes and gather data necessary to retrain models on the ground to counter sensor-based domain shifts and provide continuous updates of model weights to the satellite platform.

1. Introduction

The intensity and duration of wildland fires are growing in Canada (Jain *et al*., 2024; Hannes *et al.*, 2025; Tymstra and Flannigan, 2025). This behaviour is projected to increase globally in the years to come (Haas *et al.*, 2026), and end-users require better information faster to safeguard human health, infrastructure, and the natural environment (Sankey, 2018). Modern wildfire management in the Canadian landscape is complex and involves a wide variety of stakeholders and technical approaches. Space-based and aerial imaging payloads can provide wide swath monitoring and key scientific and operations data but for wildfire response the latency involved with turning this data into fire intelligence products can limit its potential. Early fire detection is critical to control and suppression, leading to fewer escaped fires and impacts (Johnston *et al.,* 2020). For fires that have escaped initial suppression efforts, if consolidated and derived products from aerial and satellite platforms can make it to the ground operations centre in under two hours then this provides the best chance of informing tactical decisions by fire managers. Recent advancements in using machine learning techniques on embedded platforms for disaster management, such as satellites deployed with onboard flood detection neural network models (Mateo-Garcia *et al.*, 2021), introduce new technological approaches for bringing processed wildfire detections and science data to relevant users at speed.

2. Related Work

2.1 Earth Observation for Wildfire Data

Government and commercial earth observation (EO) satellites and aircraft acquire petabytes of data with varying revisit and spatial/spectral resolution. Flagship programs such as Landsat or the EU’s Copernicus program operate multiple earth observation platforms while commercial providers, like Planet Labs, Vantor (formerly Maxar), Black Sky, and others provide satellite earth observation data products as part of their core business and companies like Ororatech and Muon Space target wildfire data products specifically. EO data is both platform and sensor specific, with examples like ESA’s Sentinel-2, NASA’s Landsat-8, ASI’s PRISMA, NASA’s Moderate Resolution Imaging Spectroradiometer (MODIS) aboard the Aqua and Terra satellites, or the Visible Infrared Imaging Radiometer Suite (VIIRS) aboard S-NPP. Space-based EO data can be combined with data from aerial assets such as NASA’s Airborne Visible/Infrared Imaging Spectrometer (AVIRIS). All these assets have different sensors and different data product outputs that have advantages and disadvantages for informing pre-suppression preparedness, detection, response and recovery. For example, Sentinel-2 with its MultiSpectral Instrument (MSI) captures 13 image bands, AVIRIS captures 224 imaging bands at high spatial resolution but has low global coverage. Landsat-8, a popular choice for wildfire detection work, has a multis-spectral Near Infrared (NIR) and Short-Wave Infrared (SWIR) imager as part of its Thermal Infrared Sensor (TIRS) instruments with lower spectral resolution than AVIRIS but higher global coverage.

Data from these sensors is then used as a resource for web platforms like NASA’s Fire Information for Resource Management System (FIRMS), the Canadian Interagency Forest Fire Centre (CIFFC) Canadian Wildland Fire Information System (CWFIS), the EU Copernicus European Forest Fire Information System (EFFIS), the Government of Western Australia’s MyFireWatch service, and others. These platforms, refine and combine the raw sensor data acquired onboard space or aerial assets and are used by both wildfire professionals and members of the public. Latency between detection by a sensor and information appearing on one of these platforms varies by geolocation, for example, NASA FIRMS targets three-hour latency for global coverage, but this varies by region, exact alignment between orbit and ground station network, and is dependent on government funding status.

https://doi.org/10.5194/isprs-archives-XLIX-B3-2026-1387-2026

2.2 Machine Learning Wildfire Applications on Ground

Machine learning techniques applied to earth observation data have two notable differentiating aspects compared to many terrestrial problems in computer vision. The data is complex, often being generated by custom sensors that can include multispectral or hyperspectral bands that span the visible to long wave thermal infrared (LWIR). The scarcity and complexity of ground truth labels necessary for supervised learning is particularly difficult to confirm for satellite data, where it often must be inferred or verified via a ground-based data source.

Arrue (Arrue *et al.*, 2000) pioneered the use of artificial neural networks for infrared image processing in conjunction with other data sources to identify wildfires. Since then, the architecture design has shifted to focus on Convolutional Neural Networks (CNNs). In Canada both Natural Resources Canada (NRCan) and the BC Wildfire Services are exploring active fire detection using CNN segmentation architectures (Crowley M. A., 2023). Most recent applications have trained CNN models on terrestrial images of fire and smoke (Sousa *et al.,* 2019**;** Barmpoutis *et al.,* 2019). Notably, Zhang (Zhang *et al.*, 2018**)** demonstrated superior performance of CNNs compared to SVM-based methods, and Yuan (Yuan *et al.*, 2018) incorporated optical flow for time-dependent information. Li et al. (Li *et al.*, 2018) utilized 3D CNNs for smoke detection, treating it as a video segmentation problem. Thangavel's (Thangavel *et al.*, 2023) work introduced several key innovations, including the first use of deep learning methods on hyperspectral data for wildfire detection. They developed a 1D-CNN architecture for detecting wildfires using the 234 bands of PRISMA hyperspectral data. The model was evaluated on three different hardware accelerators and demonstrated both time and power-efficiency, highlighting the use cases for AI-on-the-edge and potential solutions to technical challenges in the deployment.

2.3 Machine Learning Wildfire Applications on Orbit

Due to the latency bottleneck that exists between the ground and flight segment, there has been increasing interest in deploying machine learning wildfire detection algorithms directly onto the satellite to improve the time it takes for end-users to get the wildfire data. One such example is Orora Technology's FOREST-1 launched in 2022 (Spichtinger A., 2023), equipped with LWIR, MWIR, and RGB sensors and an NVIDIA Jetson that could run CNN-based fire detection onboard and provide wildfire detection notifications within 3-10 minutes instead of the 3+ hours for global ground-based systems like FIRMS. Traba (Traba *et al.*, 2026.) have shown success in onboard thermal hotspots segmentation with raw multispectral satellite imagery on representative satellite hardware. Ingesting raw Sentinel-2 NIR and SWIR data they show that a fully convolutional network can achieve a performance of Intersection over Union (IoU) of 0.988 and an F-1 score of 0.986 running in 1.45 seconds on Cubesat equivalent hardware. FireSat, led by the Earth Fire Alliance aims to use onboard processing techniques to identify hot spots 5 m x 5 m in size.

This evolution in machine learning onboard processing techniques for satellites is occurring across a larger set of application areas, including in disaster response scenarios like flood detection. Mateo-Garcia *et al.* (Mateo-Garcia *et al.*, 2021) performed in-orbit flood mapping using an onboard segmentation approach, allowing for reduced data transmission needs during downlink and demonstrating the first use of onboard deep learning for disaster response. Broadly there are now a class of satellites, such as PhiSat-2 (Melega, 2023), CogniSat-6 (Rijlaarsdam *et al*., 2025), Intuition-1 (Nalepa, 2021), and Persistence (Bonham-Carter *et al.*, 2025) specifically designed to test, characterize, and validate the kind of machine learning models and infrastructure needed to achieve robust, effective, and reliable performance for on-orbit tasks from vessel detection to land cover classification.

2.4 User-Centred Approaches to Wildfire Management

Centring users and communities is a necessary step to inform the choice of solutions and tools necessary in a changing fire regime. Tymstra (Tymstra *et al.,* 2020) note, in their review on wildfire management challenges and opportunities in Canada, that fire presents a 'double paradox' where fire suppression contributes to a fire problem that can be addressed with fire use. This leads to the need for adjusted wildfire management capacity, a greater degree of wildfire on the landscape, and the need for more resilient community practices and design, particularly at the wildland urban interface (WUI). Programs such as FireSmart, used by fire agencies across Canada, Australia, and New Zealand since 1999 and organizations like the Wildfire Resilience Consortium of Canada (WRCC) take a holistic approach that focus on empowering individuals and providing education, technology, and practices from indigenous and place-based traditional and historical knowledge.

Within the context of centring stakeholder groups and communities affected by wildfire as the primary drivers of solutions and technical requirements, machine learning-based onboard processing allows the derived value from large amounts of raw data to be moved between orbit and users on the ground within the time constraints required for making wildfire management decisions. This thereby makes it one of the three pillars in the overall solution shown in Figure 1. Johnston (Johnston *et al.,* 2020) describe the user requirements approach for the WildFireSat mission, starting with a Wildfire Management Needs Assessment that surveyed provincial and wildfire management authorities. The survey results were used to inform the Fire Management Functionalities for the mission, and within the context of tactical intelligence and decision-making data products were required as quickly as possible, preferably within thirty minutes or less, with a cut-off threshold of two hours.

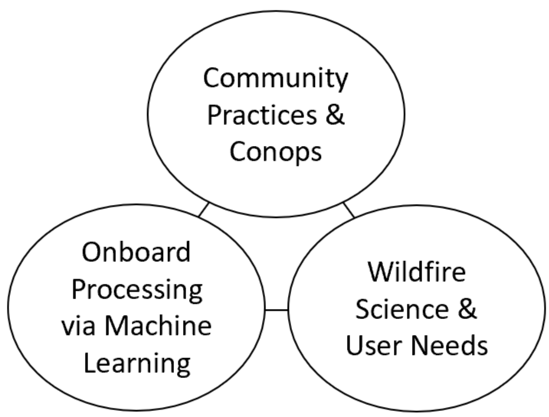

Figure 1. Three pillars to the project approach to use onboard machine learning to provide value to communities and users.

3. Method

3.1 Satellite Platform

Mission Persistence, launched on Transporter-14 on June 23, 2025, is designed to deploy, test, and maintain machine learning payloads in orbit. The flight segment consists of a 6U cubesat built by Spire Global Inc. and it hosts a Dragonfly Gecko RGB imaging payload and Xilinx Ultrascale+ multiprocessor system-on-a-chip (MPSoC) flight computer. Onboard it contains a flight software payload that runs on this MPSoC and hosts multiple neural network model architectures which can run forward inference passes onboard. The ground and flight segment of the mission are designed to measure covariate shift and to enable domain adaptation via in-flight model updates. The updates are performed using a ground segment MLOps pipeline, designed to make neural network models more robust and resilient to covariate shift over time. (Bonham-Carter *et al.*, 2025)

The mission is designed to support a multitude of different experiments, primarily ones focusing on running and updating neural networks in flight. This makes it a valuable platform for technology demonstrations of onboard Fire Band Analysis Network (FireBAN) models that are trained on existing EO datasets and then updated on orbit throughout operations. This concept of operations allows performance improvements and redefinitions of data products through continuous dialogue with end users – who get to see real on-orbit data products early and an ability to request changes. At launch there were two FireBAN model architectures, both trained for semantic segmentation, onboard the satellite.

3.2 End-User and Community Engagement

With work related to wildfire detection engaging the end users, especially the communities most affected, is essential to ensure that the solution developed addresses real and existing needs. To this end, the consortium is informed by wildfire science and operations experts who have decades of experience in the wildfire sector. Their expertise guided the development of the concept of operations and provided expert input on data product structure that would add benefit to the current structure employed by the various wildfire agencies across Canada. These agencies operate at the provincial and territorial levels and are coordinated nationally through CIFFC.

At the community level, Eagle Flight Network, as an Indigenous partner, leads efforts in indigenous community engagement and the integration of Indigenous Traditional Knowledge (ITK). The result is an approach grounded in Stewardship and Sovereignty. Their engagement strategy is built on a foundation of active listening. This approach prioritizes a deep understanding of local wildfire concerns and the operational realities of detection, prevention, and response within each community. Engagement focuses on identifying the practical limitations faced by remote regions, including assessing the information currently available to the public and community leadership to identify critical gaps in awareness and accessibility. By centring dialogue on lived experience rather than predefined technical solutions, system requirements are directly informed by frontline needs. Technical design is therefore anchored in the operational realities of remote and northern communities. Communities consistently emphasized the importance of early warning, timely delivery of information, and system reliability in areas with limited ground infrastructure and heavy reliance on satellite and aviation services. All engagement activities were undertaken with attention to Free, Prior, and Informed Consent (FPIC) and OCAP® (Ownership, Control, Access, and Possession) principles. Clear processes are established to ensure that community input directly informs system design and operational planning. This approach is guided by the principle: Stewardship Before Signal.

3.3 Dataset Creation

3.3.1 Dataset Sources

Sentinel-2 and AVIRIS were selected as primary sources for pre-launch training datasets. Using the Sentinel-2 archives we cross-listed imaging dates against known fire records to construct a multi-spectral dataset of wildfires, a subset appears in Figure 2. To increase FireBAN robustness to changes in resolution and spectral information we complemented the Sentinel-2 dataset foundation, with a second dataset created from AVIRIS hyperspectral data cubes, a first step in planning for a sensor agnostic approach.

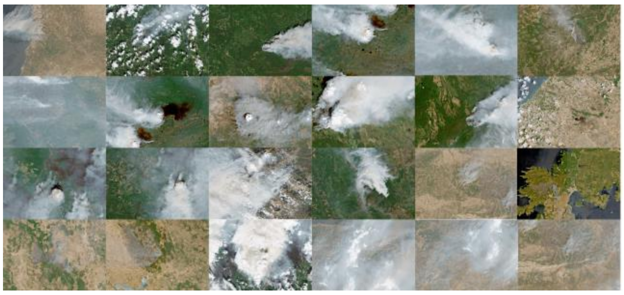

Figure 2. A mosaic of active wildfires throughout the world, captured using Sentinel-2.

Data from both sources was ingested into a machine learning operations pipeline (MLOps) that enabled data cleaning, labelling, model training and evaluation, and eventually repeated updates onboard. Preparing and cleaning the data involved determining which bands would be most relevant from each data source for wildfire detection. Infrared NIR/SWIR spectrums are crucial for this and therefore we initially focused on analysing channels 150 onwards from the AVIRIS data as these correspond to the 1800-2500mm SWIR range. Similarly, we used Band 12 from Sentinel-2 data as source for our labelling efforts.

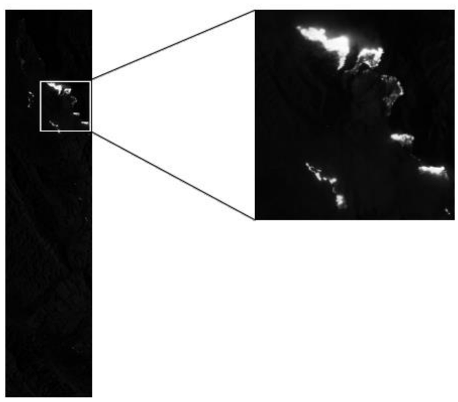

Figure 3. A composite image of the Southern California Fires on December 07, 2017, created using AVIRIS, bands 150-224.

3.3.2 Labelling the Data

The down-selected bands were analysed for their pixel intensity across all bands and to produce segmentation masks. Using the maximum intensity for each band we were able to create a greyscale image clearly showing the mask, see Figure 3 for a sample greyscale AVIRIS image. We then performed histogram analysis of the grey-scale images from AVIRIS and the band 12 images from Sentinel-2 to determine pixel intensity corresponding to fires in the data to serve as ground truth, a sample distribution can be seen in Figure 4.

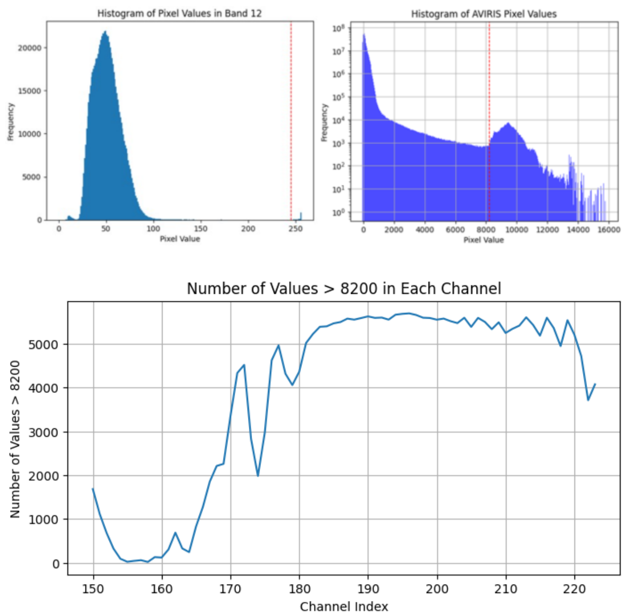

Figure 4. Sample Pixel Intensity Distribution for Sentinel-2 data (top left) and AVIRIS data (top right). Sample frequency distribution of pixel values across the different AVIRIS bands (bottom).

From the histograms we can see a clear spike in the frequency of pixels with an intensity of above 8200 for AVIRIS data or above 245 for the Sentinel-2 data. We use the intensity threshold of 8200 to further analyse the importance of each band for locating the fire. We conducted a frequency count of values surpassing the threshold for each channel, as seen in Figure 4, to further reduce the number of bands used from AVIRIS to be bands 170 onwards. We then used a basic binary threshold to create the ground truth mask for each data source. For the Sentinel-2 images we created a mask where each pixel with an intensity above 245 was set to 1 and all other to 2. The input data for the Sentinel-2 dataset was created from Band 4, 11, and 12. For the AVIRIS data the mask was created using the threshold of 8200 for setting the pixels to 1, with the input being the data cube of bands 170 onwards. Both datasets served the input and masks with width and height of 448 pixels. We collected, cleaned and labelled 136 images for the Sentinel-2 dataset and 33 images for the AVIRIS dataset.

3.4 Model Training

3.4.1 Model Selection

We selected two model architectures for training, the first being FPENet which is an extremely lightweight network architecture allowing for fast on-the-edge model inference (Liu and Yin, 2021). We chose this architecture for its ability to meet the size, weight, memory, and power requirements of a cubesat flight computer and because of its flight heritage in space missions (Cross *et al.,* 2023**)**. The second architecture was a more modern lightweight model called SwiftNet (Wang *et al.,* 2021) which was selected due to its state-of-the-art performance on various terrestrial segmentation benchmarks. Both model architectures were adapted to be agnostic to the number of input channels to support various data sources, RGB, multi- and hyper-spectral.

3.4.2 Training Setup

Given the unique problem of wildfire detection using multi- and hyper-spectral data, the datasets are relatively small, which can easily lead to degraded performance. To mitigate this, we employed standard data augmentation methods to improve the model’s robustness. The models were trained using cross-entropy-loss and the RAdam optimizer. Training was done over 100 epochs but with an early stopping mechanism enabled. Training was done using a Nvidia RTX 4090. The data used was split using an 80-10-10, train-validation-test split.

3.4.3 Optimization for Flight

The initial models deployed to the satellite were deployed pre-flight and therefore did not have any optimization done that would reduce model size or resource requirements. The MLOps pipeline supports performing model optimizations such as pruning and quantization as well as partial model retraining. These optimizations can reduce the model update size allowing for faster onboard satellite model updates. Quantization can also reduce the resource requirements of the model and further increase the inference speed of these lightweight architectures (Gholami, 2022**)**.

4. Results

4.1 Sentinel and AVIRIS Results

We measure performance on a binary classification task between “Fire” and “No Fire” pixels in the test set. When assessing the model performance, we primarily focus on the intersection-over-union (IoU) and not the accuracy as the accuracy is skewed due to the data mostly consisting of “No Fire” areas but our interest is to correctly identify the areas with fire. Table 1 shows the results of both architectures on the Sentinel-2 dataset.

Dataset	Input Bands	Model	Test Accuracy	Test IoU
Sentinel-2	4, 11, 12	FPENet	0.73999	0.37791
Sentinel-2	4, 11, 12	SwiftNet	0.95079	0.64933

Table 1. FPENet and SwiftNet Performance on Sentinel-2 Test Set

From the Sentinel-2 results we can see that FPENet underperforms SwiftNet’s ability to reliably infer pixels containing wildfire. This is likely a combination of the extreme lightweight nature of FPENet as well as the limited data available for model training. In contrast SwiftNet performs much better on the Sentinel-2 data and was therefore selected to also be trained on the more complex hyperspectral AVIRIS data. Figure 5 shows a sample output of SwiftNet predictions compared to the ground truth labels created.

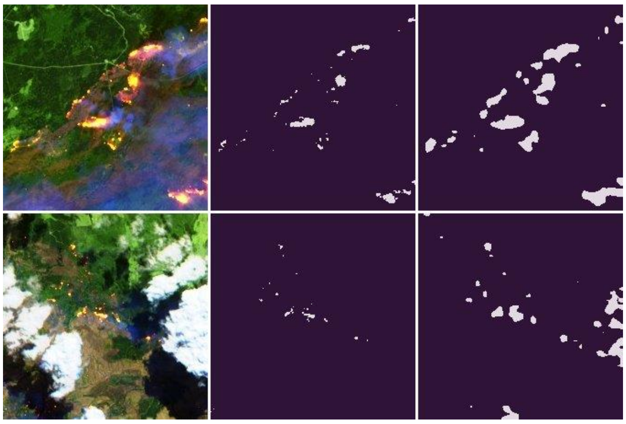

Figure 5. Sentinel-2 sample predictions, Input (left), Ground Truth (middle), SwiftNet Prediction (right).

The predictions show a smoothed fire mask which is to be expected from a segmentation neural network when compared to the jagged ground truth labels created by the binary thresholding. Some additional analysis is needed to determine if the smoothed predictions can be entirely attributed to the model performance or if the thresholding method used to create the labels underrepresents the data and the machine learning model is able to infer some non-linearity as seen in Figure 6.

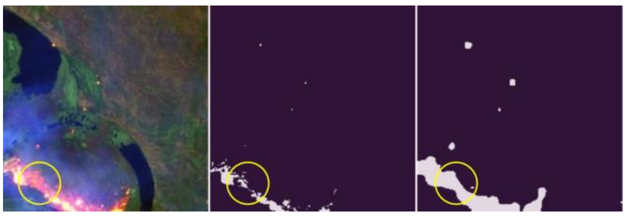

Figure 6. AVIRIS false colour (left), ground truth (middle), FireBAN prediction (right). Notice the sparse nature of the ground truth in contrast to the false colour and SwiftNet prediction.

When training SwiftNet on the AVIRIS data we performed an ablation study to ensure that there are no confounding variables influencing the model performance. This study analysed the impact of spectral bands on model performance. Specifically, we investigated performance variations under three scenarios: using all 224 bands of the hyperspectral data, using the top 34 most correlated bands, 190-224 (see Figure 4) and exploring the efficacy of four distinct three-band subsets derived from the correlated data. Performance across the different band combinations explored are itemized in Table 2.

Dataset	Input Bands	Test Accuracy	Test IoU
AVIRIS	All	0.97244	0.64607
AVIRIS	190-224	0.96756	0.67708
AVIRIS	222-224	0.9441	0.55217
AVIRIS	210-212	0.95736	0.74964
AVIRIS	200-202	0.94045	0.54725
AVIRIS	190-192	0.9464	0.7258

Table 2. Ablation study on different spectral bands when using the SwiftNet backbone on the AVIRIS dataset.

When looking at the three-band subsets, one might assume a trend of performance as we progress into higher bands; however, the results reveal inconsistencies in IoU. These performance differences likely arise from disparities across bands, attributed to the unique attenuation characteristics of each band and the variability in null values within the bands signal. A more interesting finding is that FireBAN can effectively propagate all bands into a sufficient segmentation mask, without requiring a domain expert to use dimensionality reduction techniques to remove irrelevant spectral bands. Focussing on correlated bands yields a 5% improvement in relative performance, while mitigating the risks associated with using a smaller selection of bands.

4.2 Operational Wildfire Imaging

The sensor platform onboard the test satellite is most sensitive to wildfire smoke signatures, which are verified against hotspot detections from platforms like CWFIS, FIRMS and MyFireWatch. Images are captured in bursts of ten images at 1 Hz frequency. an example of a wildfire smoke plume signature in images collected onboard in December 2025 over northeastern Australia is shown in Figure 7.

5. Future Work

5.1 Collecting and Verifying Training Data On-Orbit

Currently the validated wildfire data from satellite operations is limited but already demonstrates a sensor induced domain gap

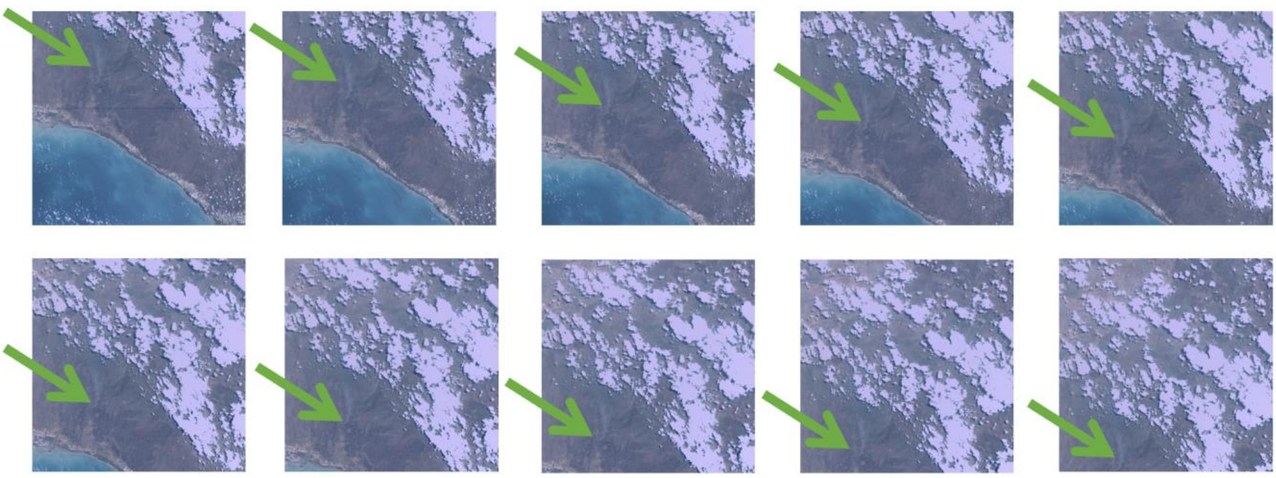

Figure 7. Tasking acquisition over northeastern Australia showing a series of ten images, each with a wildfire smoke plume signature as indicated by the green arrow.

https://doi.org/10.5194/isprs-archives-XLIX-B3-2026-1387-2026

between the multi- and hyper-spectral data pre-launch training data and RGB images captured by the platform post-launch. This is expected and provides a test case for the use of MLOps to update the model architectures in flight based on the latest training data, requiring an adaptation of the labelling approach.

The current approach detailed in section 3.4.3 using pixel intensity threshold is limited in the presence of clouds or thick smoke. NIR and SWIR cannot penetrate thicker clouds or smoke and those will produce pixels of similar intensity as the IR detected from the fire. A more refined approach for filtering out clouds and other atmospheric artifacts is inspired by the work of Singh (Singh, *et al.*, 2025**)**, where Landsat-8 images were used as a baseline for a better labelling approach for Sentinel-2 data. Based on this method we utilized both SWIR and the NIR band, 11, 12 and 8 respectively, from Sentinel-2 to create a preliminary fire-line data product. Figure 8, shows the overall improved labelling approach. This fire-line product is then compared to the aerosol band, 1, and any white pixels that appear in both are classified as clouds since the IR cannot penetrate through those thicker clouds. All other data in the aerosol band is then classified as smoke; however, this is refined by only counting connected swaths of aerosols as smoke if they are near a fire location. The fire-line product is filtered for orange-red pixels to determine the fire source locations.

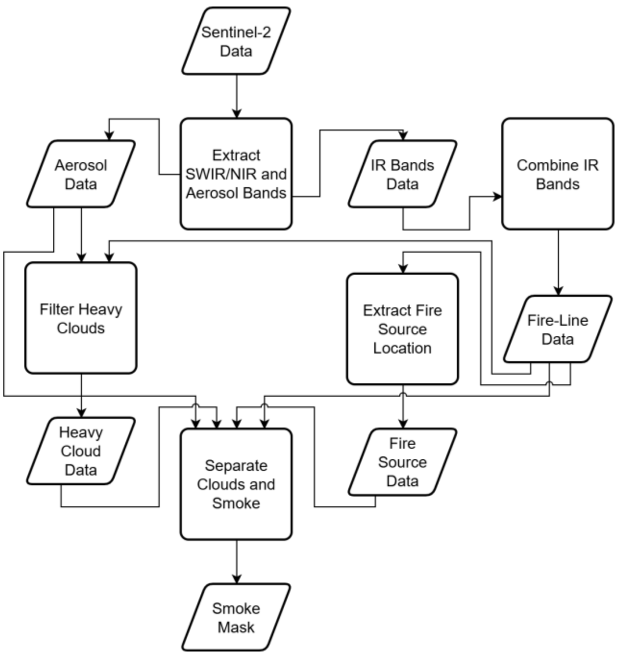

Figure 8. Improved Sentinel-2 data labelling pipeline.

This labelling approach creates several different intermediate data products all of which can be used to train a FireBAN model depending on the use case. A sample of the main data products (i.e. the fire line and the smoke mask) of this labelling approach can be seen in Figure 9. There are still further improvements to be made to this algorithm as currently, thick heavy smoke is also classified as clouds as it will appear in both the aerosol as well as the fire-line data products.

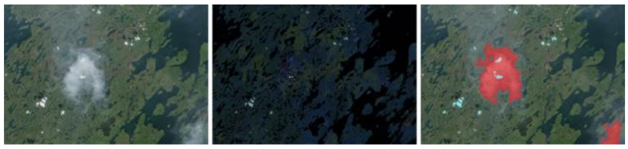

Figure 9. Main data products output from improved labelling approach.

5.2 Ground Segment Data Processing

Any images captured by the satellite are primarily downlinked unprocessed, any data products produced onboard by various machine learning models are also downlinked along with the raw data. The data downlinked then needs to be colour corrected. In addition, several applications utilizing earth observation images require them to be georeferenced. Therefore, we are currently in the process of adding a georeferencing pipeline to our existing data pipeline. Figure 10 shows the outputs of the data processing steps. This georeferencing step would also significantly improve the value our fire detection models generate by ensuring accurate latitude and longitude of any fire detected onboard by FireBAN models.

Figure 10. Original Image with wildfire (left), colour corrected image (middle), georeferenced image overlayed on Sentinel-2 image (right).

6. Conclusion

We demonstrate a concept for developing onboard machine learning satellite capabilities in the context of wildfire user needs and community practices. Future efforts will focus on upgrading onboard models and performing demonstrations with community stakeholders to characterize the utility of data products generated onboard the satellite instead of on the ground, forming experimental benchmarking to inform new satellite mission concepts of operations, including novel observing systems that coordinate satellite fire monitoring via heterogeneous constellations (Chien *et al.*, 2025).

Acknowledgements

The authors acknowledge that this program/project/activity is undertaken with the financial support of the Canadian Space Agency (CSA).

References

Arrue, B.C., *et al.*, 2000. An intelligent system for false alarm reduction in infrared forest-fire detection. *IEEE Intelligent Systems*, 15(3):64–73, doi: 10.1109/5254.846287.

Barmpoutis, P., *et al.,* 2019. Fire Detection from Images Using Faster R-CNN and Multidimensional Texture Analysis. In *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, volume 2019-may, pages 8301 8305. Institute of Electrical and Electronics Engineers Inc. doi: 10.1109/ICASSP.2019.8682647

Bonham-Carter, B. *et al.*, 2025. Advancing AI Reliability for Space Missions through Machine Learning Operations, *Proceedings of the 18th International Conference on Space Operations (SpaceOps 2025)*, Montreal Canada.

Chien, S., *et al.,* 2025. *Multi-Asset New Observing Systems Flight Demonstration*, Proceedings of the 18th International Conference on Space Operations (SpaceOps 2025), Montreal Canada 26-30 May 2025.

Cross, M., *et al.,* 2023. Enabling Autonomy and Operations for Lunar Surface Missions: An Overview of Demonstrated Capabilities. *Conference Proceedings of the 17th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*, Scheltema, Leiden (The Netherlands)

Crowley, M. A., 2023. Canada's WildFireSat Mission. ITU Webinar on "Fighting wildfires with AI-powered insights" NRCan April 19th, 2023. https://www.itu.int/en/ITU-T/webinars/20230419/Pages/default.aspx

Gholami, A., 2022. A Survey of Quantization Methods for Efficient Neural Network Inference. *Low Power Computer Vision* 1st Edition Imprint Chapman and Hall/CRC. https://doi.org/10.1201/9781003162810-13

Haas, O., *et al.*, 2026. Wildfires on a changing planet. *Nat Commun* 17, 1599 (2026). https://doi.org/10.1038/s41467-025-68176-4

Hannes, C., *et al.*, 2025. Fire regime changes in Canada: an update. *Canadian Journal of Forest Research.* **55**: 1-11. https://doi.org/10.1139/cjfr-2025-0209

Jain, P., *et al.,* 2024. Drivers and Impacts of the Record-Breaking 2023 Wildfire Season in Canada. *Nat Commun* **15**, 6764. https://doi.org/10.1038/s41467-024-51154-7

Johnston, J., *et al.*, 2020. Development of the User Requirements for the Canadian WildFireSat Satellite Mission. Sensors (Basel). 20(18):5081. doi: 10.3390/s20185081

Li, X., *et al*., 2018. 3D Parallel Fully Convolutional Networks for Real-time Video Wildfire Smoke Detection, in *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 1, pp. 89-103, Jan. 2020, doi: 10.1109/TCSVT.2018.2889193.

Liu, M., Yin, H., 2021. Feature Pyramid Encoding Network for Real-time Semantic Segmentation, *Image and Vision Computing* Volume 111, 104195 https://doi.org/10.1016/j.imavis.2021.104195.

Mateo-Garcia, G., *et al.,* 2021. Towards global flood mapping onboard low cost satellites with machine learning. *Sci Rep* 11, 7249. https://doi.org/10.1038/s41598-021-86650-z

Melega, N., 2023**.** The ESA Φ-Sat-2 mission: An AI enhanced multispectral cubesat for earth observation, in *Proc. AIAA/USU Conf. Small Satellites*, in Next Pad - Res. Academia, SSC23-WIV-09, 2023.

Nalepa, J., 2021. Towards on-board hyperspectral satellite image segmentation: Understanding robustness of deep learning through simulating acquisition conditions, *Remote Sens.*, vol. 13, no. 8, Art. no. 1532. https://doi.org/10.3390/rs13081532

Rijlaarsdam, D., *et al*., 2025. The Next Era for Earth Observation Spacecraft: An Overview of CogniSAT-6," in *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 18, pp. 2450-2463, doi: 10.1109/JSTARS.2024.3509734.

Sankey, S., 2018. Blueprint for wildland fire science in Canada (2019-2029). *Publication of the Canadian Forest Service and Natural Resources Canada*. Catalogue No.: Fo134-12/2018E-PDF

Singh, H., *et al.*, 2025. Active wildfire detection via satellite imagery and machine learning: an empirical investigation of Australian wildfires. *Nat Hazards* 121, 9777–9800. https://doi.org/10.1007/s11069-025-07163-w

Sousa, M. J., *et al.* 2019, Wildfire detection using transfer learning on augmented datasets. *Expert Systems with Applications*, page 112975, sep 2019. doi: 10.1016/j.eswa.2019.112975

Spichtinger, A., 2023. On-Board Wildfire Detection. ITU Webinar on "Fighting wildfires with AI-powered insights" Orora Technologies. April 19th, 2023. https://www.itu.int/en/ITU-T/webinars/20230419/Pages/default.aspx

Thangavel, K., *et al.,* 2023. Autonomous Satellite Wildfire Detection Using Hyperspectral Imagery and Neural Networks: A Case Study on Australian Wildfire. *Remote Sensing*, 15(3), 720.

Traba, C., *et al.*, 2026. Towards onboard thermal hotspots segmentation with raw multispectral satellite imagery. *International Journal of Applied Earth Observation and Geoinformation* Volume 146, 105095. https://doi.org/10.1016/j.jag.2026.105095

Tymstra, C., *et al.,* 2020. Wildfire management in Canada: Review, challenges and opportunities. *Progress in Disaster. Science* Volume 5, 100045 https://doi.org/10.1016/j.pdisas.2019.100045

Tymstra, C., Flannigan, M., 2025. Canada's Landscape Fires Spreading into Uncharted Territory. Environ. Res. *Lett*. 20 (2025) 111003 doi: 10.1088/1748-9326/ae1495

Wang, H. *et al.*, 2021. SwiftNet: Real-time Video Object Segmentation, 2021 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Nashville, TN, USA, 2021, pp. 1296-1305, doi: 10.1109/CVPR46437.2021.00135.

Yuan, J., *et al.*, 2018. Detection of Wildfires along Transmission Lines Using Deep Time and Space Features. *Pattern Recognition and Image Analysis,* 28(4):805–812, oct 2018. doi: 10.1134/S1054661818040168.

Zhang, Q. X., *et al.*, 2018. Wildland Forest Fire Smoke Detection Based on Faster R-CNN using Synthetic Smoke Images. In *Procedia Engineering*, volume 211, pages 441–446. Elsevier Ltd, 2018b. doi: 10.1016/j.proeng.2017.12.034.