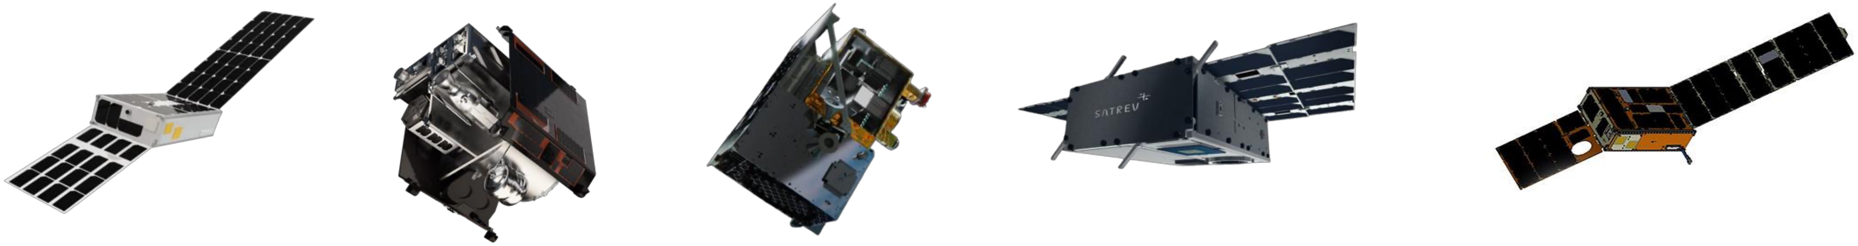

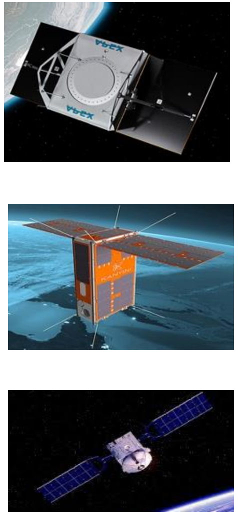

Rapid Multi-Mission Deployment of Convolutional Neural Network and Spectral Algorithm Flight Software

Itai Zilberstein[1], Alberto Candela[1], Steve Chien[1],
David Rijlaarsdam[2], Tom Hendrix[2], Leonie Buckley[2], Aubrey Dunne[2],
Vishesh Vatsal[3], Adithya Kothandhapani[3], Arvind Subramanian[3], Harmit Vyas[3], Abhinav Jayaswal[3], and Vijay Singh Purohit[3]

[1]Jet Propulsion Laboratory, California Institute of Technology, USA
[3]Ubotica Technologies, Ireland
[3]Hyspace Technologies (Skyserve), India

Jet Propulsion Laboratory
California Institute of Technology

Portions of this research were carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). This work was supported by the Earth Science and Technology Office (ESTO), NASA

Overview

- JPL AI Group partnering with new space companies to demonstrate rapid deployment of data analysis FSW to multiple LEO spacecraft with AI hardware

Satellite: CogniSAT-6
Provider: Ubotica/Open Cosmos
Status: Flown and executed

Satellite: ION SCV-004
Provider: Hyspace/D-Orbit
Status: Flown and executed

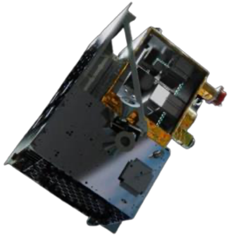

Satellite: YAM-6
Provider: Hyspace/Loft Orbital
Status: Delivered to provider

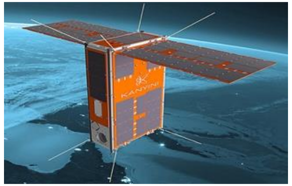

Satellite: Kanyini-1
Provider: SmartsatCRC
Status: Planned

Satellite: Aries SN1
Provider: Ubotica/Apex
Status: Planned

Satellite: SOWA-1
Provider: SatRev/Hyspace
Status: Planned

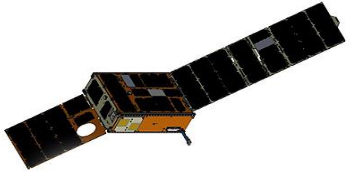

Satellite: Phi-Sat-2
Provider: ESA/Open Cosmos
Status: Planned

Satellite: Crypto-2
Provider: Aptos Orbital
Status: Planned

Motivation: Why do data analysis at the edge?

1. Data insights for rapid response
    - Volcanoes, floods, wildfires, algal blooms, ….
    - Intersatellite-links (ISL) for near-instantaneous alert

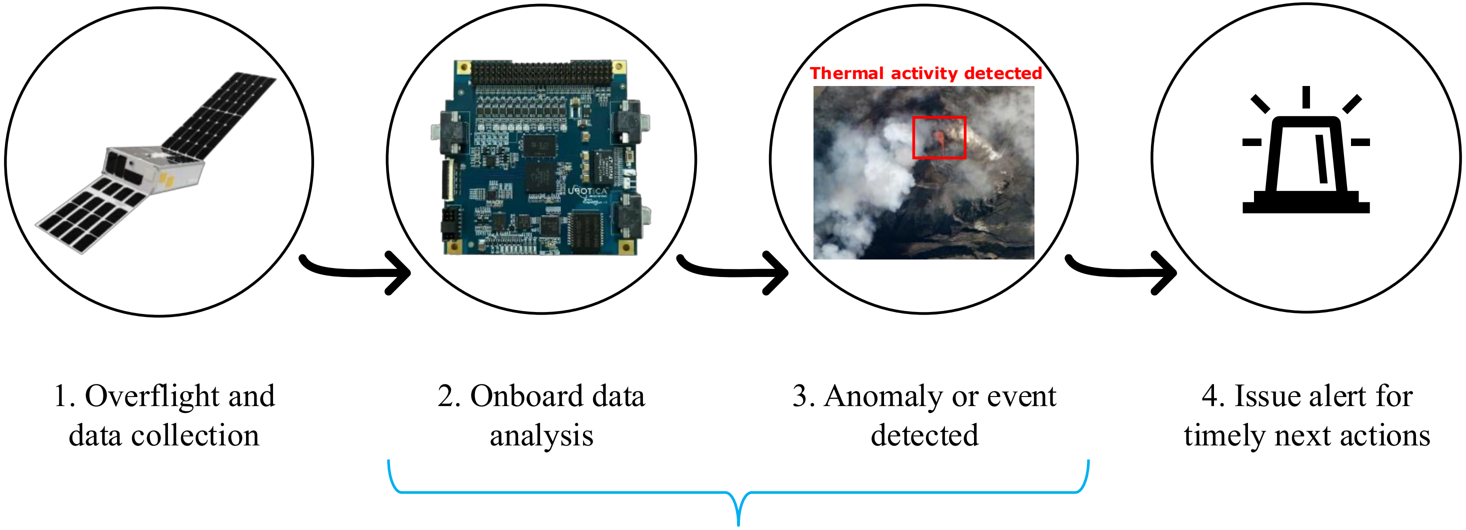

1. Overflight and data collection

2. Onboard data analysis

3. Anomaly or event detected

4. Issue alert for timely next actions

Focus of this work

Motivation: Leveraging knowledge onboard

2. Optimization of resources and observations

- Reactive:
  - Acquire data and reject if bad
  - Save storage and downlink

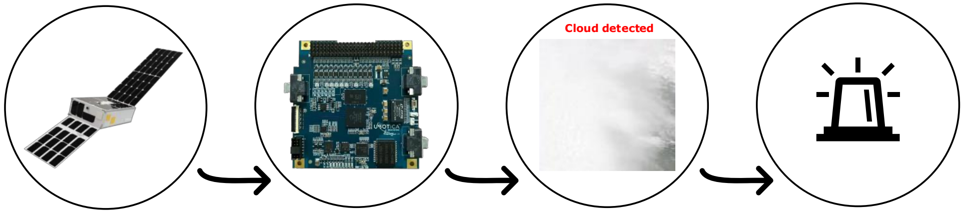

1. Overflight and data collection
2. Onboard data analysis
3. Unusable data identified
4. Offload data and/or alert

- Proactive (*Dynamic targeting*):
  - Look-ahead data
  - Saves:
    - Storage
    - Downlink
    - Time/energy (slewing)

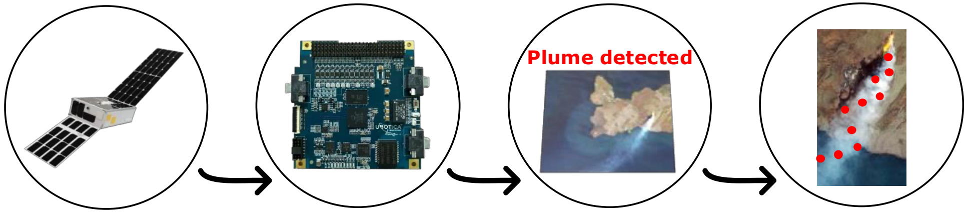

1. Collection of lookahead data
2. Onboard data analysis
3. Event detection
4. Collection of precise data

What’s been done at JPL before?

- AI FSW in operation:
    - EO-1, AVIRIS/EMIT, IPEX, ASTERIA, AEGIS, M2020, CADRE …
    - First flight of ML from AIG - ASE 2004
    - Extensive development, verification, and operation cost
- Automated *ground* workflows for rapid response (sensorwebs)
    - Download data, analyze, and trigger follow-up actions

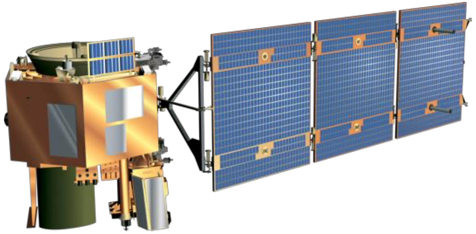

EO-1

Perseverance Rover

CADRE

Sensorweb concept

AI Hardware in Space

- VPUs for computer vision, image signal processing, CNN execution
- ISL for persistent comms
- Increasingly capable CPUs
- Increase in RAM

Intel Myriad 2

Intel Myriad X

NVIDIA Jetson TX2i

Software Capabilities

- Increasingly friendly FSW environments
- Availability of modern software libraries

GDAL

scikit learn

ONNX

TensorFlow

CONDA

Keras

OpenVINO

pandas

OpenCV

NumPy

Rapid Development and Deployment

- **Challenge**: each spacecraft has varied
  - Instruments/data products
  - Software versions and edge hardware
  - RAM, uplink, and other data volume limitations
- Also, different science applications for same spacecraft

Input data → Preprocessing → Model execution

RGB Snapshot Camera @ 50m GSD
vs
Hyperspectral Pushbroom @ 5m GSD

Python 2.7 vs Python 3.11

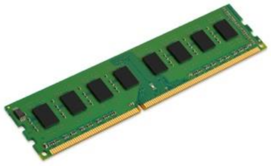

50 MB RAM vs 2 GB RAM

Myriad X vs Myriad 2

Rapid Development and Deployment

- **Solution**: highly (re)configurable software

Parameterized Core Software Suite

TensorFlow

{JSON}

CONDA

Tailored Software for Specific Spacecraft

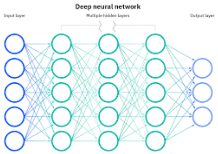

V&V and Deployment

Spacecraft Specific Information

YAML

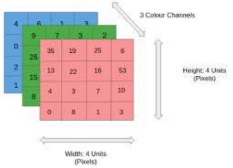

35	19	25	6
13	22	16	53
4	3	7	10
0	8	1	3

Core Data Analysis Software Suite

- CNNs and Spectral Algorithms
  - Parameterized based on
    - Specific satellite (e.g. available operations, preprocessing, file formats)
    - Input data (e.g. dimensions, bit depth, resolution)
    - Application (e.g. cloud detection, surface water mapping)
- Memory-safe pre/postprocessing scripts
  - Tilers, scene statistics, band alignment
  - Limited dependencies

Output:

Algorithm ready for onboard integration

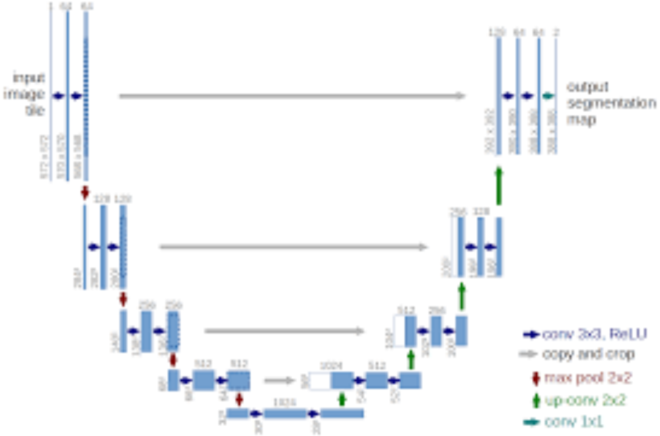

Input:

Configuration file
Training data (e.g. labeled scenes, spectral signatures)

Training, compilation, conversion

Convolutional Neural Networks

- Example:
    - CNNs for semantic segmentation
        - Clouds, surface water, thermal activity, algal blooms, + more
    - Developed 2 U-Net Architectures: Xception and UAVSAR
        - Fixed model size of ~4 MB
- Huge space of CNNs architectures and applications to deploy onboard

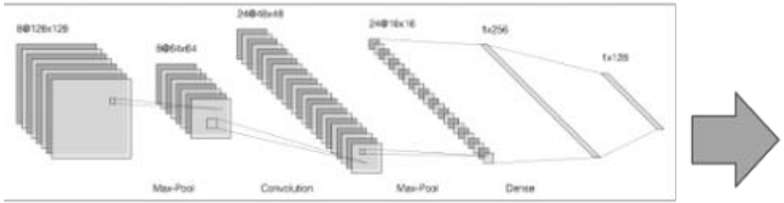

Deep Learning
Convolutional Neural Network

Object Detection
example: volcanic eruption

Image Segmentation
example: cloud mask

Spectral Algorithms

- Signature detection and unmixing of high dimensional data
- Algorithms:
  - Spectral angle mapper (SAM)
  - Matched filters (MF)
  - Reed-Xiaoli anomaly detector (RX)
- Engineered to leverage AI hardware onboard

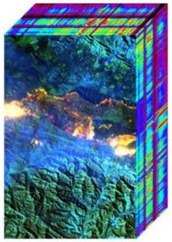

False color image from AVIRIS data cube acquired on Dec 5 2017. The front face shows: (red) active fires at 2250 nm; (green) surface at 1650 nm; and (blue) smoke at 1000 nm.

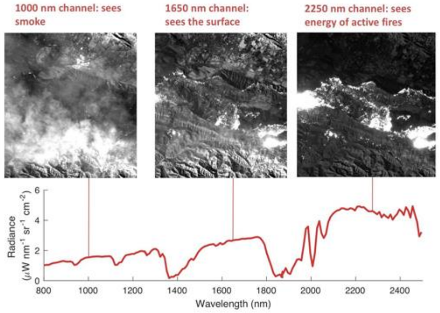

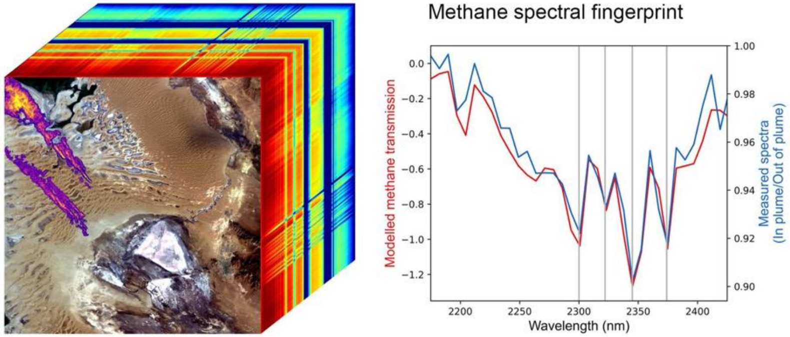

CogniSAT-6 Completed Demonstrations

- CNNs for cloud, surface water, thermal activity
- 2 models for each application, 30+ total executions

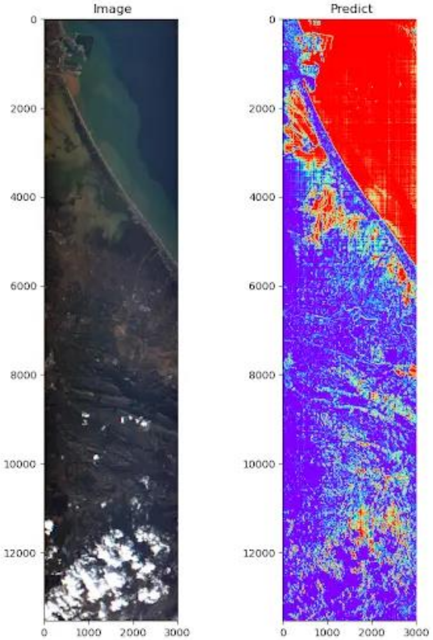

11/02/24:

Observed 21% of area near Valencia was flooded

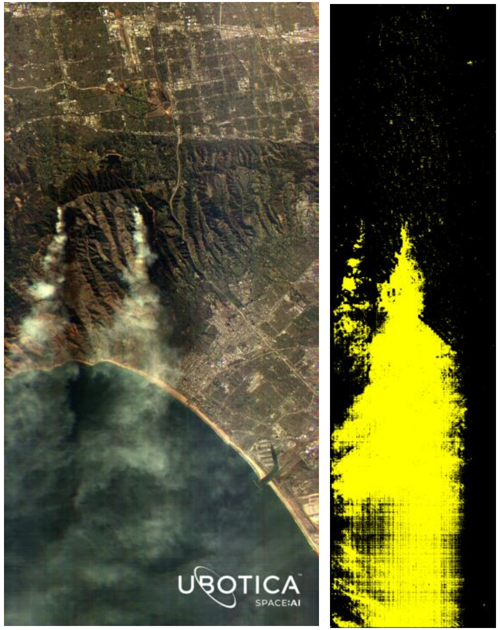

01/11/25:

Classified active plumes from Palisades fire

Includes imagery from CogniSAT-6/HAMMER, 2025, Ubotica. All rights reserved.

CogniSAT-6 Current Status

- Spectral Algorithms flight in Spring 2025

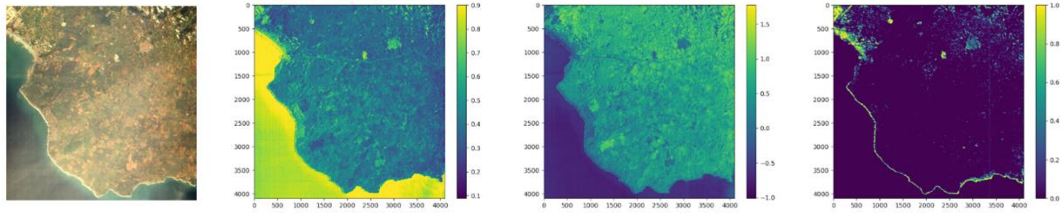

HyperScape Scene

SAM: detection is smaller value

MF: detection is larger value

RX: anomaly is larger value

Vegetation mapping via spectral algorithms. *Includes imagery from CogniSAT-6/HAMMER, 2025, Open Cosmos Limited. All rights reserved.*

- Dynamic targeting demonstration in Spring/Summer 2025

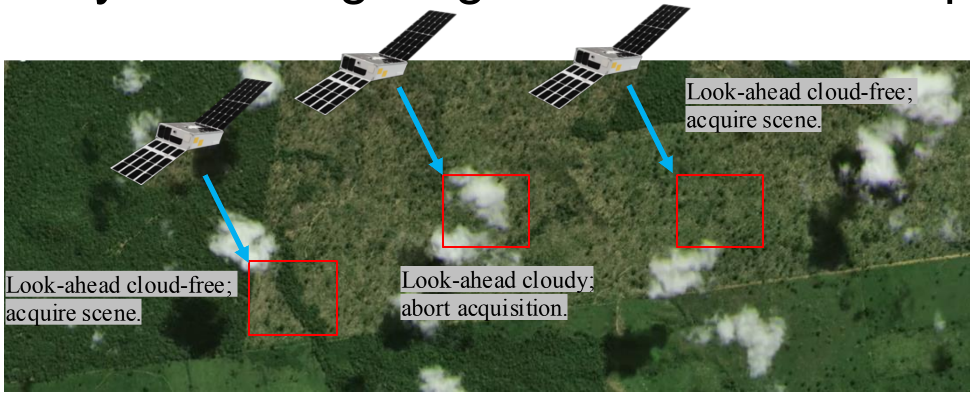

Look-ahead cloud-free; acquire scene.

Look-ahead cloudy; abort acquisition.

Look-ahead cloud-free; acquire scene.

CogniSAT-6 will actively avoid taking cloudy observations by slewing its sensor forward and analyzing a look-ahead image prior to near nadir acquisition

ION SCV-004 and YAM-6 Current Status

- Software development completed – dozens of flight applications
    - CNNs and spectral algorithms for clouds, water, vegetation, urban detection
- As of 2/10/25; execution of 9 models (spectral & CNN) on ION SCV 004
- More flights in spring 2025

ION SCV-004 Xception Cloud Classifier (onboard execution)

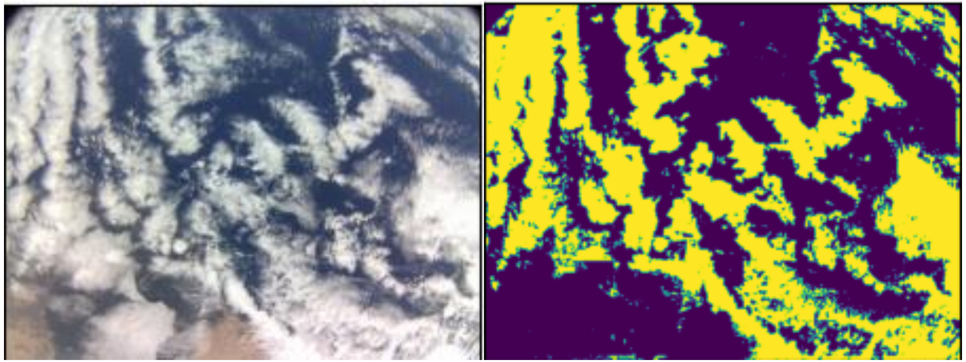

ION SCV-004 RX Anomaly (onboard execution)

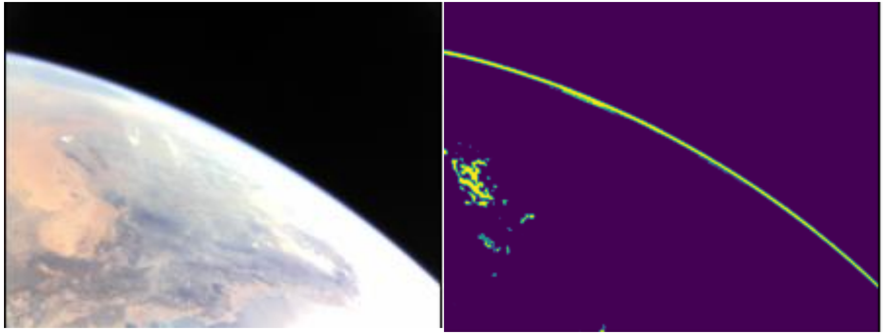

Upcoming Demonstrations

- Deployment to more spacecraft part of NASA/ESTO

Satellite: Aries SN1
Provider: Ubotica/Apex

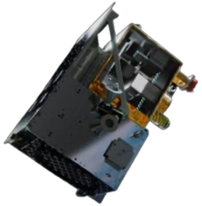

Satellite: YAM-6
Provider: Hyspace/Loft Orbital

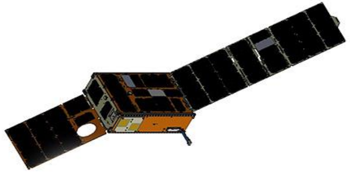

Satellite: Phi-Sat-2
Provider: ESA/Open Cosmos

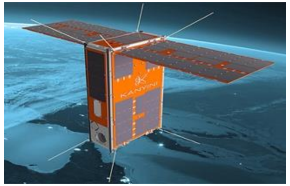

Satellite: Kanyini-1
Provider: SmartsatCRC

Satellite: Crypto-2
Provider: Aptos Orbital

+ more….

- FAME:
  - Largest in-space demonstration of AI
  - 50+ spacecraft involved
  - Demonstrate multi-asset coordination

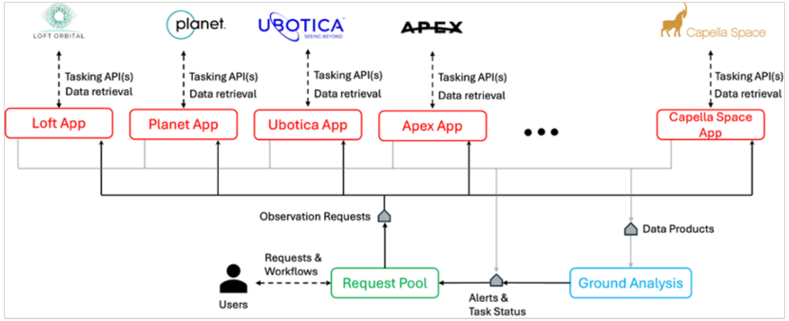

Tasking API(s)
Data retrieval

Loft App

Planet App

Ubotica App

Apex App

•••

Capella Space App

Observation Requests

Data Products

Users

Requests & Workflows

Request Pool

Alerts & Task Status

Ground Analysis

A Future Vision of FSW for Data Analysis

- Leverage existing languages & libraries for rapid development
- End-users (e.g. scientists, consumers) create workflows for data analysis
  - Plug-and-play Jupyter notebooks


Artist's Concept

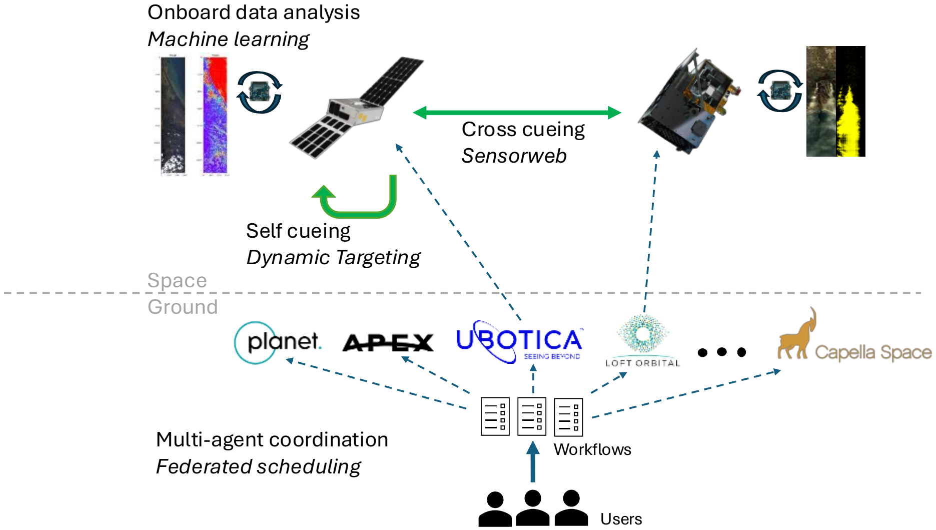
Onboard data analysis
Machine learning

Cross cueing
Sensorweb

Self cueing
Dynamic Targeting

Space
Ground

Multi-agent coordination
Federated scheduling

Workflows

Users

Conclusion

- Onboard edge processing for image processing improving drastically
- Demonstrating rapid development and deployment of CNN & Spectral Algorithm FSW to numerous spacecraft
- Future of data analysis FSW can be agile, efficient, and innovative and improve the return of Earth-observing satellites and other spacecraft

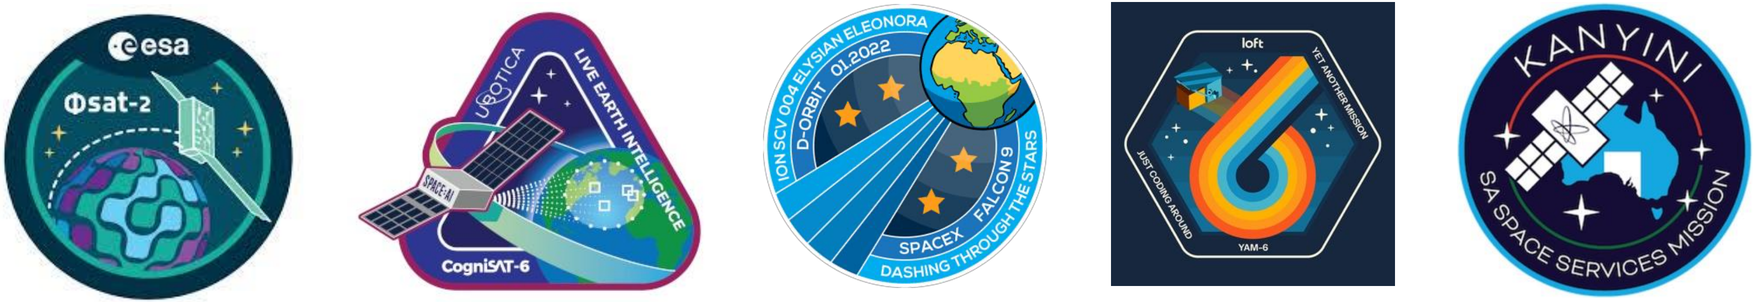

Jet Propulsion Laboratory
California Institute of Technology

jpl.nasa.gov

Portions of this research were carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). This work was supported by the Earth Science and Technology Office (ESTO), NASA

Rapid Multi-Mission Deployment of Convolutional Neural Network and Spectral Algorithm Flight Software

Itai Zilberstein*[1], Alberto Candela[1], Steve Chien[1], David Rijlaarsdam[2], Tom Hendrix[2], Léonie Buckley[2], Aubrey Dunne[2], Vishesh Vatsal[3], Adithya Kothandhapani[3], Arvind Subramanian[3], Harmit Vyas[3], Abhinav Jayaswal[3], and Vijay Singh Purohit[3]

[1]Jet Propulsion Laboratory, California Institute of Technology, CA, USA
[2]Ubotica Technologies, Ireland
[3]Hyspace Technoglogies (SkyServe), India

1 Abstract

The Jet Propulsion Laboratory is deploying onboard machine learning and spectral analysis capabilities to numerous spacecraft. These deployments leverage edge AI hardware for rapid analysis and insight to reduce reaction times. One such spacecraft is CogniSAT-6/HAMMER (CS-6) (Ubotica/Open Cosmos) which carries a Myriad X Vision Processing Unit (VPU) [1] for edge computer vision, image signal processing, and neural network execution [6]. Additionally, in collaboration with SkyServe.ai we are deploying to the ION SCV-004 spacecraft (D-Orbit) (Myriad X VPU) and YAM-6 spacecraft (Loft) (Jetson TX2i [2]) and also ground-comparable regular CPUs. These spacecraft provide virtual environments for deployment of languages and libraries such as Python, NumPy, and Tensorflow.

By analyzing data onboard, spacecraft are able to rapidly obtain knowledge from data - enabling rapid response to detected phenomena and reduction in data volume. We perform this onboard inference through spectral signature detection and image segmentation using CNN and other techniques. Image analysis consists of semantic segmentation using adaptations of the Xception and UAVSAR models [5] (both U-Net [7] deep CNN architecture). These models are tailored for deployment on flight hardware by ensuring the feasibility of operations, reducing model size, and embedding preprossessing operations to reduce CPU computation. We engineer spectral algorithms such as the Spectral Angle Mapper, matched filters, and the Reed-Xioali anomaly detector to leverage the AI acceleration onboard, a novel approach to deploying these algorithms. We identify operations feasible of executing on AI accelerated hardware and wrap these as neural network operations. We target numerous Earth science applications ranging from the detection of clouds and volcanic activity to flood and surface water mapping as well as land-use classification.

Development begins with the engineering of algorithms within the constraints of flight hardware. We require software that is memory-safe and efficient while performing operations on gigabytes of data. The models are amenable to the specifics of different instruments such as dimensions, bit depth, and wavelengths. Through stretching and interpolation, we can calibrate data products across instruments to train CNNs on larger datasets. When executing models, we perform these preprocessing operations onboard as part of the application. These applications are then validated on ground hardware prior to flight. As of October 2024, the first in-orbit executions of these models have successfully completed on CS-6 [9]. Flight demonstrations on ION SCV-004 and YAM-6 are expected in the fall of 2024. We are in earlier stages with an additional four spacecraft: Aries SN1 (Ubotica/APEX), Kanyini-1 (SmartsatCRC), SOWA (SatRev, Hyspace), and Phi-Sat-2 (ESA).

In future demonstrations, the knowledge obtained onboard will be leveraged to perform intelligent observations using *dynamic targeting* [3, 8] and *multi-asset coordination* [10]. In dynamic targeting data from a look-ahead view is analyzed in real-time to identify targets of interest that drive near-nadir observations. In multi-asset coordination, knowledge from one spacecraft is communicated to another one to drive observations across the networked spacecraft. Flight demonstrations of dynamic targeting are in development for CS-6 which has the ability to point forward to obtain a look-ahead view and communicate via inter-satellite links [4]. Through software advancements, we can improve the return of Earth-observing satellites.

---

*Contact: itai.m.zilberstein@jpl.nasa.gov.

Acknowledgments

Portions of this work were performed by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). This work was supported by the NASA Earth Science and Technology Office (ESTO). Government sponsorship acknowledged.

References

[1] Intel Movidius Myriad X Vision Processing Unit. Link, accessed on 2024-10.

[2] NVIDIA Jetson TX2. Link, accessed on 2024-10.

[3] A. Candela, J. Swope, and S. Chien. Dynamic targeting to improve Earth science missions. *Journal of Aerospace Information Systems*, 20(11):679–689, 2023.

[4] Steve Chien, Itai Zilberstein, Alberto Candela, David Rijlaarsdam, Tom Hendrix, Aubrey Dunne, Aragon Oriol, and Miquel Juan Puig. Flight of dynamic targeting on the CogniSAT-6 spacecraft. In *International Symposium on Artificial Intelligence, Robotics and Automation in Space*, November 2024.

[5] E. Dunkel et al. Benchmarking deep learning models on Myriad and Snapdragon processors for space applications. *Journal of Aerospace Information Systems*, 20(10):660–674, 2023.

[6] David Rijlaarsdam, Tom Hendrix, Pablo T Toledano González, Alberto Velasco-Mata, Léonie Buckley, Juan Puig Miquel, Oriol Aragon Casaled, and Aubrey Dunne. The next era for Earth observation spacecraft: An overview of CogniSAT-6. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2024.

[7] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015.

[8] Hiroshi Suto, Fumie Kataoka, Nobuhiro Kikuchi, Robert O Knuteson, Andre Butz, Markus Haun, Henry Buijs, Kei Shiomi, Hiroko Imai, and Akihiko Kuze. Thermal and near-infrared sensor for carbon observation Fourier transform spectrometer-2 (TANSO-FTS-2) on the Greenhouse gases Observing SATellite-2 (GOSAT-2) during its first year in orbit. *Atmospheric Measurement Techniques*, 14(3):2013–2039, 2021.

[9] Itai Zilberstein, Alberto Candela, Steve Chien, David Rijlaarsdam, Tom Hendrix, Leonie Buckley, and Aubrey Dunne. Demonstrating onboard inference for Earth science applications with spectral analysis algorithms and deep learning. In *International Symposium on Artificial Intelligence, Robotics and Automation in Space*, November 2024.

[10] Itai Zilberstein, Ananya Rao, Matthew Salis, and Steve Chien. Decentralized, decomposition-based observation scheduling for a large-scale satellite constellation. *Journal of Artificial Intelligence Research*, 82:169–208, 2025.