Jet Propulsion Laboratory
California Institute of Technology

Benchmarking Deep Learning On a Myriad X Processor Onboard the International Space Station (ISS)

Léonie Buckley (Ubotica) and Emily Dunkel (JPL)

Flight Software Workshop 2022

Onboard Deployment of Deep Learning Models

- Future space missions will need more powerful onboard autonomy to meet mission and science objectives
  - Time delay between earth and space system can be prohibitive to mission goals
  - Deep learning is state of the art in computer vision [DL]
  - Can commercial off the self computer systems withstand harsh space environments?
- Objective: Demonstrate deployment of deep learning to a Myriad X processor on HPE’s Spaceborne Computer-2 payload onboard the ISS

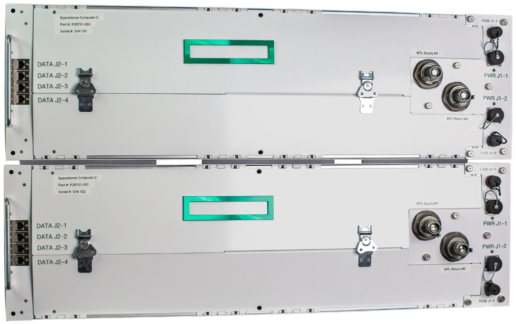

HPE’s Spaceborne Computer-2

Myriad X

Team

- Ubotica
  - Léonie Buckley
  - Dr. Jose Luis Espinosa-Aranda
  - Juan Romero-Cañas
  - Elena Hervas-Martin
  - ML algorithms: Elena Hervas-Martin
- Jet Propulsion Laboratory, California Institute of Technology
  - Dr. Emily Dunkel
  - Jason Swope
  - Dr. Zaid Towfic
  - Dr. Steve Chien
  - Dr. Damon Russell
  - Dr. Joseph Sauvageau
  - Dr. Douglas Sheldon
  - ML algorithms: Dr. Kiri Wagstaff, Steven Lu, Dr. Emily Dunkel, Dr. Michael Denbina, Dr. Zaid Towfic
- Hewlett Packard Enterprise
  - Dr. Mark Fernandez
  - Carrie Knox

Outline

- Myriad X Processor and Model Porting
- ISS Experiment
- Machine Learning Models and Myriad Results
- Prior Flights: Machine Learning Inference Onboard
- Future Work and Conclusions
- References

Myriad X Processor

16 VLIW SHAVE vector cores
2 Deep neural network processing units: NCE
Dedicated Imaging and Vision hardware accelerators

Up to 4 TOPS
1 TOPS of dedicated Neural Networks Compute

Software reconfigurable hardware–software platform

Ideally suited for in-orbit
Low thermal budget
Low power budget (~1.5 W)

Previous generation Myriad 2 flew on ESA's Phi-Sat [ESA]

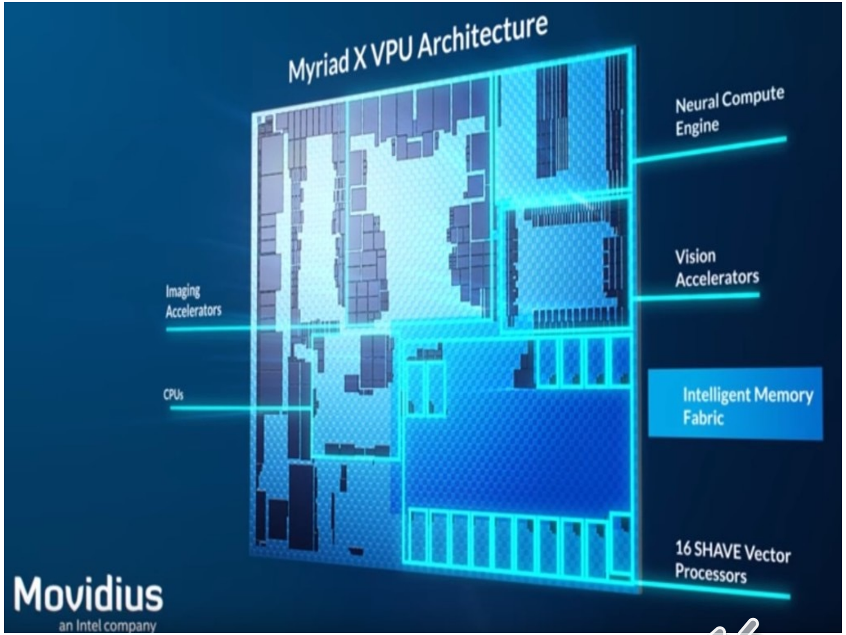

Credit: Intel Movidius

Porting the models to Myriad X

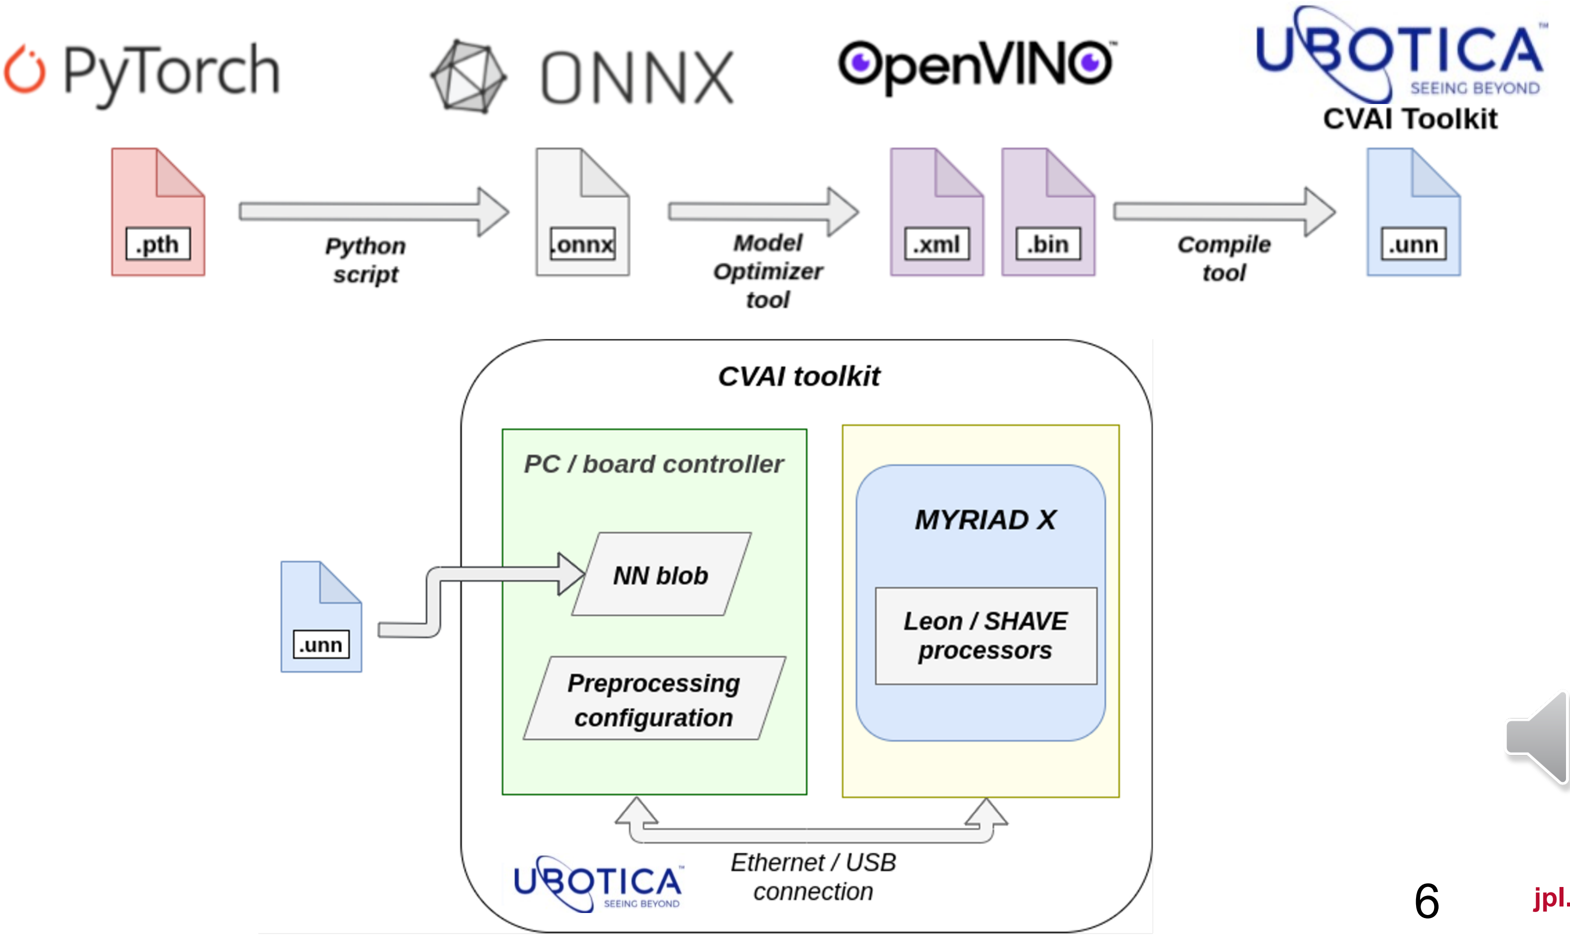

Jet Propulsion Laboratory
California Institute of Technology

ISS Experiment

Overview

- Objective: Demonstrate in space deep learning inference for future missions
  - Onboard processing could allow for autonomous targeted data collections and downloads, and decision making

- ISS Experiment
  - HPE's Spaceborne Computer-2 (SBC-2) delivered to the ISS onboard Cygnus NG-15: February 20, 2021
    - Powered on: May 12th 2021
  - Two Intel Myriad X VPU's have been delivered and integrated with SCB-2
    - Hosting of Myriad processors is enabled by HPE's SBC-2
    - A set of deep learning algorithms have been run onboard the ISS
  - Two Qualcomm Snapdragon 855 HDK's have also been delivered and integrated with SBC-2
    - See presentation by Jason Swope (JPL)

- Sponsored by the Earth Science Technology Office (ESTO)

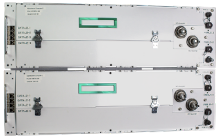

SBC-2

+

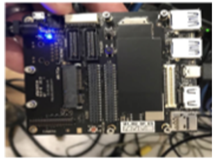

Snapdragon

+

Myriad

All SBC-2 images credit HPE

Experiment Setup

1. JPL + Ubotica Develops and Tests

HPE:
Chippewa Falls, WI

HPE Ground Testbed

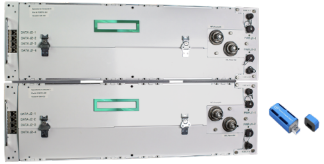

2. JPL Tests

HPE Flight Testbed (Ground)

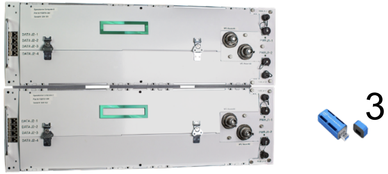

3. HPE Tests

4. HPE Tests

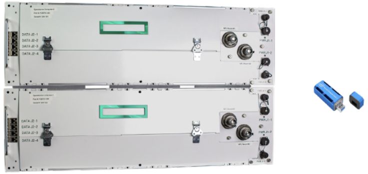

SBC-02 on ISS

5. Results → JPL

Jet Propulsion Laboratory
California Institute of Technology

Machine Learning Models and Results

Machine Learning Models Ported and Run on ISS

Objective: Deploy machine learning models for inference on a Myriad X onboard HPE's Spaceborne Computer-2 onboard the ISS

We will show a range of the models that can be deployed

- Mars HiRISE Classifier
- Mars MSL Classifiers (2)
- UAVSAR Flood Mapping
- Ship Segmentation
- DDR Test

Mars HiRISE: Deep Learning for Landmark Classification

- What: Mars surface feature classification using orbital imagery
- Applications: autonomous data collection, targeted downloads, commanding space assets, science interpretation
- Details:
  - Data: Imagery collected by the HiRISE instrument onboard the Mars Reconnaissance Orbiter (MRO) [PDS]
  - Model: Transfer learning from AlexNet [Mars1, Mars2, DL]

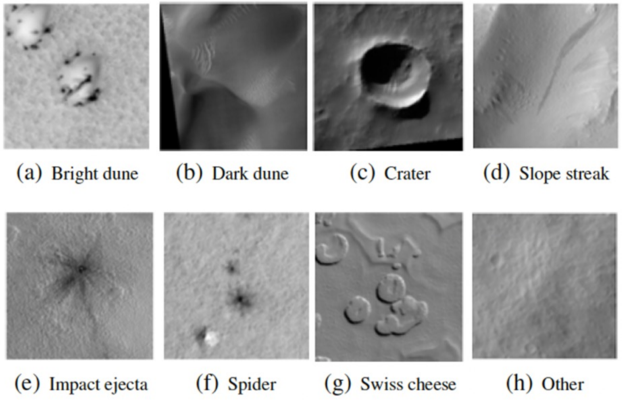

Example of HiRISE Classes

Mars MSL1: Deep Learning for MSL Imagery

- What: Mars rover image classification
- Applications: onboard navigation, science interpretation
- Details:
  - Data: Collected by the Mast Camera and Mars Hand Lens Imager instruments on MSL Curiosity rover [PDS]
  - Model: Transfer learning from AlexNet [Mars1, Mars2, DL]

Drill Hole

Sun

Mars MSL 2: Deep Learning for Object Classification

- What: Rover part classification
- Applications: spacecraft health analysis
- Details:
  - Second classifier in a chain that classifies images from MSL Curiosity rover [PDS]

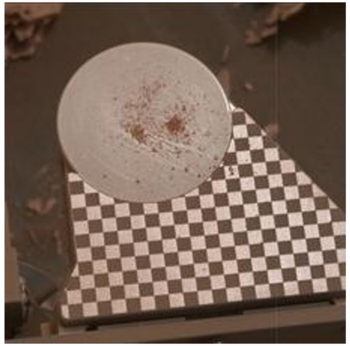

Observation Tray

Rear Rover Deck

Results for the 3 Mars Classifiers

Model	Model Size (MB)	Number images	Quantization Discrepancy*	Incorrect bits in-flight	Inference Time Per Image (ms)	FPS
HiRISE	121	1,793	2 (0.11%)	0	16.14	61.95
MSL 1	121	600	3 (0.5%)	0	16.09	61.11
MSL 2	121	1305	3 (0.2%)	0	16.077	62.2

*Compared to laptop run

HiRISE on PC - 56.9 ms/17.5 FPS

UAVSAR: Deep Learning for Flood Mapping

What: Pixel-wise flood classification of Uninhabited Aerial Vehicle Synthetic Aperture Radar (UAVSAR) imagery

Applications: flood mapping, alert generation

Details:
Data: UAV polarimetric SAR data from Hurricane Harvey
real-time = 5.13 patches/second
Surface water extent model
Model: UNET-6 fully convolutional model [UAVSAR, UNET]

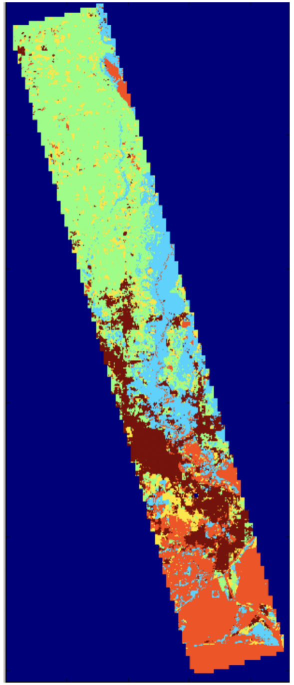

Hurricane Harvey Classified Image

UAVSAR: Quantization Discrepancy and Timing

Model	Model Size (MB)	Number images	Quantization Discrepancy: # missed pixels (out of 5058*2323 pixels) Full Classification / Binary Classification	Incorrect bits in-flight	Inference Time Per Image (ms)	FPS*
UAVSAR	4.3	3,633	184,820 (1.56%) / 82,519 (0.70%)	0	6	166.6

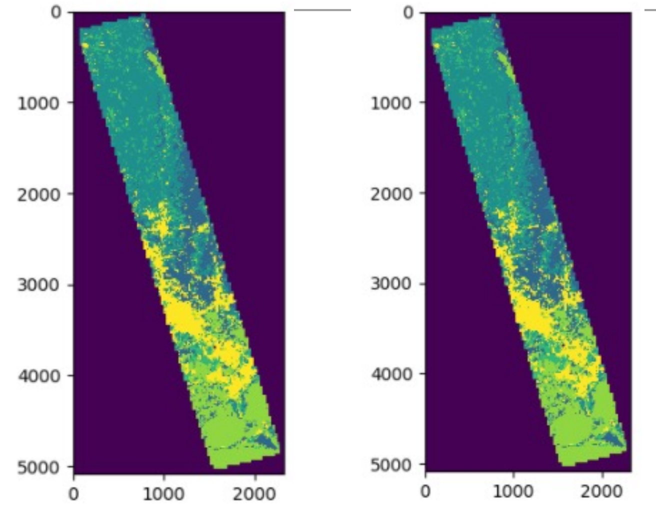

CPU-Keras output v Myriad X output

*200 ms per image
1.3 images/second to meet real time

25 image patches per second on Linux

Ship Segmentation: Detection in Satellite Imagery

- What: Fast ship detection in satellite imagery
- Applications: situational awareness
- Details:
  - Data: Kaggle Airbus dataset
  - Model: UNET + MobileNetV2 backbone, 8 SHAVES [UNET, Mobile]
- Myriad timing:
  - Model Size 14.8 MB
  - 98 ms / image, 10.2 FPS
  - images of size 320 x 320 pixels

Airbus Ship Detection Dataset
https://www.kaggle.com/c/airbus-ship-detection/data

Myriad X DDR test

Write pattern to DDR, read back and calculate memory bit errors

Provides full coverage of 512MB DDR

No in-flight errors recorded thus far

- Test 1 - Write address value to address
  - 0x80000000 to 0x80000000
- Test 2 - Write XOR of address value to address
  - 0x7FFFFFFF to 0x80000000
- Test 3 - Marching test of 63 rounds
  - Check flipping of bits from 0 – 1 and 1 – 0. Same pattern written to each address each round

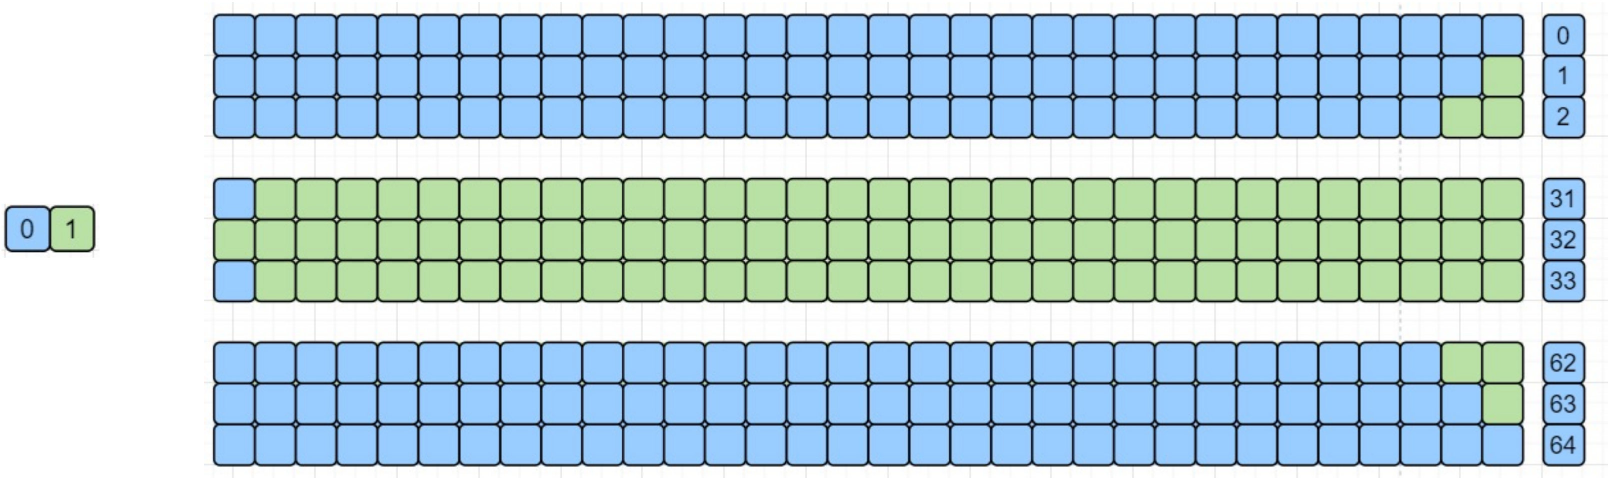

Prior Flights: Machine Learning Inference Onboard

- Autonomous Sciencecraft Experiment (ASE on EO-1)
  - Flew SVM 2005-2017
  - Flew Random Forest, Salience, Bayesian Thresholding 2017
  - Onboard image analysis for intelligent downlink and autonomous retargeting

- Intelligent Payload Experiment (IPEX)
  - CubeSat mission
  - Flew Random Forest, Salience 2013-2014

- PhiSat-1 (2020 - 2021)
  - CubeSat mission from the European Space Agency [ESA]
  - Flew Myriad 2 (previous generation VPU), automatic discard of cloudy imagery

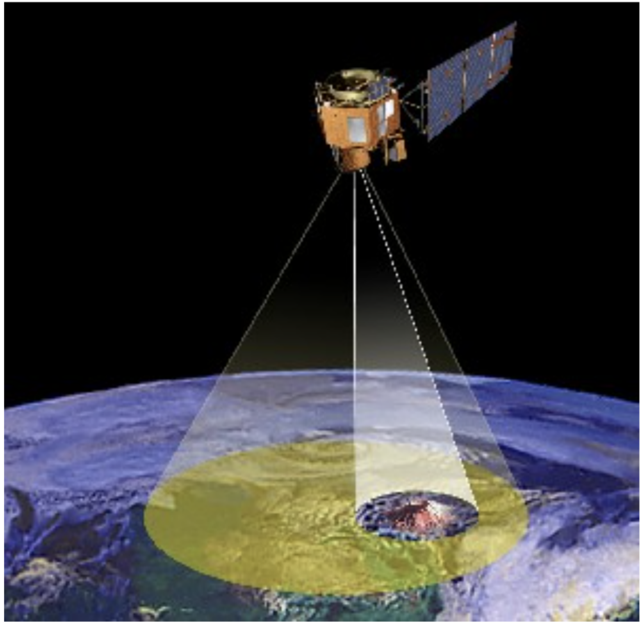

Earth Observing-1 (EO-1) Mission
(ai.jpl.nasa.gov)

Future Work / Conclusions

- Continually benchmarking new deep learning models on the Myriad X processor on SBC-2 onboard the ISS
  - Upcoming models include:
    - Volcano eruption detection
    - Viable road classification for disaster relief
- We have demonstrated fast and accurate inference with the Myriad X
  - Step toward new era of more complex and powerful processing with edge computing

References

Intel Movidus X Processor

https://www.intel.com/content/www/us/en/products/details/processors/movidius-vpu.html

HPE's Spaceborne Computer-2

https://www.hpe.com/us/en/compute/hpc/supercomputing/spaceborne.html

Machine Learning References

[DL] Krizhevsky, Sutskever, Hinton. 2012. ImageNet classification with deep convolutional neural networks. Advances in Neural Information Processing Systems 25, 1097-1105.

[PDS] https://pds-imaging.jpl.nasa.gov/search

[Mars1] Wagstaff, Lu, Dunkel, Grimes, Zhao, Cai, Cole, Doran, Francis, Lee, Mandrake, 2021. Mars Image Content Classification: Three Years of NASA Deployment and Recent Advances. Innovative Applications of Artificial Intelligence, 33.

[Mars2] Wagstaff, Lu, Stanboli, Grimes, Gowda, Padams. 2018. Deep Mars: CNN Classification of Mars Imagery for the PDS Imaging Atlas. The Thirteenth AAAI Conference on Innovative Applications of Artificial Intelligence. 7867-7872.

[UAVSAR] Denbina, Towfic, Thill, Bue, Kasraee, Peacock, Lou, Flood Mapping Using UAVSAR and Convolutional Neural Networks. IEEE Xplore 2020

[UNET] Ronneberger, Fischer, Brox, 2015. U-net: Convolutional Networks for Biomedical Image Segmentation. International Conference on Medical Image Computing and Computer-assisted Intervention, 234-241.

[Mobile] Sandler, Howard, Zhu, Zhmoginov, Chen, 2018. MobileNetV2: Inverted Residuals and Linear Bottlenecks. CVPR 2018.

[ESA] https://www.esa.int/Applications/Observing_the_Earth/Ph-sat

Jet Propulsion Laboratory
California Institute of Technology

jpl.nasa.gov