

Digital Coast for Climate Resilience from Drones: Automation and Scaling from Satellite and Drone Collections

David Hughes, Ph.D.

Oak Ridge National Laboratory

Matthew McCarthy, Ph.D.

Oak Ridge National Laboratory

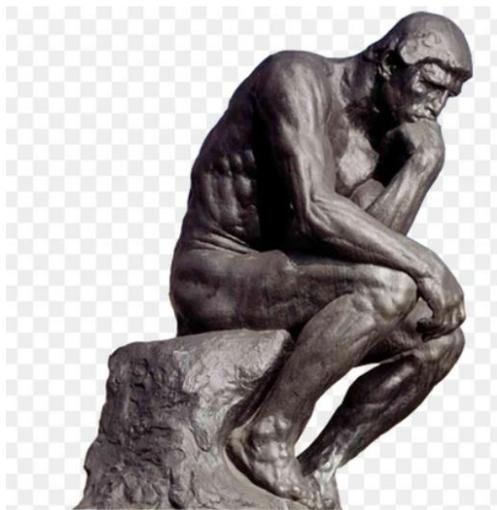
International Society for Digital Earth (ISDE)

26 September 2023

How to map coastal regions effectively?

Drones!

Satellites!



ORNL is DOE's largest science and energy laboratory

\$1.9B budget

4,500+ employees

3,200+ research guests annually

\$750M modernization investment

Nation's largest materials R&D portfolio

Oak Ridge Leadership Computing and Remote Sensing Facilities

Nation's most powerful open science computing facility

World's most intense pulsed neutron source

World-class research reactor

Nation's most diverse energy R&D portfolio

Managing billion-dollar U.S. ITER project



Toward automated, accurate remote sensing workflow

Most analysis pods consist of research scientists, image scientists, image analysts, geospatial scientists, and geospatial analysts.

They are focused on:

1. Operating in disconnected environments
2. Landcover detection, accurate terrain extraction using mostly imagery and maybe some video
3. Automation over a large scale of imagery/video acquisitions

They have a large amount of imagery/video collected and can't/don't lay eyes on all data coming in, sampling.

Many research groups would profit from edge compute over constrained comms, automation with regards to image quality [what data is going to yield results] and object detection/wide-area-search [what data seems to have objects of interest in them].

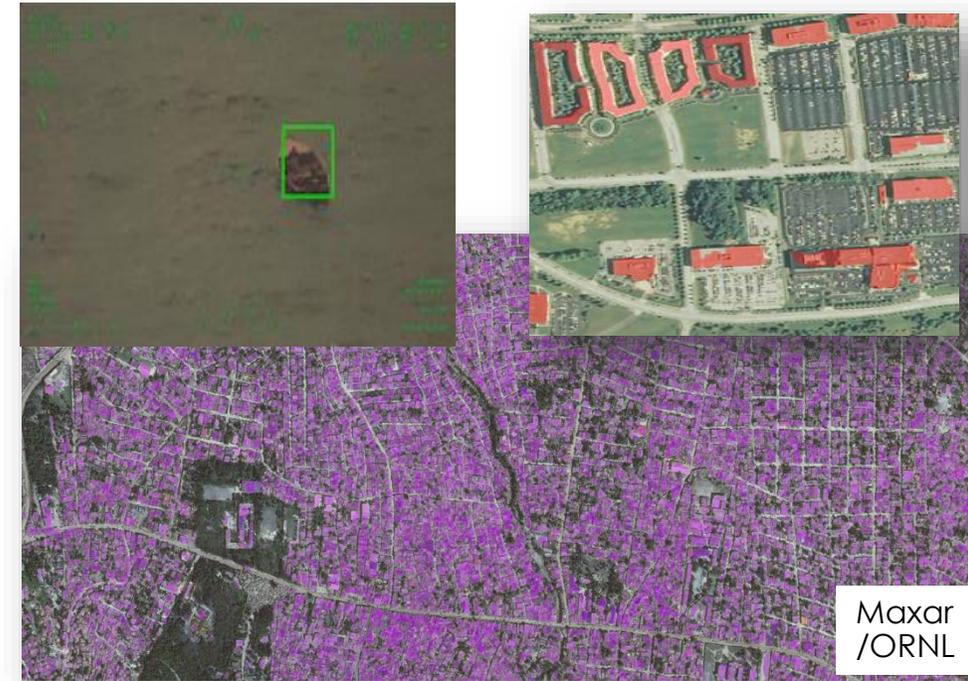
We are going to show some examples of:

- Satellite-derived efforts for the coastal environment, namely bathymetry from spectra and IOPs
- Drone engineering for disaster assessment in constrained communications environments for an aspect of the infrastructure
- Automated pre-processing developed by ORNL
- Mensuration, and wide-area-search to help with automation and timely reporting.
- The ORNL image pipeline



Automated Feature Extraction (AFE)

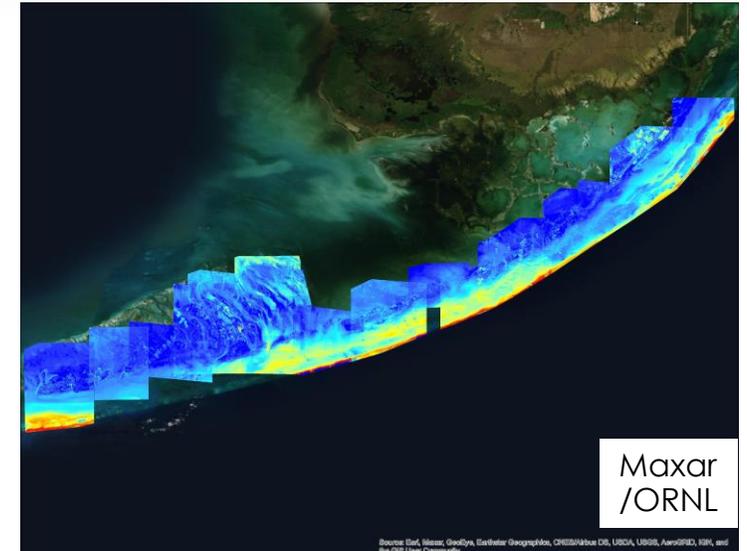
- GeoAI driven mapping of built and natural environments (near-shore characterization, machine learning relevant to maritime, buildings, roads, and other features) with astonishing clarity and speed
- Emerging applications in surveillance, change monitoring, attribution, and collateral damage estimation
- Leveraging ORNL expertise in High Performance Computing, Machine Learning, Remote Sensing, Data science, Computer Vision, and Geoinformatics engineering



At global scale: Large scale feature mapping



The "only true exascale machine" on the TOP 500 list.



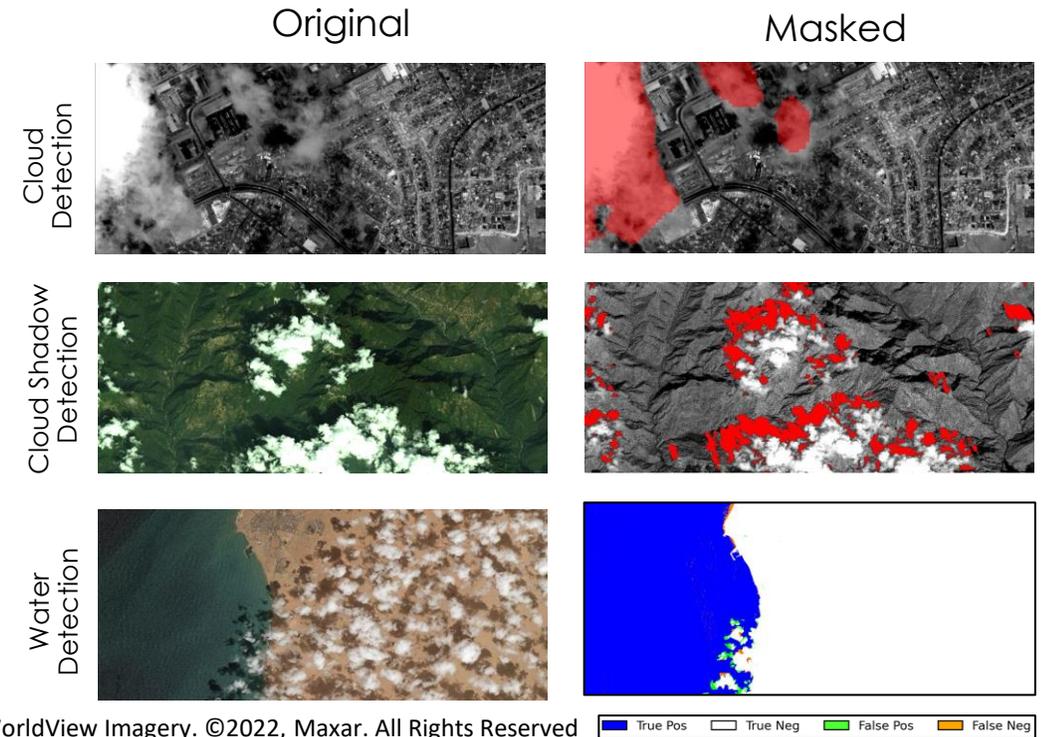
HPC, large scale, high-resolution bathymetry extraction

What we have found important about edge computing

Most sponsors have determined that they need actionable intelligence within 24 to 48-hours.

To meet that deadline:

1. Scalable data collection.
2. Automation to reduce the imagery/video acquisitions to detected features of interest.
3. Ability to function in hostile post-event environments including destroyed local communications infrastructure.

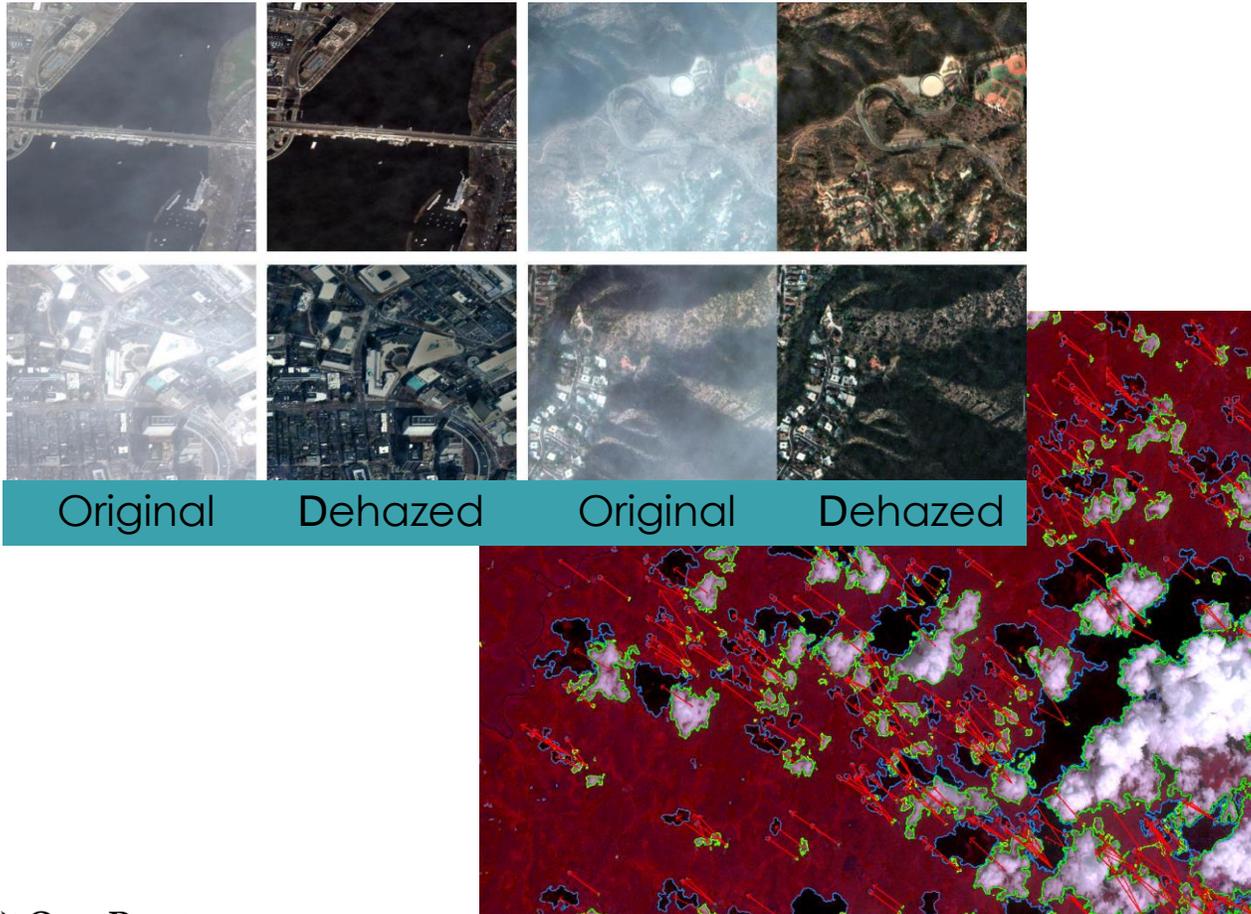


ORNL has years of experience:

- Delivering automated ML/AI pre-processing solutions.
- Deploying analyst-assist tools to sponsor locations.
- Achieving timely reporting of wide-area measurement and searches through edge-computing and automation.
- Developing an image pipeline producing analyst-ready products.

Reducing the Processing Load

Quality assessment reduces the need to process vast swaths of useless imagery



Quality aspects currently considered

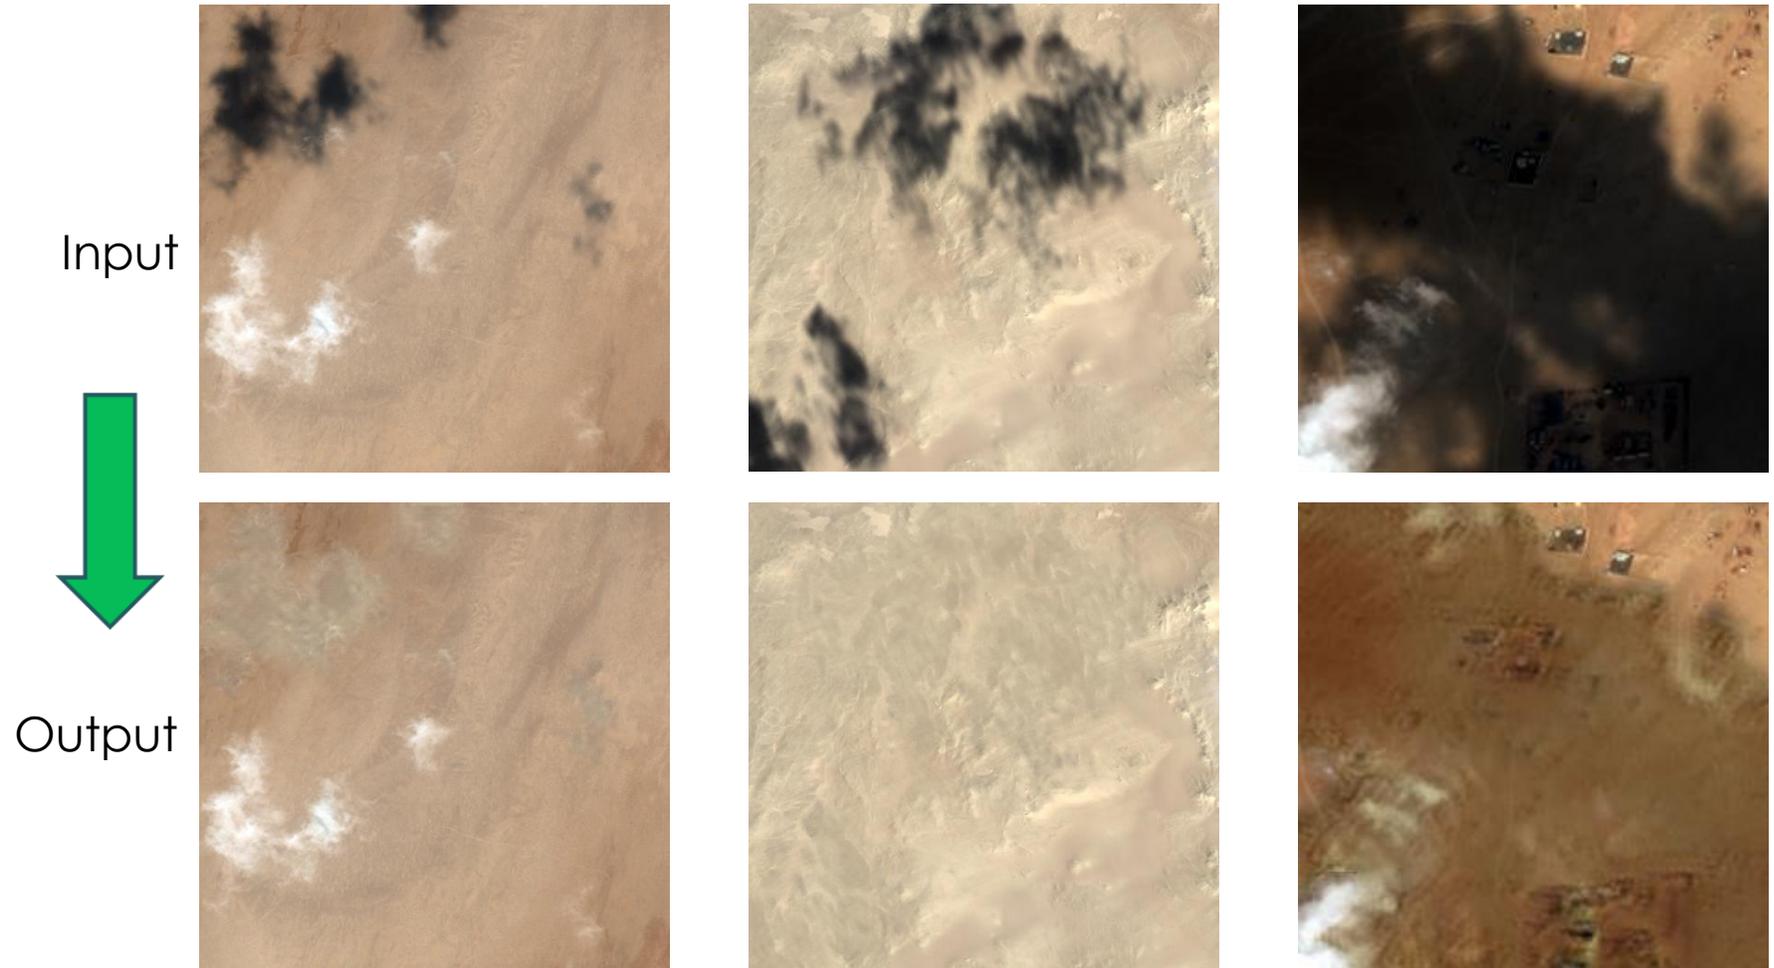
- Cloud detection
- Haze detection and mitigation
- Water detection (MSI) / glint correction
- Shadow detection and mitigation
- Image anomaly mapping
- HPC Atmospheric compensation

The effects of poor image quality

- Erroneous statistics for detection algorithms
- Poor confusion matrix representation
- Invalid timings to process

“The Dark Arts”: Shadow detection and mitigation

- Objective: preserve spectra
- Machine learning (cycleGAN-based algorithm)
- Increase interpretability of imagery



WorldView Imagery. ©2022, Maxar. All Rights Reserved
Byung Park, ORNL

Satellite-based Debris Measurement and Wide Area Outage



Figure 17a. Norman Regional Moore Medical Center. (Left) Greyscale image. (Right) Derived digital surface model. Both pre-tornado.

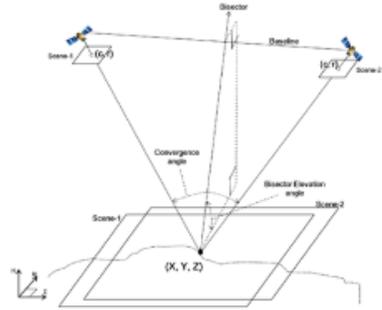
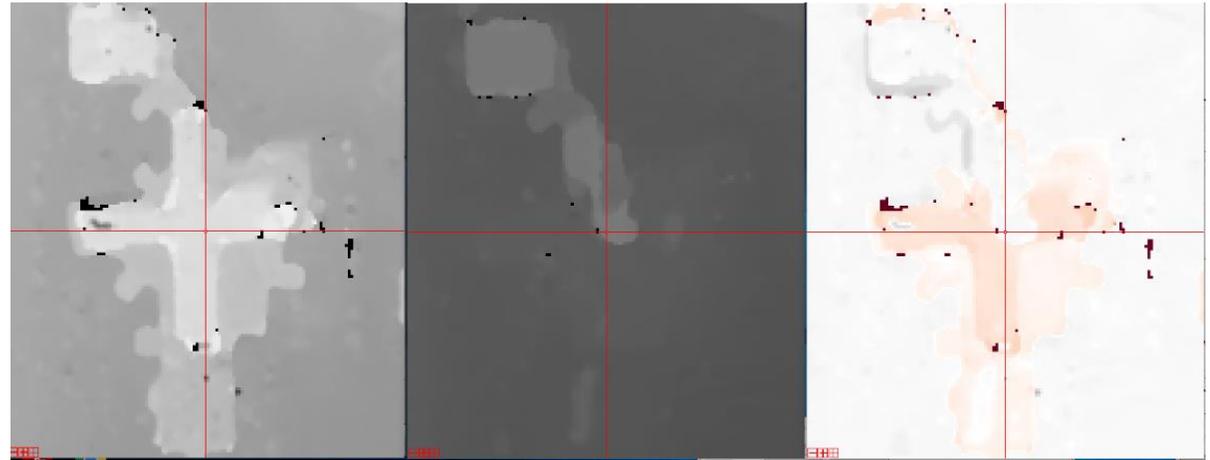


Figure 2. Illustration of optical pair detection by measurement of the hybrid "convergence angle" sensor. The viewing angle of the sensor with respect to the ground is 30 degrees.



In right image, darker reds denote reduced heights and darker greys denote increased heights. White is no change.

Hospital DSM pre-tornado (l),
Hospital DSM post-tornado (c),
Hospital debris estimate (r).

U.S. DEPARTMENT OF ENERGY
Office of Cybersecurity, Energy Security, and Emergency Response

Wide Area Outages

Wide Area Outage Detection is used to identify areas that may be experiencing power outage using the Day/Night Band of the Visible Infrared Imaging Radiometer Suite (VIIRS) onboard the Suomi National Polar-Orbiting Partnership (Suomi NPP) spacecraft.

[Login](#)

NOTICE TO USERS: This is a Federal computer system and is the property of the United States Government. It is for authorized use only. Users (authorized or unauthorized) have no explicit or implicit expectation of privacy. Any or all use of this system in violation of this system may be intercepted, monitored, recorded, copied, stored, inspected, and disclosed to authorized site, Department of Energy, and law enforcement personnel, as well as authorized officials of other agencies, both domestic and foreign. By using this system, the user consents to such interception, monitoring, recording, copying, auditing, inspection, and disclosure at the discretion of authorized site or Department of Energy personnel. Unauthorized or improper use of this system may result in administrative disciplinary action and civil and criminal penalties. By continuing to use this system you indicate your awareness of and consent to the terms and conditions of the Terms of Use. **LOG OFF IMMEDIATELY** if you do not agree to these terms and conditions.

Actionable Analysis from Satellites and UAS

Project Overview:

- DOE CESER provided funding for ORNL to generate scalable damage assessment and outage products from satellite and UAS. Products are required to be deployed through EAGLE-I™ to support response and recovery of energy infrastructures.

Project Technology:

- Satellite-based broad area outage and restoration mapping from satellite (VIIRS) nighttime lights
- Debris estimation at approx. 1-meter spatial resolution from 2-view satellite acquisitions.
- **Automated utility pole damage assessment using UAS and machine learning**



WorldView Imagery.
©2022, Maxar. All Rights Reserved

Autonomous UAS Lessons Learned

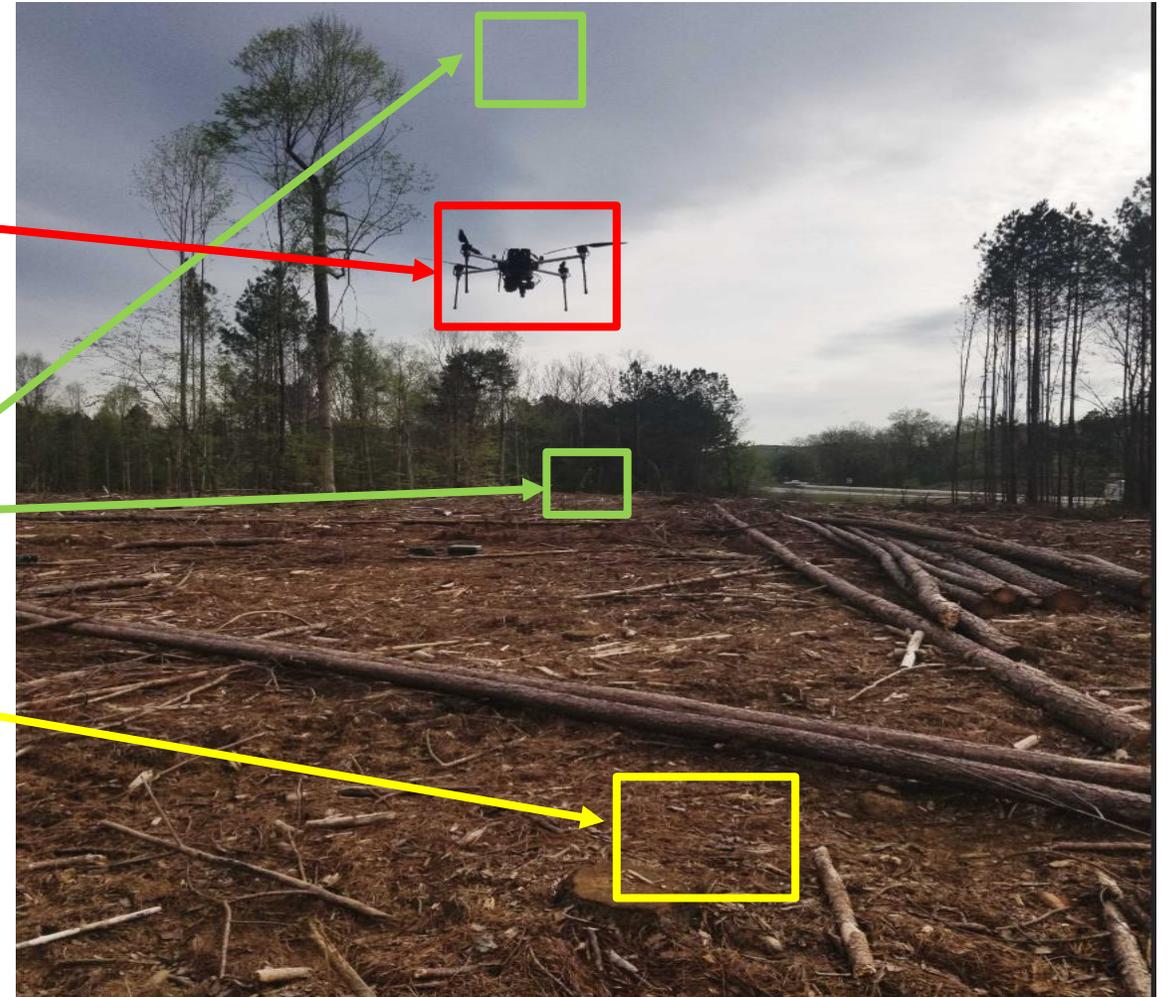
1) UAS platform

- 1) Processor
- 2) Compute
- 3) Machine Learning

2) Global command and control

3) Geopositioning

Lexie Yang, Jordan Bowman
Jerry Kirk, Jairus Hines, Dakota Haldeman
Orrin Thomas
David Hughes, Darrell Roddy

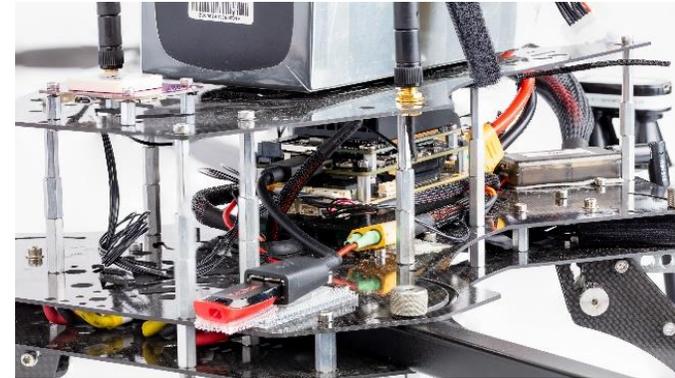


Custom airframe: Cornelia

The prototype aircraft (figure below) was named “Cornelia,” after the first wife of the Roman emperor Julius Caesar. Modifications include the addition of the PixC4 compute module, camera, and the Multimodal Autonomous Vehicle Network (MAVNet) cellular command and control software.

Specifications:

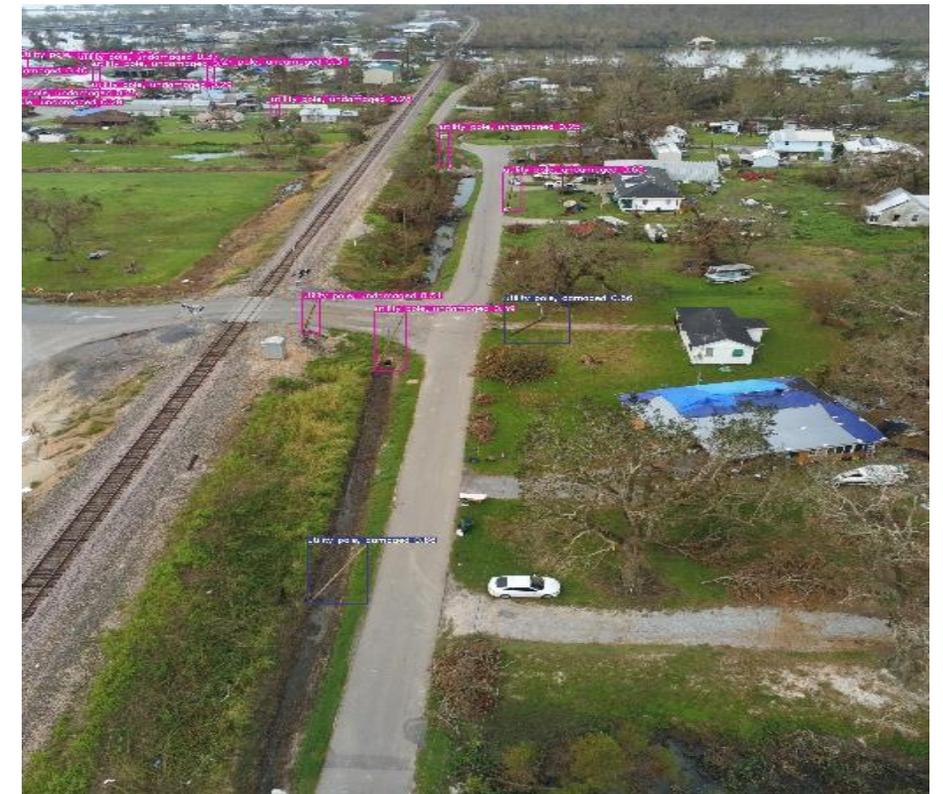
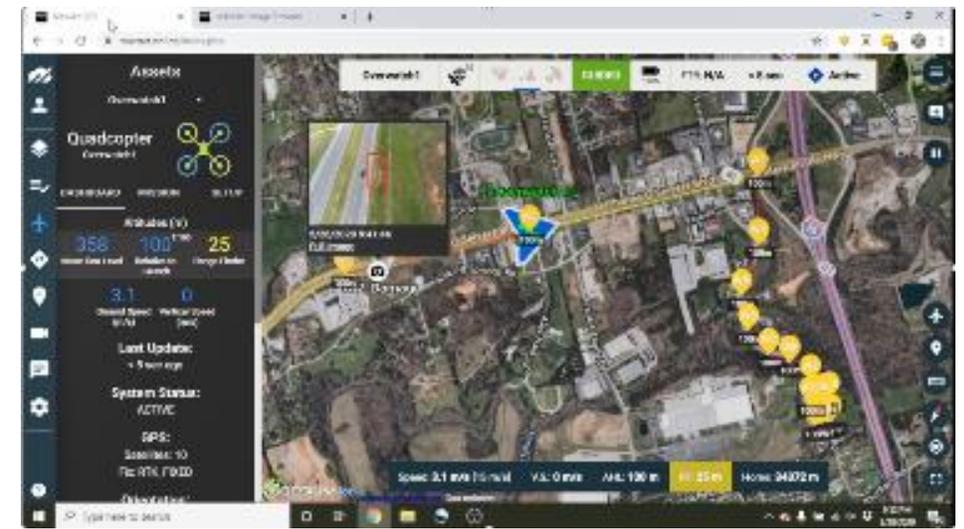
- Fully autonomous operations
- No need for local communications
- Weight: 3,895 g (with battery and payload)
- Wingspan: ~44 in.
- Battery: 6S (24 V); 10,000 mAh lithium polymer
- Flight time: 25 min
- Range: 4–5 mi



Modified aircraft “Cornelia” prior to test flight. (Source: Carlos Jones, ORNL.)

Edge Computing Solution Framework

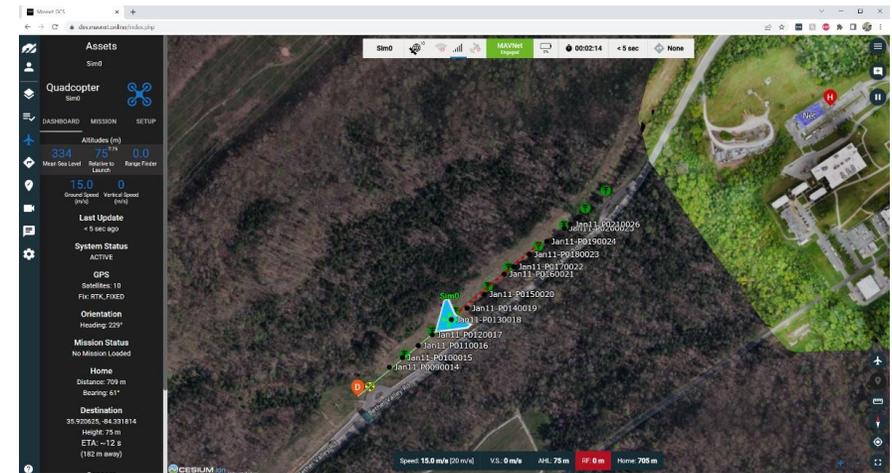
- Custom airframe capable of real-time operations.
 - Onboard imagery processing
 - Reports back over non-local constrained comms.
 - Fully autonomous
- Alternative near real-time operations using COTS
- Model Training, Validation, Testing, and deployment on edge devices
- Web based control and monitoring (MavNet)
- With good cellular connections, it's near real time mostly because we fly very slow and take pictures at distance intervals (20 meters). If we fly faster, there will be more lag. One object chip per 10 seconds at the low-ish quality we are using now.



Hurricane Ida: machine learning detection/classification on the edge

Communications and Global Command and Control

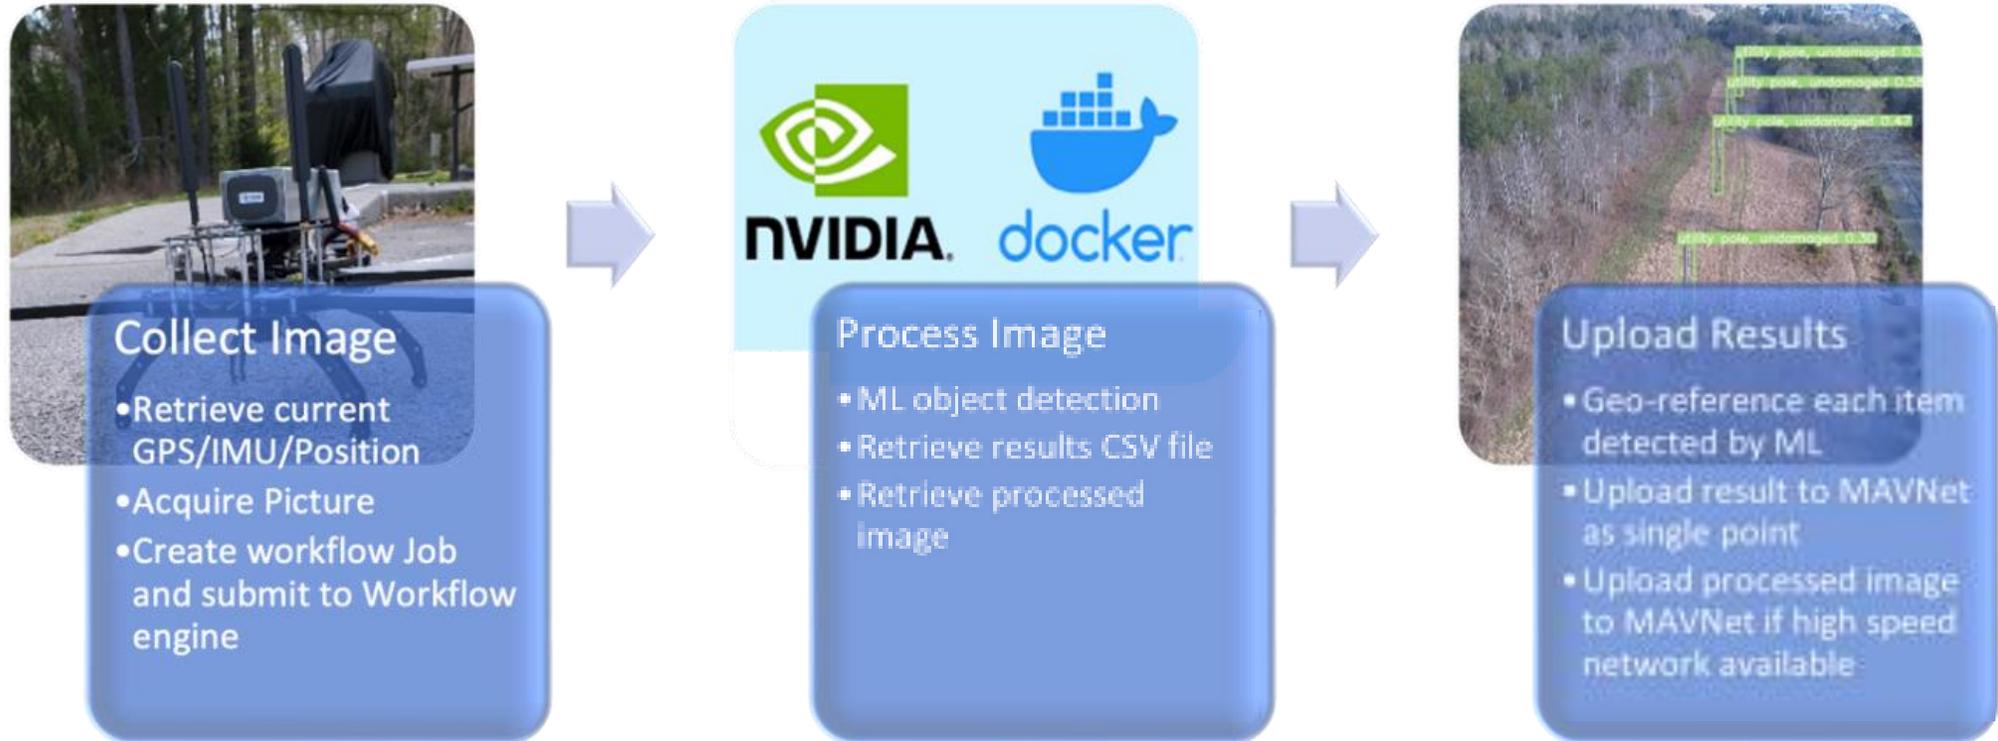
- Our satellite communication is based on [Iridium Short Burst Data](#) (SBD), small ~270byte packets that can be sent ~ 1-3 messages per minute when there are satellites above (about 80% of the time).
- MAVNet is a suite of hardware and software technologies developed at ORNL that maintains connectivity anywhere by allowing mesh networks, cellular/5G, and satellite communication modems to co-exist, automatically falling forwards or backwards among them. Beyond line-of-sight control.
- Key to this project, the MAVNet system allows for the UAV's operation through a web-based ground control system (GCS). The GCS is the operator's primary interface with the aircraft and permits the operator to see real-time data acquired from the UAV.



MAVNet GCS with sample detections. The green icons represent undamaged infrastructure.

Onboard imagery processing

Prototype utility pole detection, geolocation, and classification (damage or undamaged).



Model Selection and Training

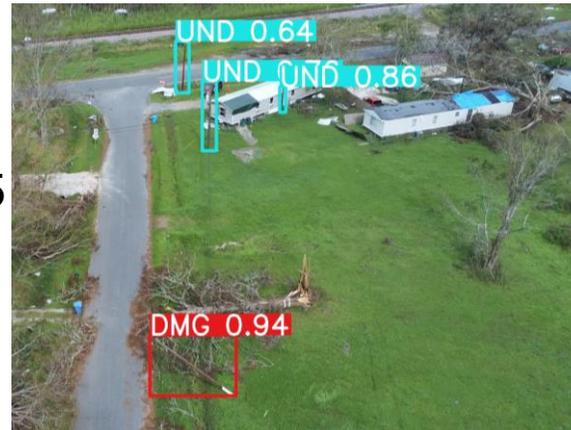
- YOLO object detection models were selected due to their good balance of speed and accuracy
 - YOLOv3: final version of YOLO made by original authors
 - YOLOv5: more recent iteration of YOLO
 - Model comes in multiple sizes with different tradeoffs between speed and accuracy
- Model training was performed in two stages
 - Pre-training: learn to detect utility poles
 - Fine-tuning: distinguish damaged and undamaged poles

Input



YOLOv5
Large

YOLOv5
Medium



YOLOv5
Small

Examples of the predictions from three YOLOv5 models. DMG in red boxes: damaged utility pole; UND in cyan boxes: Undamaged utility pole (Hurricane Ida).

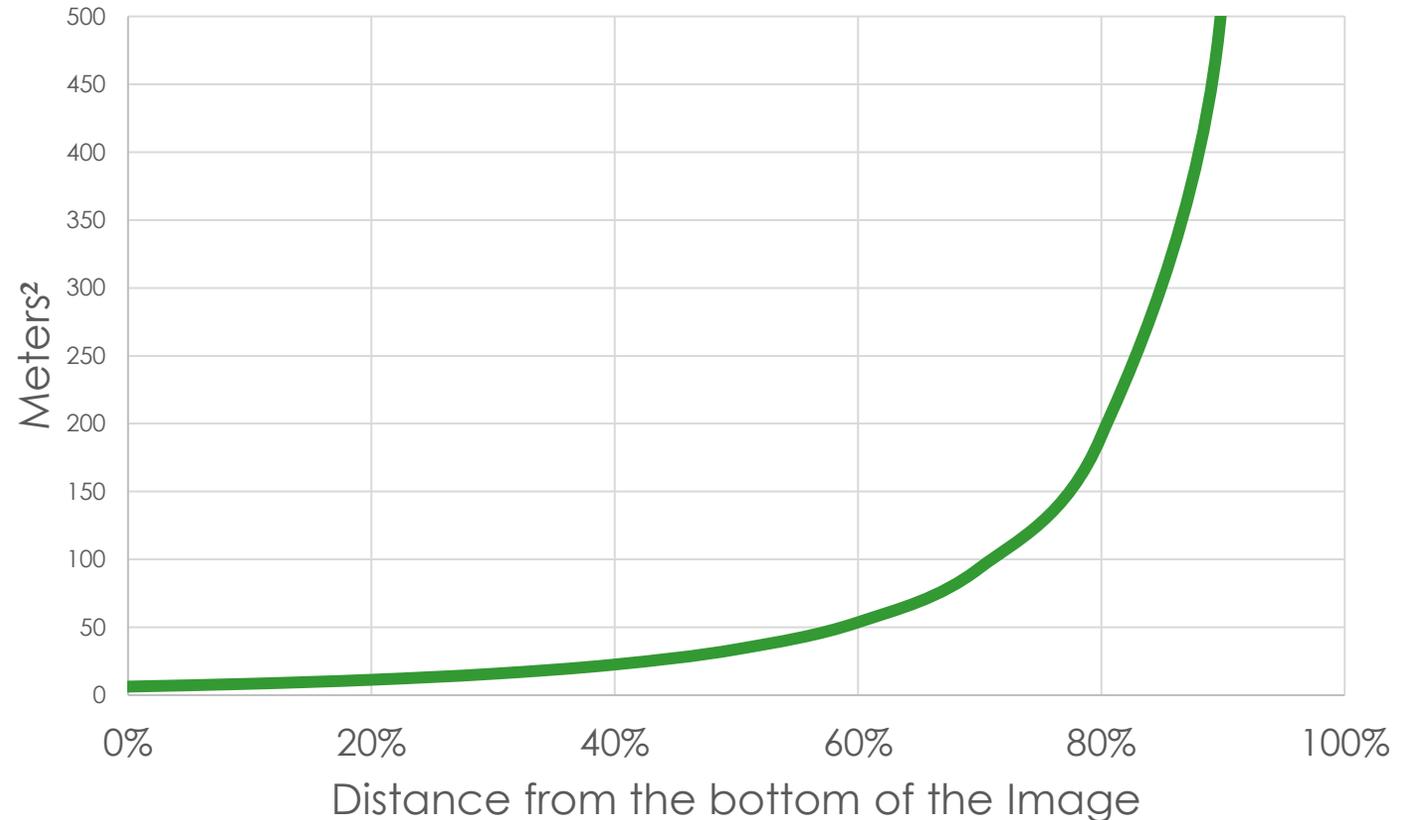
Geopositioning: Needs Improvement

Most edge detected features are not usefully geolocated. The uncertainty explodes in the top half of the frame because the images are oblique.

We do not plan to point the camera more-nadir because the oblique views are better for the object detection. Instead, future work, we plan to track the features through multiple frames.

Multiple looks at each feature would improve the ML classification and fix the geopositioning—provided there is at least one look in the bottom half of the frame.

Area of Error Ellipse



Edge computing timings and accuracy

Object detection model size and inferencing speed.

Framework	Number of parameters	Inferencing speed (seconds per image)
YOLOv5 Large	46,149,064	2.7
YOLOv5 Medium	20,879,400	1.35
YOLOv5 Small	7,027,720	0.48
YOLOv3	61,508,200	1.03

The system delivered an F1 score of 0.65 operating with a 2.7 s/frame processing speed with the YOLOv5 large model and an F1 score of 0.55 with a 0.48 s/frame with the YOLOv5 small model. These numbers will improve as we get more training data.

Geolocation uncertainty in the bottom half of the frame was ~8 m, mostly driven by error in camera pointing measurement.

With additional training data to improve performance and detect additional types of features, a fleet of similar drones could autonomously collect actionable post-disaster data.

A



B



C



D



Take-home for the researcher/engineer

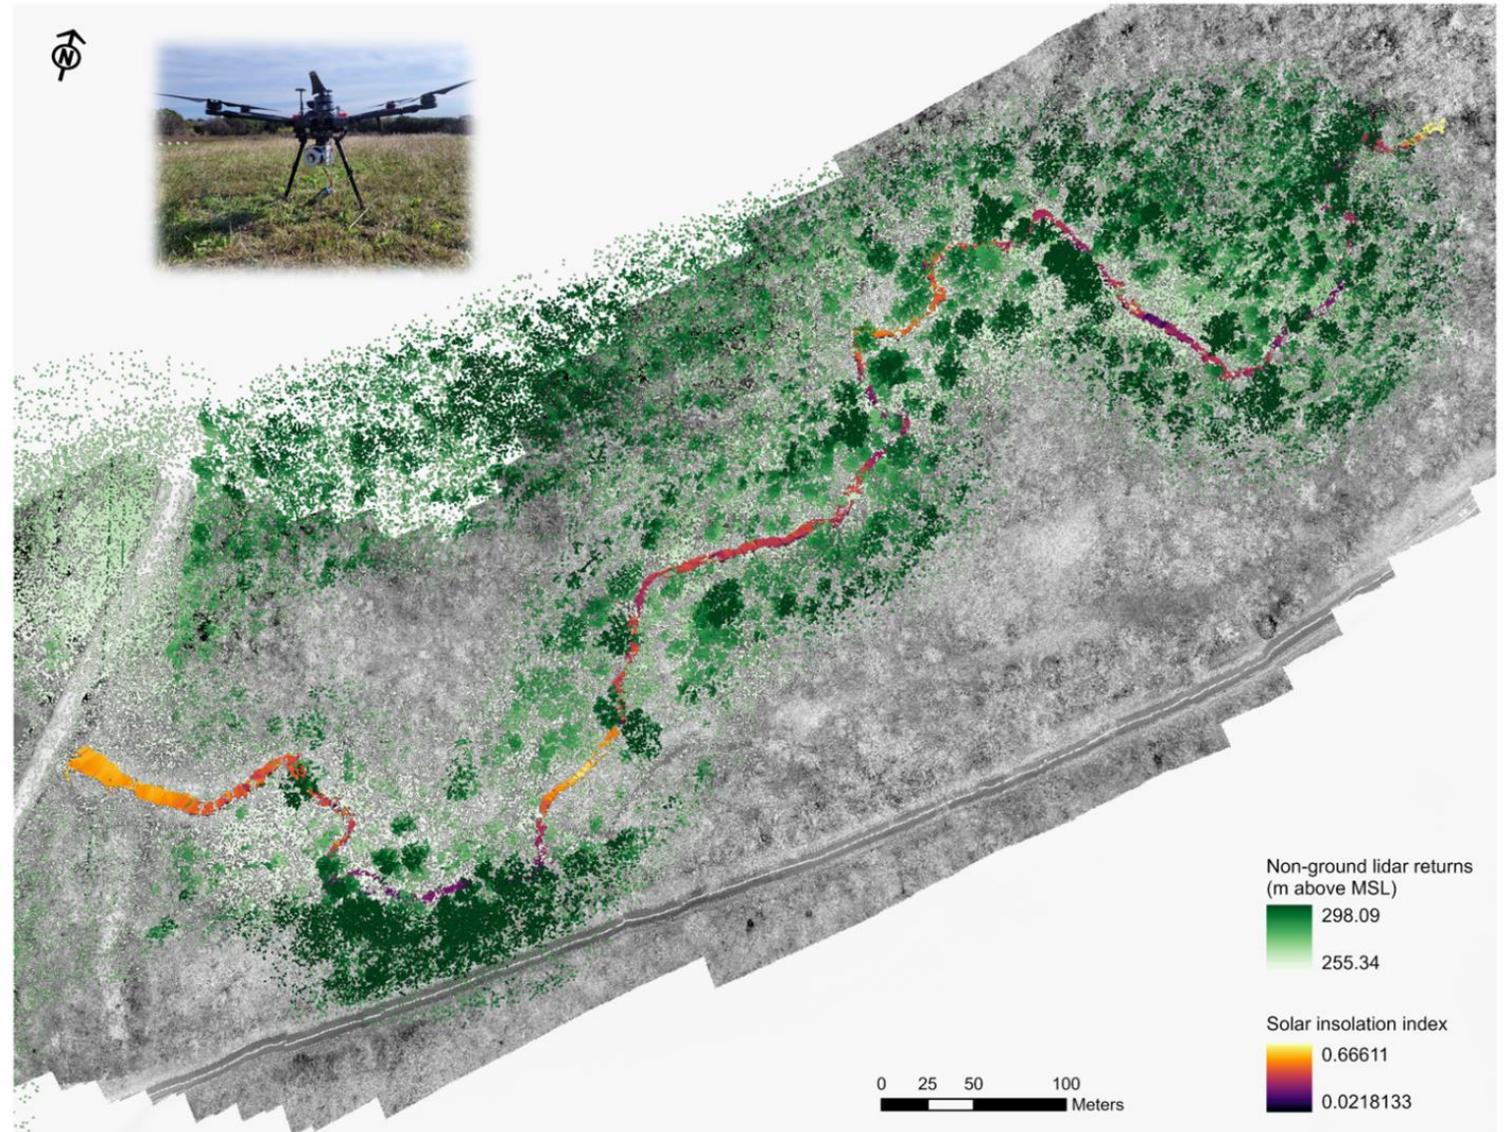
- Demonstrated a prototype UAS capable of edge, real-time object detection and autonomous operations without local communications networks (using satellite).
- Built cooperative relationships with national, state, and local entities to collect further training data and test real-time support (e.g. ESF-12 embedding and EAGLE-I TM integration).

What's next and potential collaboration?

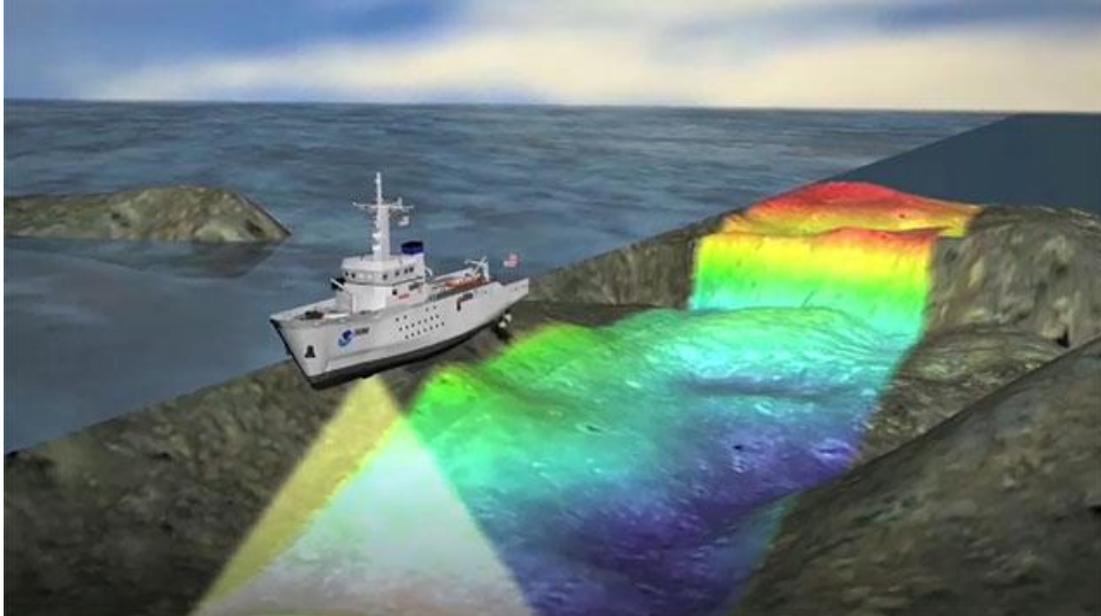
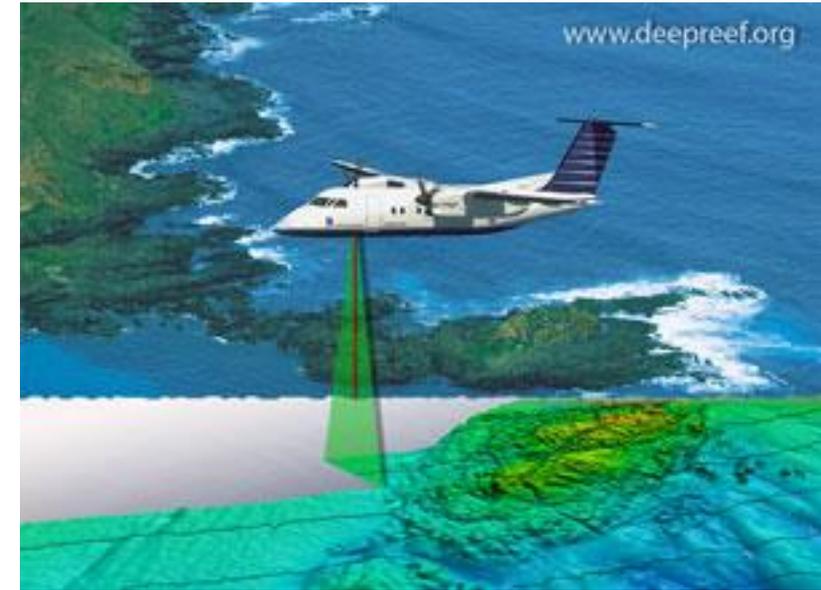
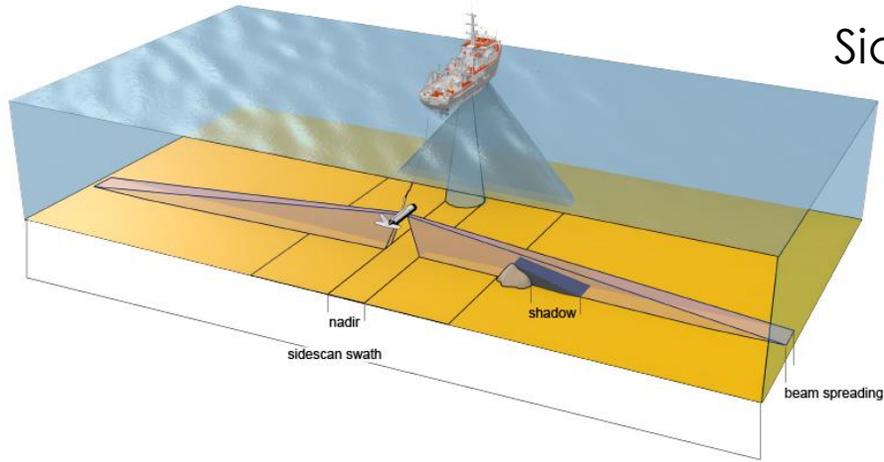
- Investigate object detection from different camera pointing angles to support more UAS platforms and use data collected for other purposes.
- Handle multiple assessments of same observable during flight intelligently.
- Research into adapting to other platforms, onboard processing hardware, other disaster observables (e.g. fire, flood,...)

ORNL Drone collections toward mercury methylation hotspots based on periphyton/vegetation communities/types

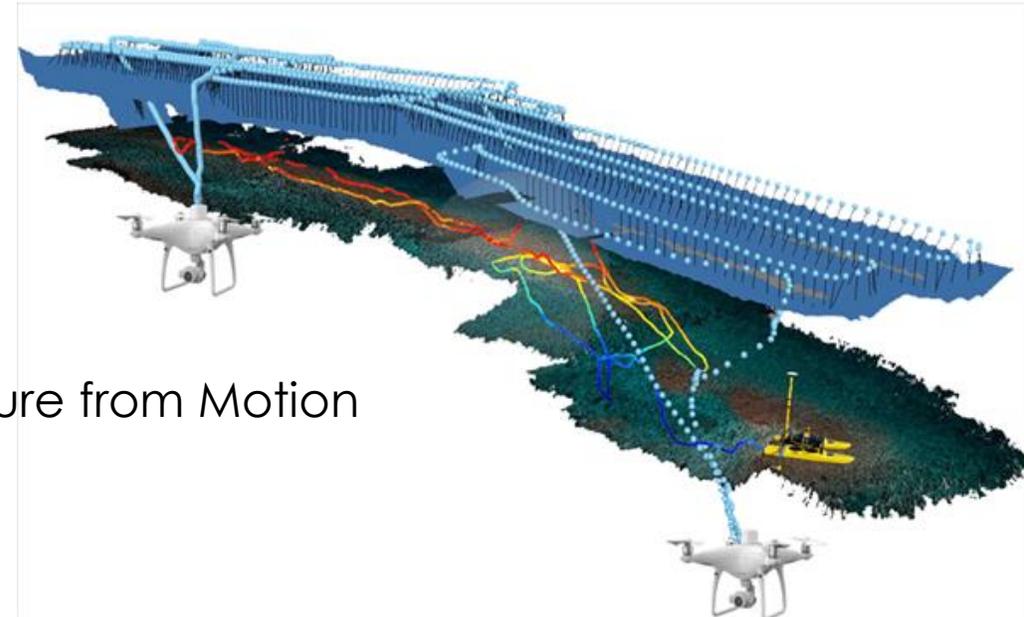
- Non-ground lidar point cloud returns and solar insolation index. Data collected on March 1, 2023 in Bear Creek. Non-ground returns shown in green, with darker green indicating taller trees. Solar insolation index shows estimate of solar exposure throughout stream channel, with higher values indicating higher exposure. Phoenix Lidar onboard the DJI M600 shown in inset.
- Chris DeRolph (derolphcr@ornl.gov)



Coastal Bathymetry Background



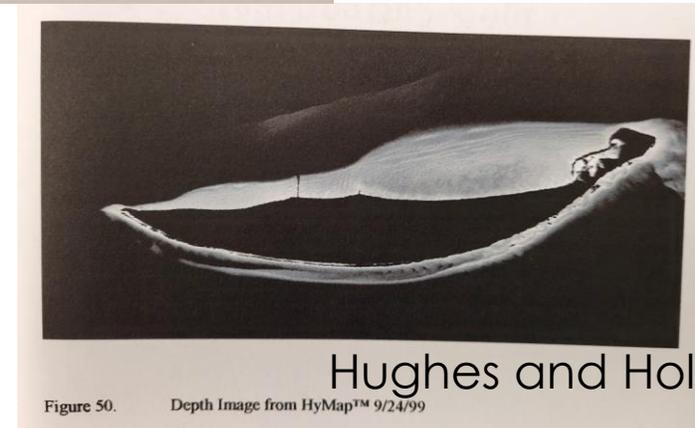
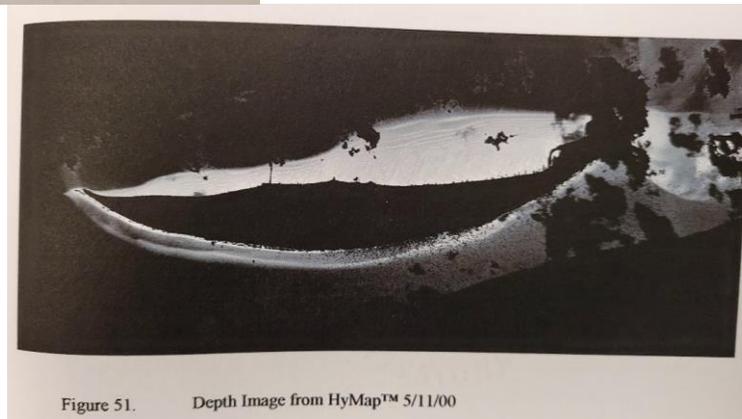
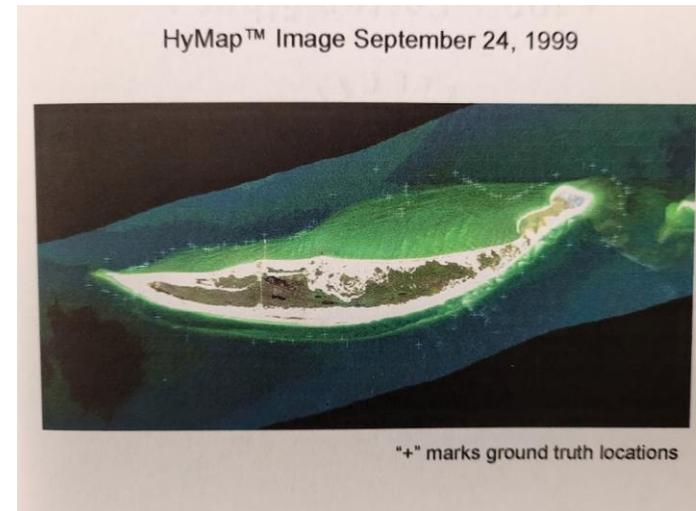
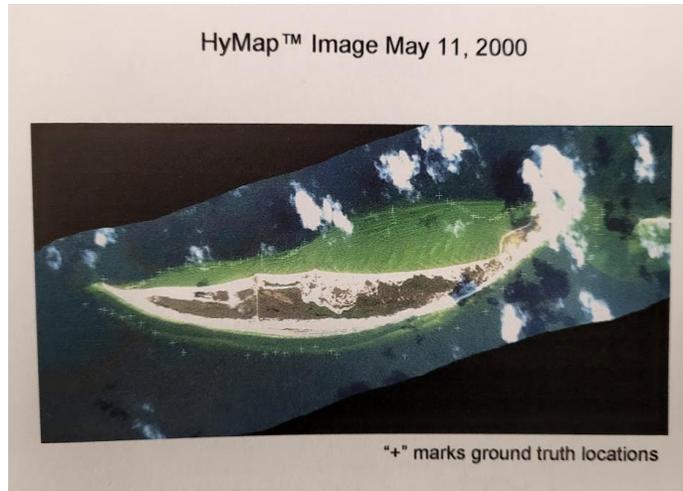
Multibeam Sonar



Structure from Motion

Coastal Bathymetry Background

More than 20 years ago, we were investigating airborne hyperspectral imagery and its use for bathymetry and coastal constituent retrieval.

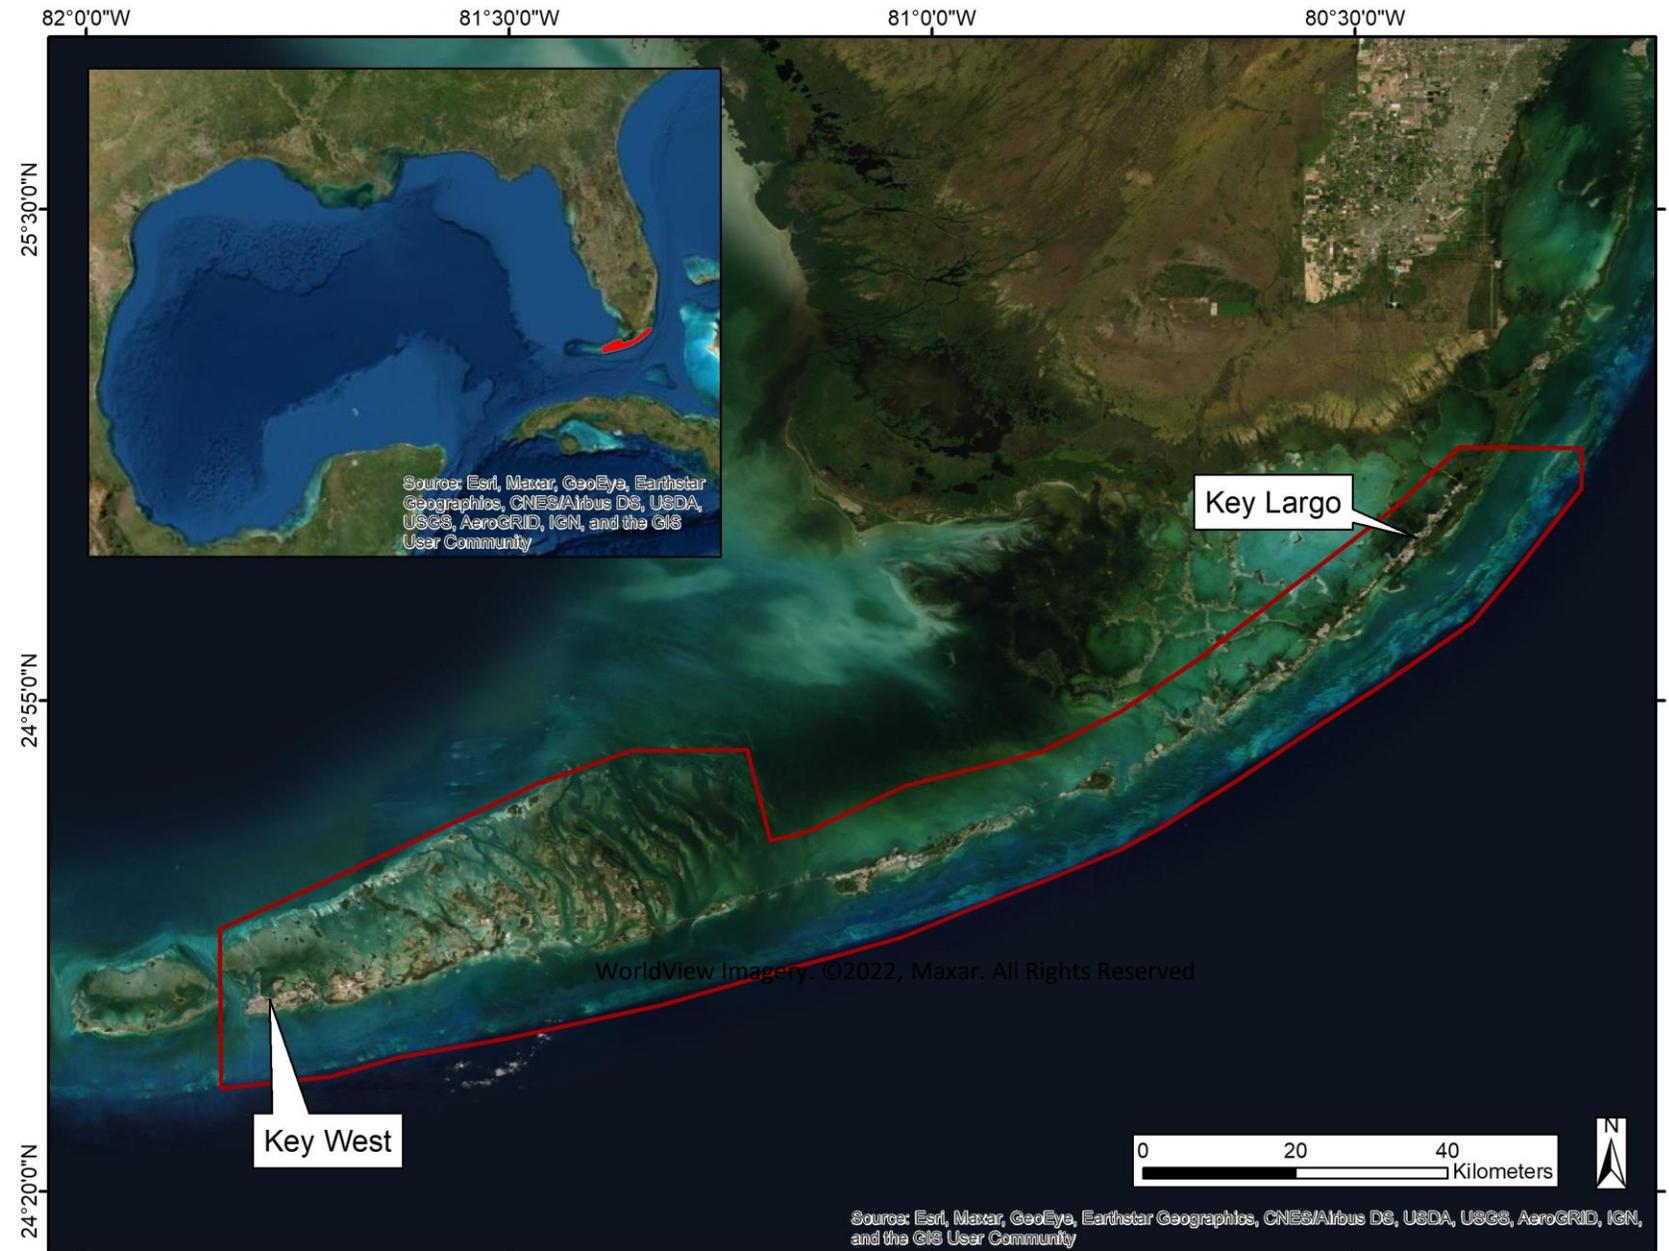


Hughes and Holyer, 2000

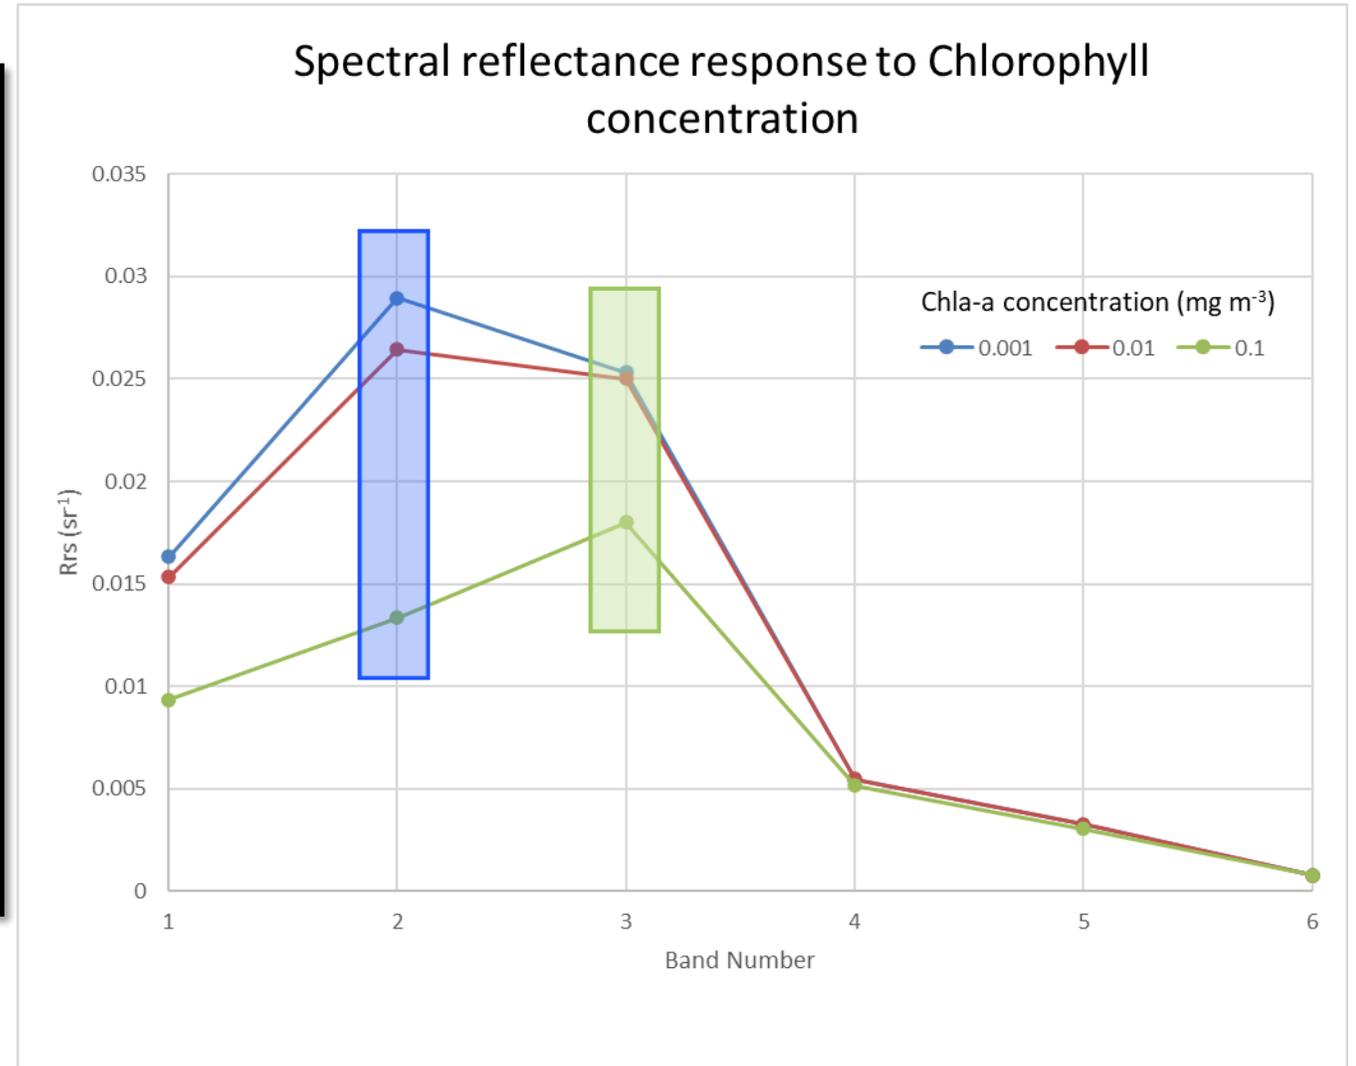
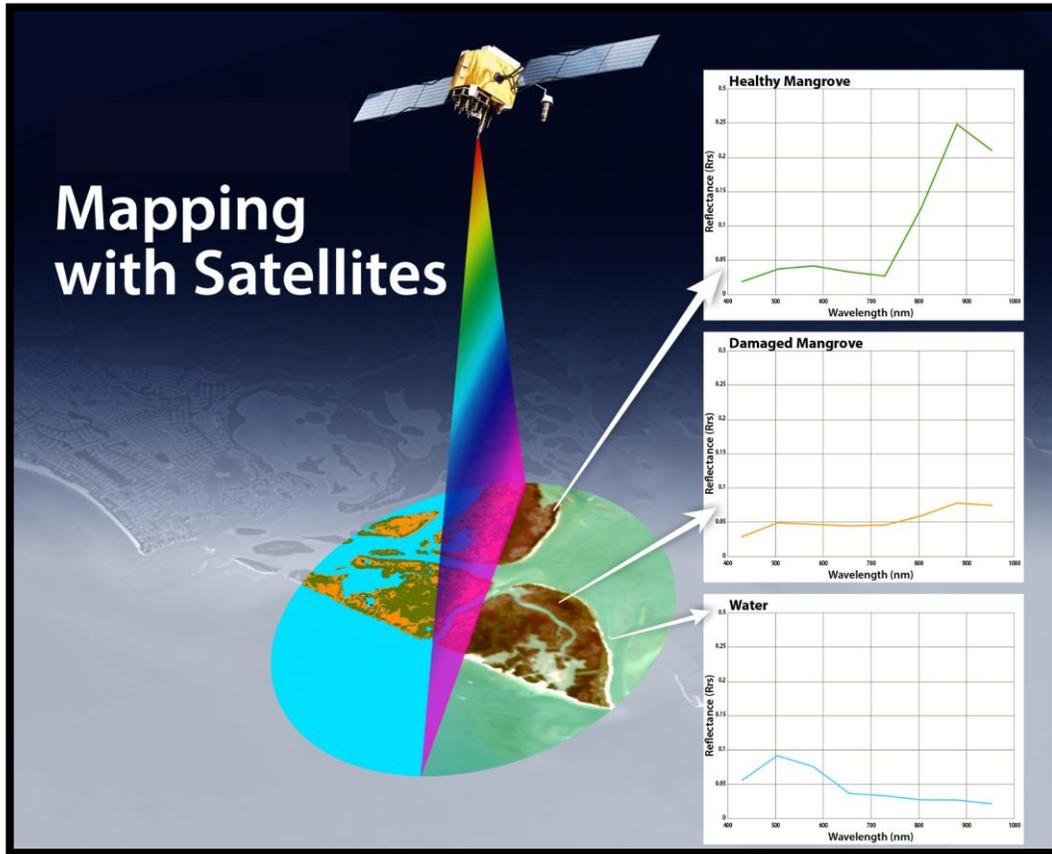
Study Area

Matt McCarthy (lead),
David Hughes

- Florida Keys
 - 3,700 km²
 - Shallow reef
 - Sandy bottom
 - Steep reef edge
 - Case I water



SDB: Light Attenuation



Satellite-Derived Bathymetry (SDB) Method:

- Band-ratio approach:

$$\text{Depth} = m1 * \frac{\ln(1000 * \text{Blue})}{\ln(1000 * \text{Green})} - m0$$

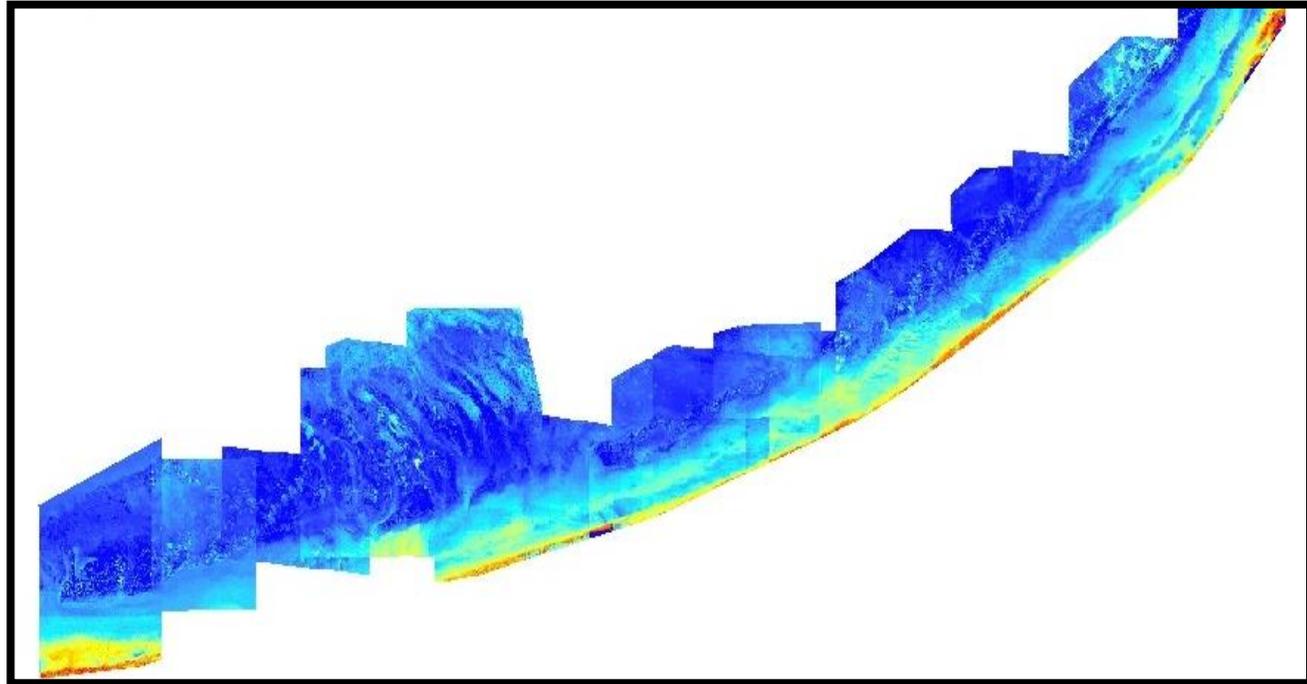
Stumpf et al. (2003)

Objective:

- Develop automated SDB pipeline that:
 - Estimates Chlorophyll-a concentration
 - Calculates coefficients
 - Maps bathymetry

Approach

1. Radiometrically calibrate
2. Correct for atmosphere
3. Correct for sunglint (as needed)
4. Correct for surface reflectance
5. Identify optically deep water (ODW)
 1. Estimate turbidity (chlorophyll concentration)
6. Derive tuning coefficients
7. Apply band-ratio algorithm
8. Postprocessing
 - Mosaic mapped tiles (ArcMap)

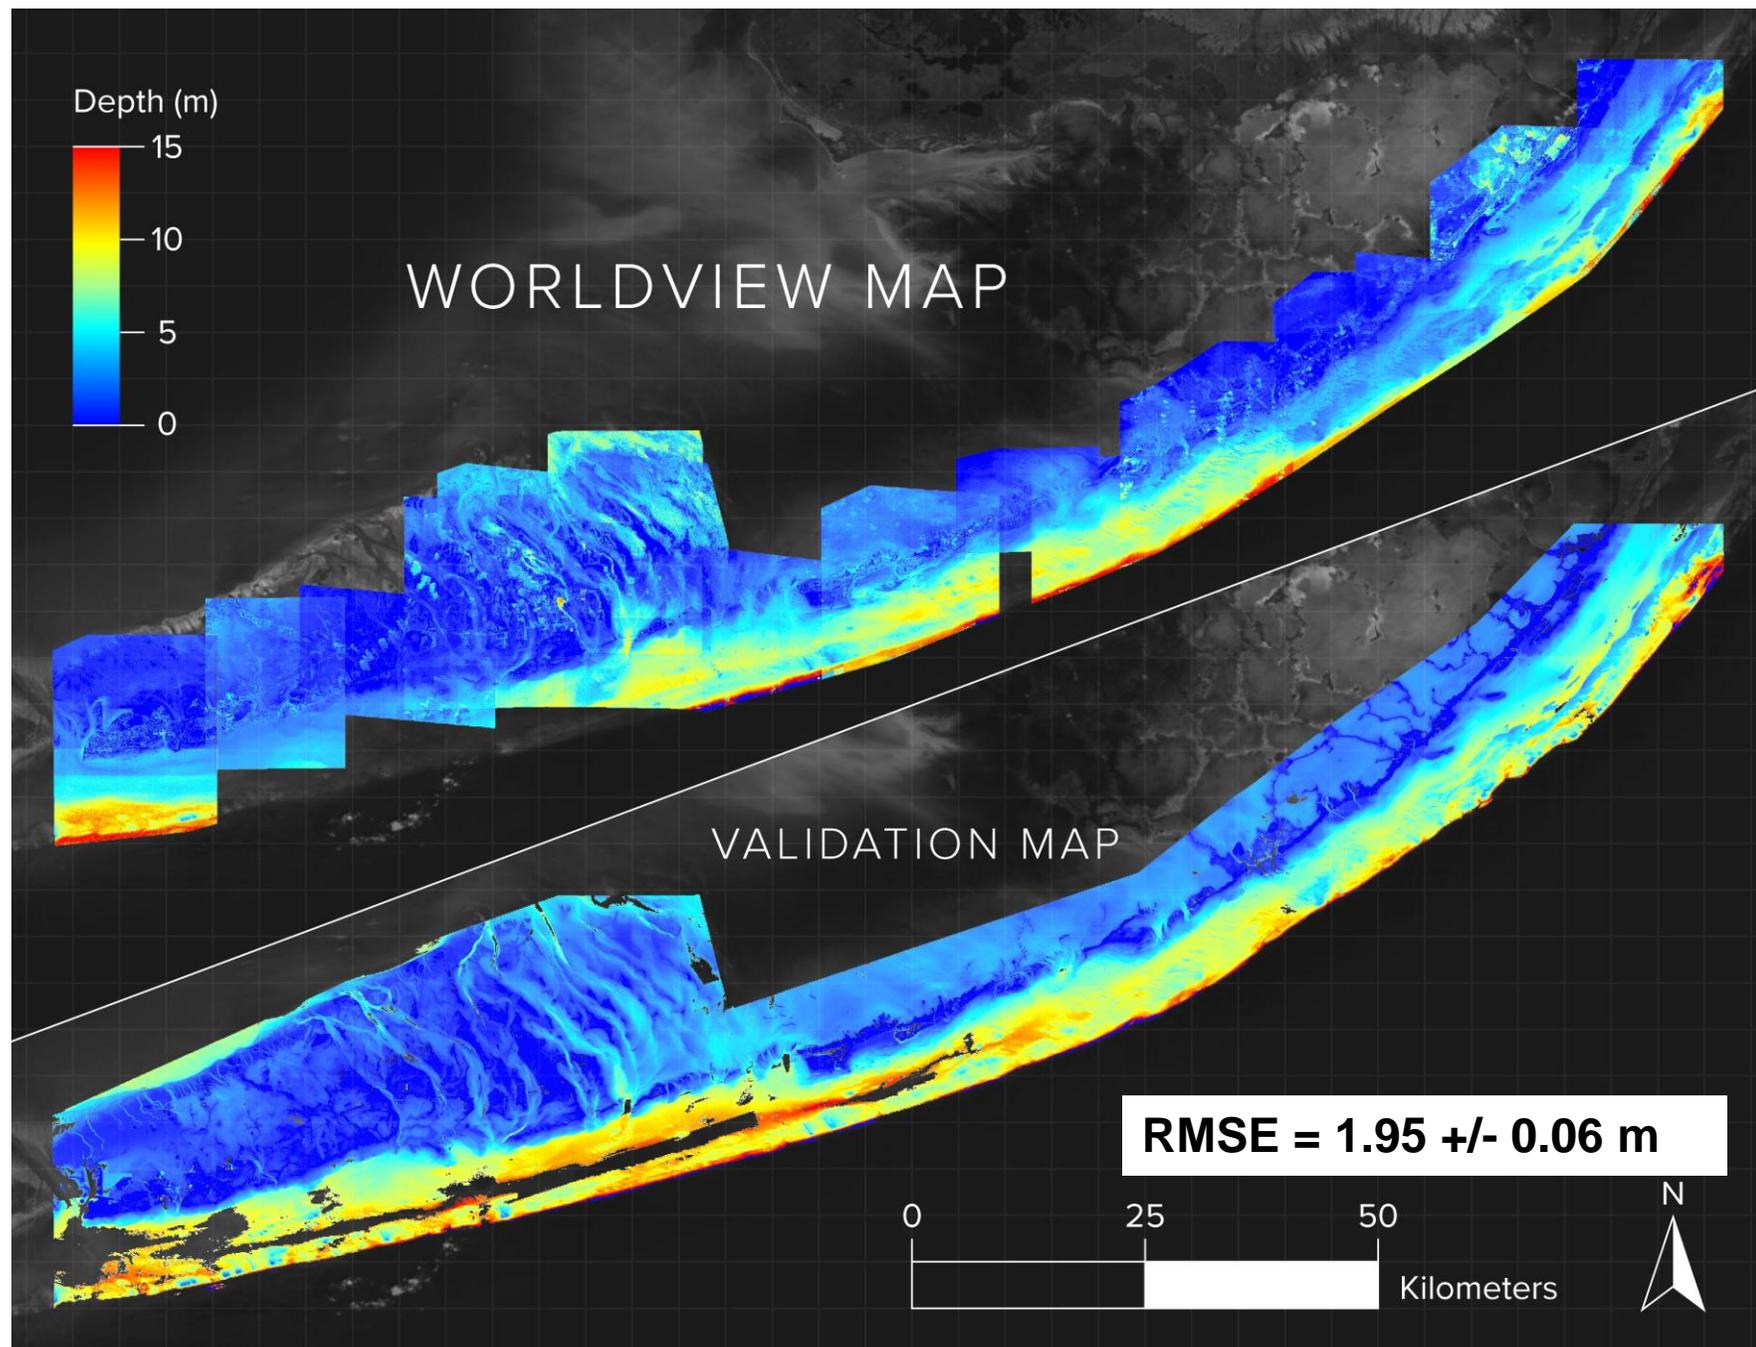


Results

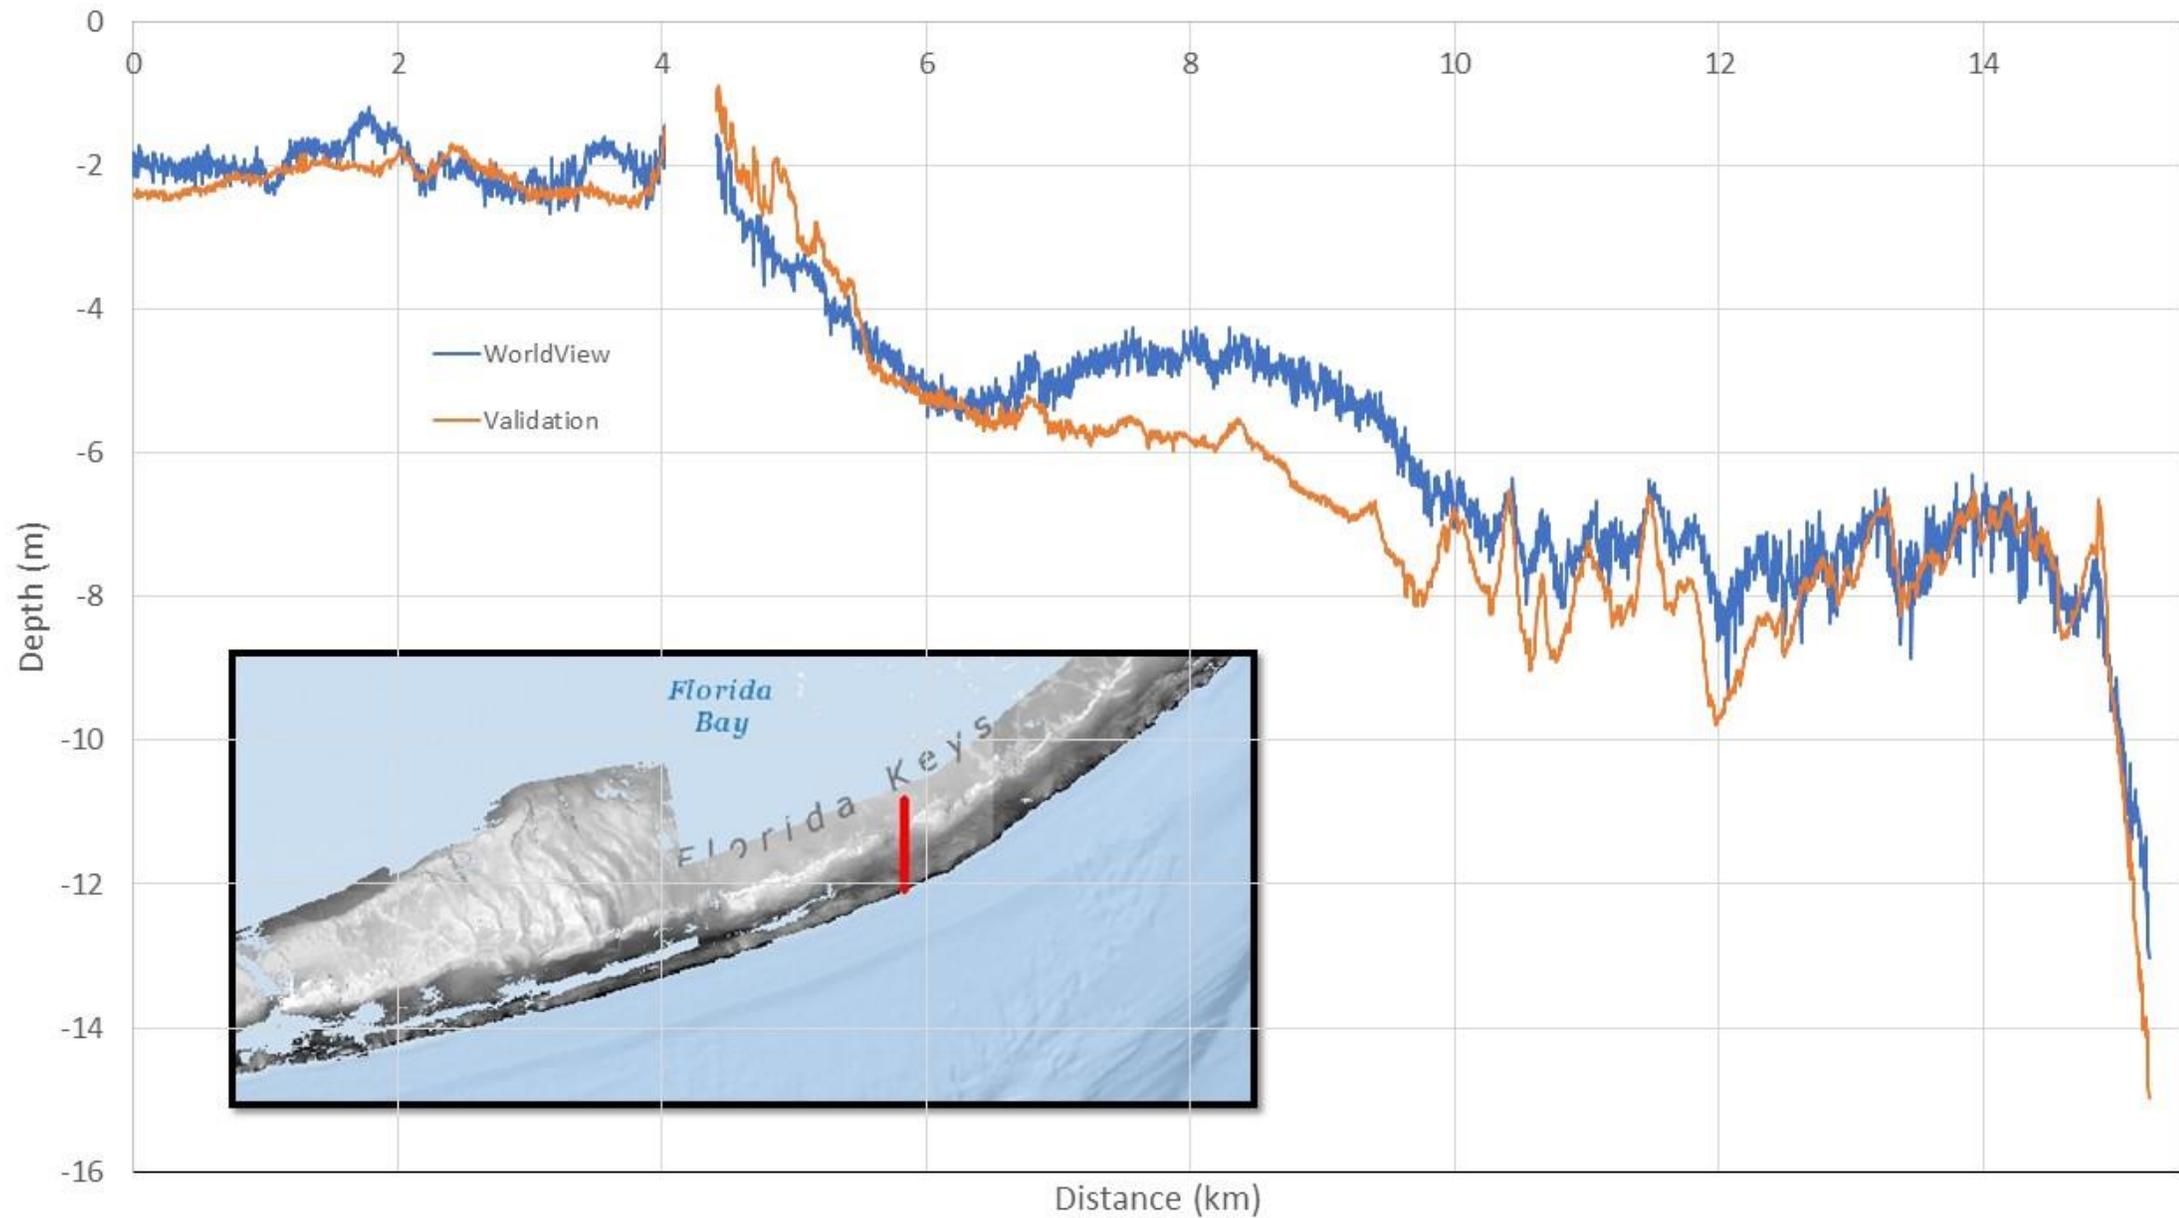
Depth	RMSE (m)	n
All	1.95	635,081
<12m	1.90	629,169
<10m	1.85	620,387
<8m	1.73	583,292
<6m	1.52	500,988
<4m	1.28	397,549
<2m	1.16	173,431

Processing efficiency:

- ~47 seconds per image
- 27 minutes total
- 138 km² mapped per minute

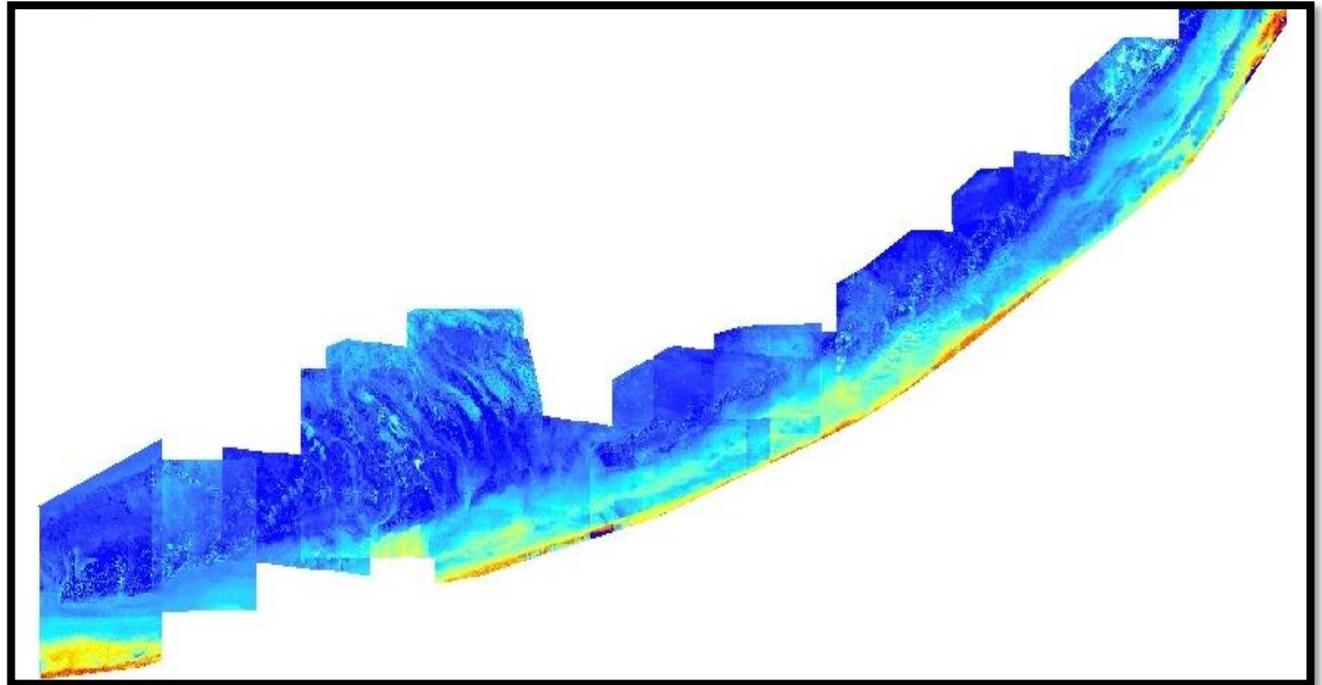


Results



Summary

- Accuracy (RMSE) = 1.95 m
 - Li et al. (2019): 1.22-1.86 m
 - Kerr et al. (2018): 0.89-2.62 m
 - Zhu et al. (2020): 3.829 m
 - DEM: 1.96 m
- Speed
 - < 1-minute per image
 - > 100 km² per minute
- Future work:
 - **Atmospheric correction enhancement**
 - **IOP modeling enhancement**



Thank you for your time!

We look forward to learn difficult problems regarding satellite, UAS, airborne, terrestrial remote sensing and computing that we could help solve and how we can help the community of interest!

David Hughes (presenter / team lead UAS edge, hughesdc@ornl.gov)

Matthew McCarthy (team lead on bathy, mccarthymj@ornl.gov)

Chris DeRolph (team lead on ecology project, derolphcr@ornl.gov)