Full Lifecycle of EO-based AI Foundation Models: Lessons Learned from Real-World Deployment

Yawen Zhang, Ph.D., Allen Institute for Artificial Intelligence
https://allenai.org/earth-system

Ai2

Ai2 is a Seattle-based nonprofit AI research institute founded in 2014 by the late Paul G. Allen. We build breakthrough AI to solve the world's biggest problems.

AI for the Planet

Ai2

AI for the Planet

Skylight

State-of-the-art AI to fight illegal fishing

Wildlands

AI to enhance forest resiliency and management

Satlas

Computer vision to monitor our changing planet

EarthRanger

Real-time data to protect wildlife and ecosystems

Climate Modeling

Better data and technology to tackle global challenges

Outline

1. Full Lifecycle of EO-based Foundation Model
    a. Skylight
    b. Lessons learned from real-world deployment
2. What is next?
    a. Earth System Platform (ESP)
    b. Helios - ESP foundation model

What is Skylight?

A maritime analyst tool for identifying suspicious behavior that may be illegal or non-compliant with fisheries and other maritime regulations.

VISION: Healthy, productive, and resilient oceans where targeted monitoring and enforcement actions support transparency and effective governance of marine resources.

SKYLIGHT
A product of Ai2

Website: https://www.skylight.global/
Email: support@skylight.global

Ai2

Skylight: Global-Scale Real-Time Vessel Detection

VIIRS Sentinel1 Sentinel2 Landsat

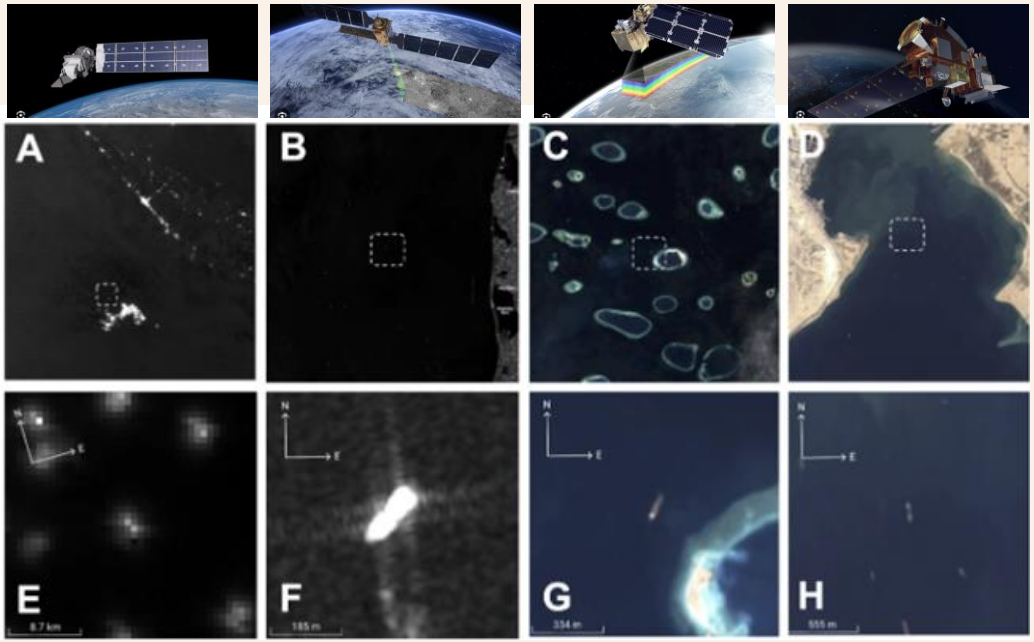

- **Reliable** (99.99% uptime)
- **Highly performant** (80+ accuracy)
- **Timely** (real-time / low latency)
- **Highly resource efficient** (runs on small GPUs, T4s)
- **Open-source** (model + APIs)

Covering all Exclusive Economic Zones (EEZ) of every country.

Already in use by hundreds of organizations worldwide.

Ai2

Beukema et. al. (NeurIPS 2024)

Skylight Demo

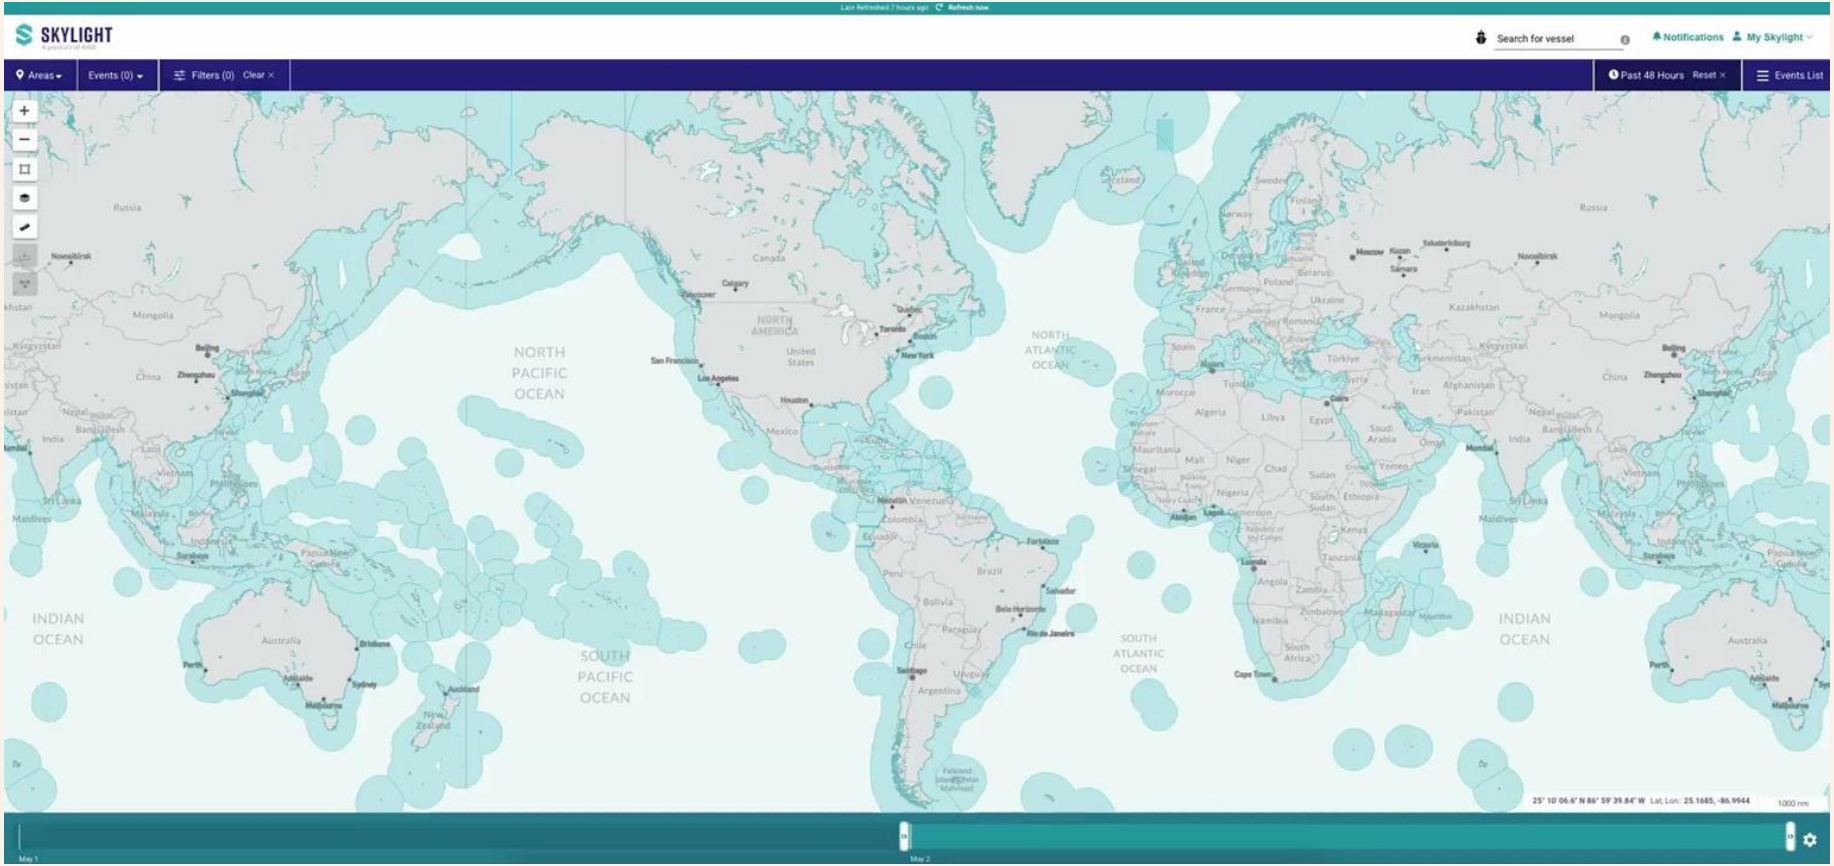

Ai2

In-Depth: Full Lifecycle of the Skylight Landsat Model

Image Sampling & Annotations

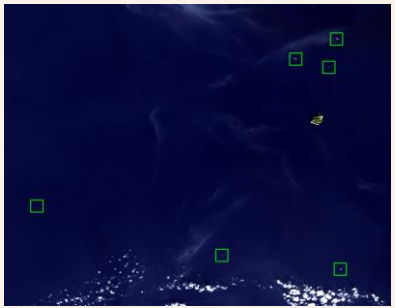

18K vessel labels (split into train/validation)

Model Development & Training

Image Tokens

↓

Pretrained Swin Transformer

↓

Object Detector

↓

Vessel Detections

Model Evaluation

Offline Validation Results:
F1-Score: 78.8%

Model Deployment

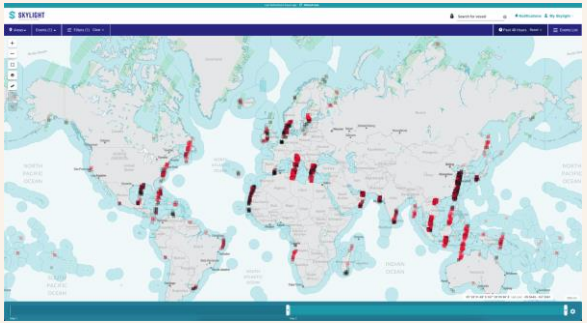

Skylight Platform

Lesson Learned: Offline ≠ Online Performance

Massive gap between **offline** (R&D) and **online** (production) performance

Landsat model v0 - Detector only (**18K** vessel labels)

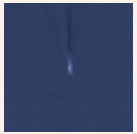
Vessel

False Positives!


Island


Whitecap


Iceberg


Cloud

Landsat model v1 - Detector + **Classifier** (**2K** vessel & False Positive labels)

+ **Evaluation by Experts** (88.0% Good, 10.3% Bad, 1.7% Unsure)
+ **Deployed** in the Skylight Platform

In-Depth: Full Lifecycle of the Skylight Landsat Model

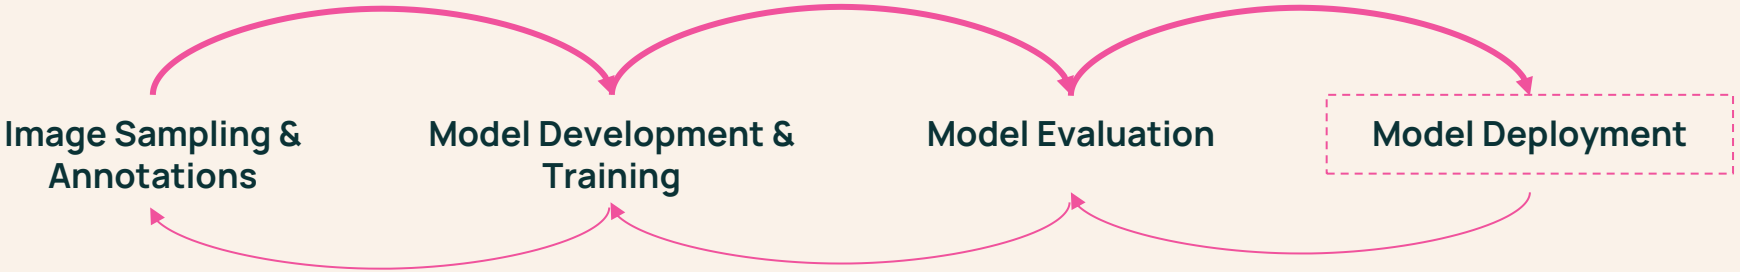

- **Training samples:** diverse, expert-labeled
- **Evaluation:** a wide range of examples and scenarios
- **User feedback:** model improvement

Lesson Learned: Development to Deployment

Key Practices for Reliable ML Deployment

- **Model version control**
- **Continuous integration**: scenario checks
- **Model as an API**: containerized FastAPI application

Changelog

- `v0.0.1` : Initial model release. Offline evaluation metrics reported.
- `v0.0.2` : Improved pansharpening, added error message to LandsatResponse.
- `v0.0.3` : Increase crop size.
- `v0.0.4` : Add Prometheus timers.
- `v0.0.5` : Fix for new version of rslearn.
- `v0.0.6` : Fix bug with RGB crops.
- `v0.0.7` : Fix Docker container bug.

Offline Scenario Checks

Mode	Status
True Positive - Dense vessels	✅ Pass
True Positive - Sparse vessels	✅ Pass
False Positive - Icebergs	✅ Pass
False Positive - Clouds	✅ Pass
False Positive - Whitecaps	✅ Pass

Integration tests

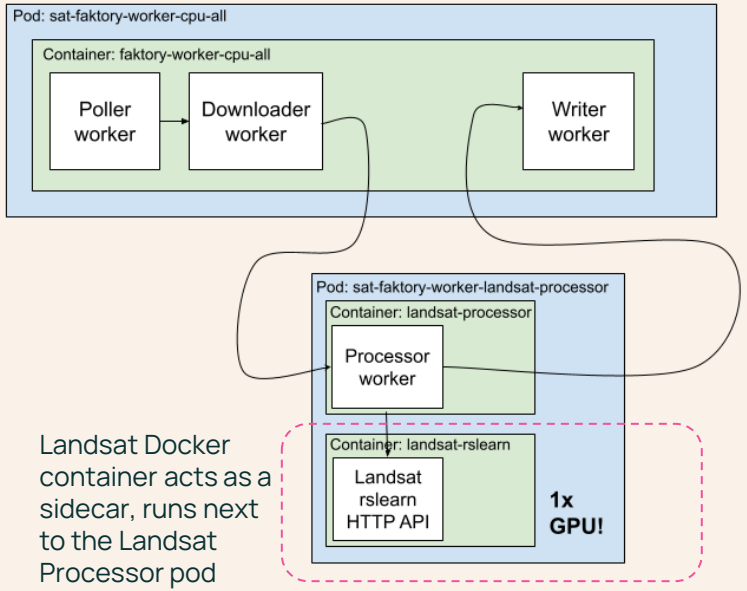

Landsat Docker container acts as a sidecar, runs next to the Landsat Processor pod

AI for the Planet

Skylight

State-of-the-art AI to fight illegal fishing

Wildlands

AI to enhance forest resiliency and management

Satlas

Computer vision to monitor our changing planet

EarthRanger

Real-time data to protect wildlife and ecosystems

Climate Modeling

Better data and technology to tackle global challenges

More applications…

Ai2

What is Next?

Earth System Platform (ESP)

Earth System

- Automated deforestation detection
- High resolution global ecosystem analysis
- Predicted human-wildlife conflicts
- Forecasted impacts on marine ecosystems
- Scalable predictions of wildfire risk
- Automated land use and crop health monitoring
- Super-resolution climate forecasts

Global scale, real-time, multimodal data

Our foundation models are capable of leveraging satellite imagery, sensor data, radar, climate forecasts, time series, and more

A fully integrated AI solution

We provide a scalable, end-to-end system—from data ingestion to model deployment—delivered through open, intuitive APIs and interfaces

Empowered partners worldwide

Makes it fast and frictionless for partners to build powerful solutions suited to their geographies and communities–no AI expertise required

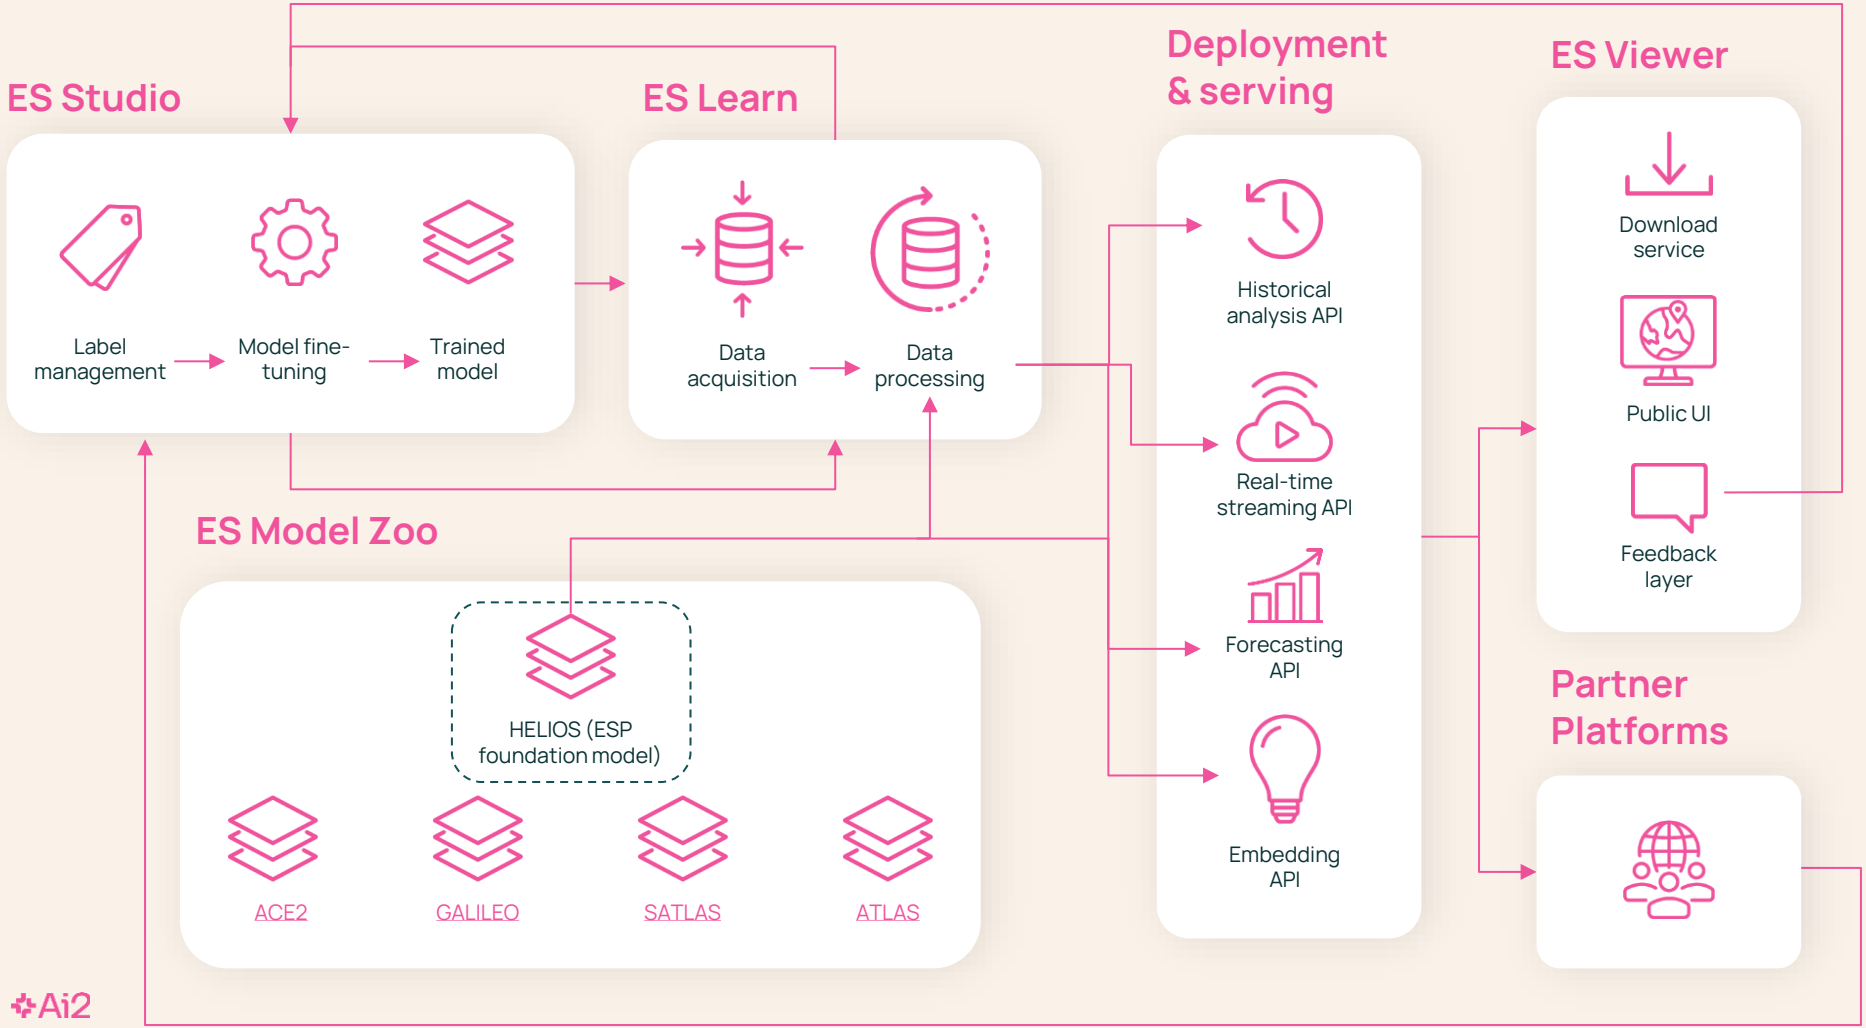
ES Studio
Label management
Model fine-tuning
Trained model
ES Learn
Data acquisition
Data processing
Deployment & serving
Historical analysis API
Real-time streaming API
Forecasting API
Embedding API
ES Viewer
Download service
Public UI
Feedback layer
Partner Platforms
ES Model Zoo
HELIOS (ESP foundation model)
ACE2
GALILEO
SATLAS
ATLAS
Ai2

Helios - ESP Foundation Model

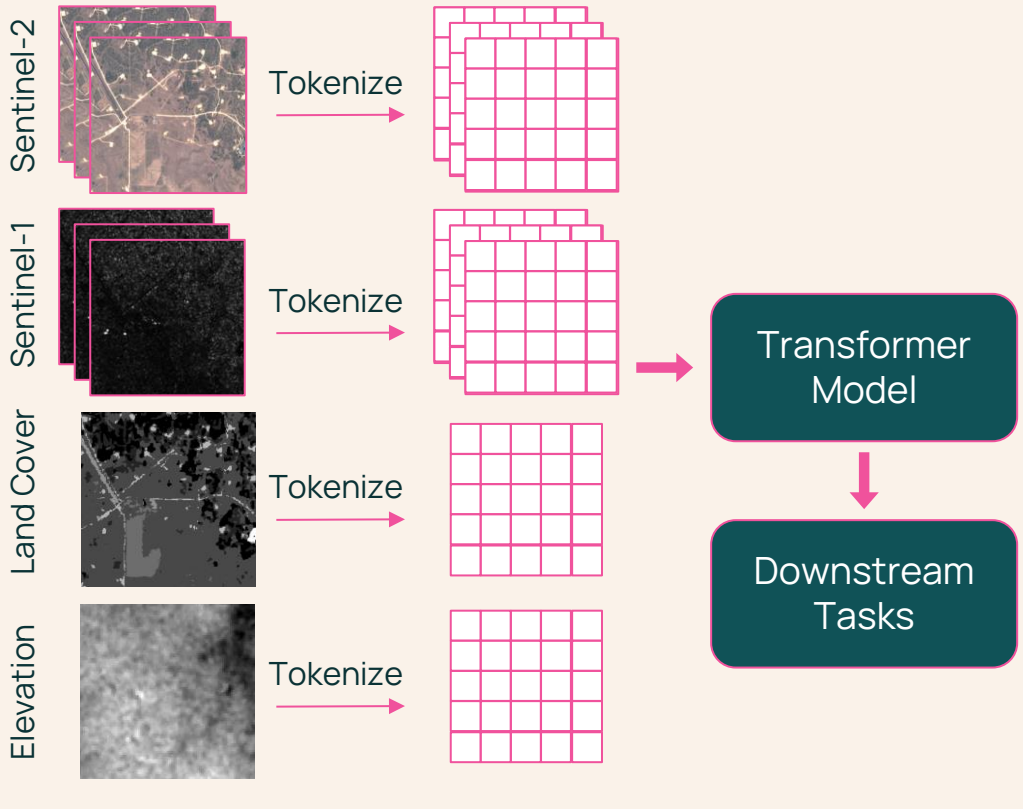

- Globally diverse samples for pretraining
- Multimodal, multi-resolution inputs
- Self-supervised learning
- Multiple model sizes
- Flexible image & patch size
- Support various downstream tasks:
  - Image (time-series) classification
  - Image (time-series) segmentation
  - Pixel (time-series) classification
  - Change detection
  - Object detection
  - ...

Stay tuned!

Forest Loss Monitoring

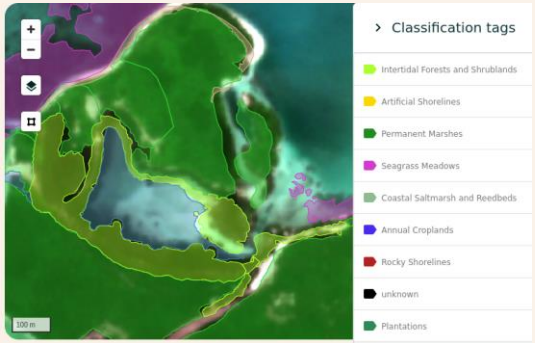

Ecosystem Type Mapping

IUU Fishing Detection

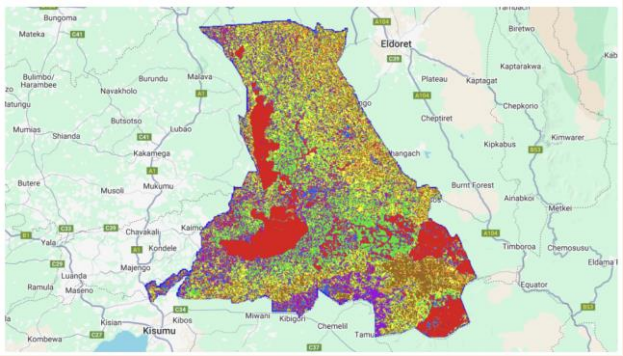

Smallholder Agriculture:
Crop Type Mapping

Global Mangrove Monitoring

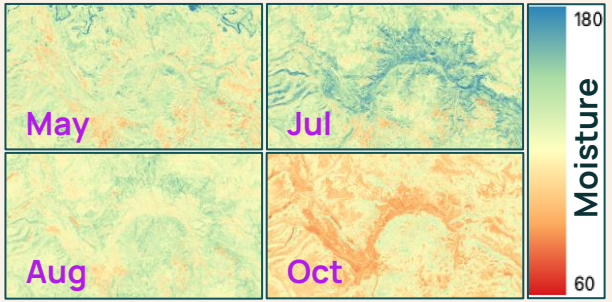

Wildfire Risk Assessment

Ai2

Thank you.
Questions?

https://allenai.org/earth-system

earthsystem@allenai.org

Ai2