TESSERA

Zhengpeng(Frank) Feng

Oct 1st , 2025

with C. Atzberger, S. Jaffer, J. Knezevic, S. Sormunen, R. Young, M. C. Lisaius, M. Immitzer, T. Jackson, J. Ball, D. A. Coomes, A. Madhavapeddy, A. Blake and S. Keshav

Department of Computer Science and Technology
Energy and Environment Group
zf281@cam.ac.uk

Why Observe Earth from Space?

Monitor

- Land use
- Biodiversity

Assess status of

- Crops
- Soils
- Forests

Quantify

- Forest degradation
- Deforestation
- Carbon stock

Remote Sensing Allows Monitoring at Scale

- Continuous monitoring of land surface
- Different, complementary sensors: optical, radar, lidar

Lots of public satellite data!

10 million
Landsat
TM / ETM+ / OLI

31 million
Sentinel-2
L1C

2 million
Sentinel-1
GRD

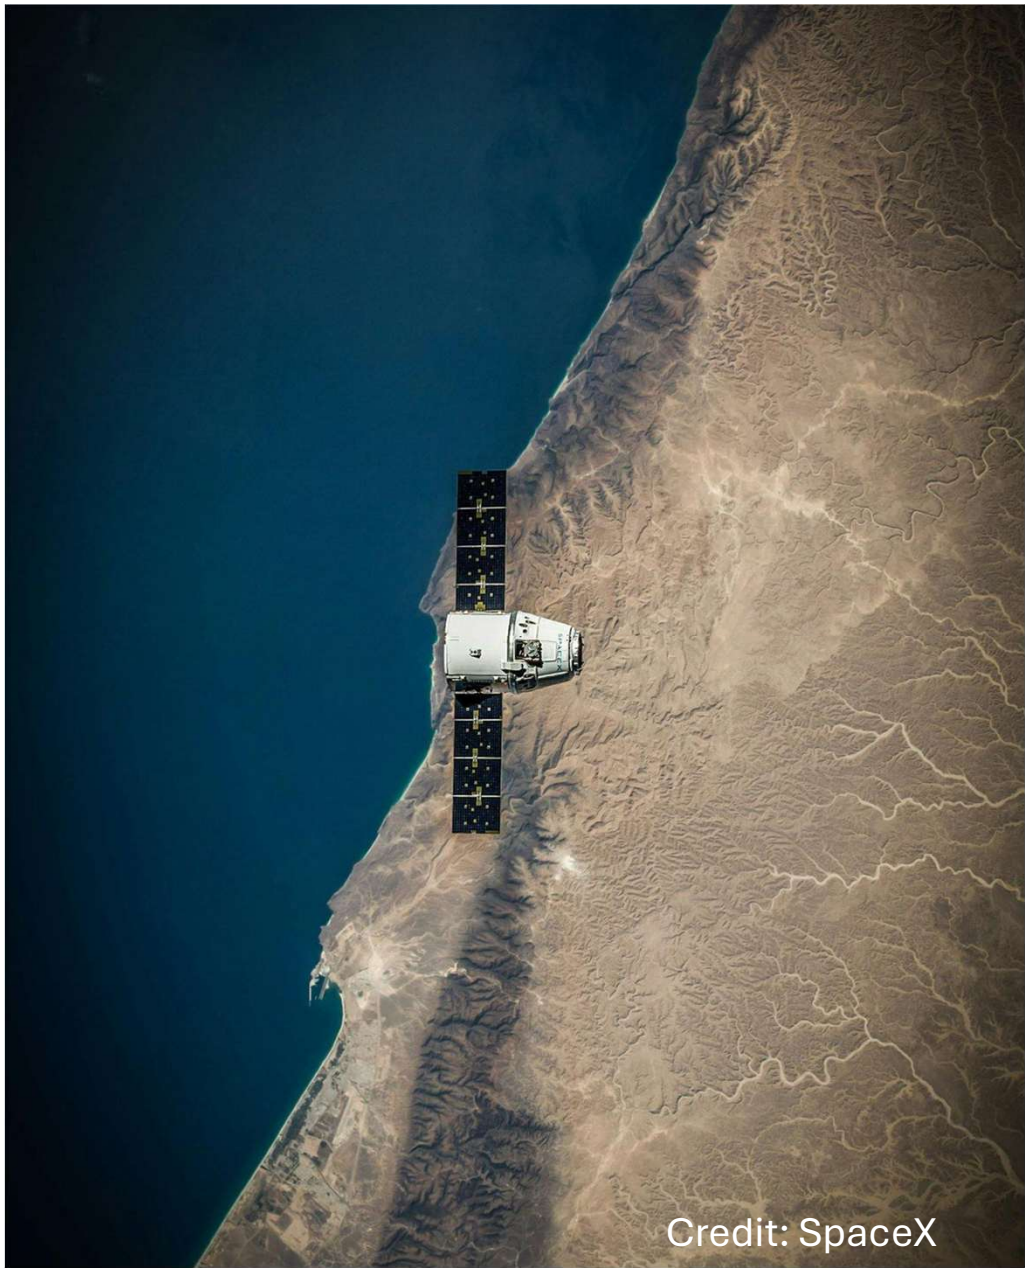

Challenges

- Clouds
- Changes in lighting
- Varying time gaps
- Sensor fusion
- Temporal compositing

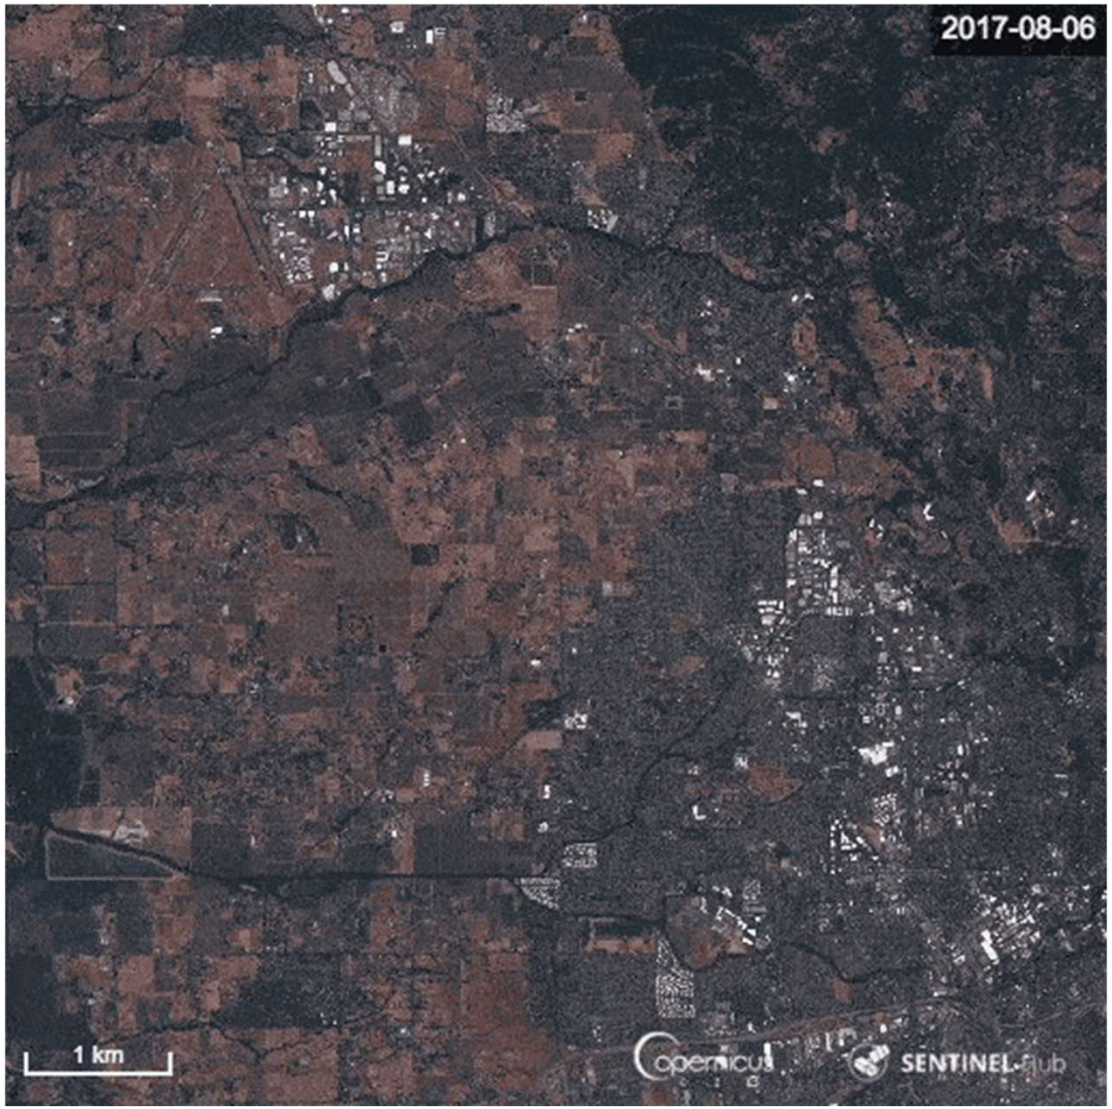

https://www.geographyrealm.com/new-tool-sentinel-hub-lets-create-time-lapse-gifs-satellite-imagery/

The Temporal Signal is Critical

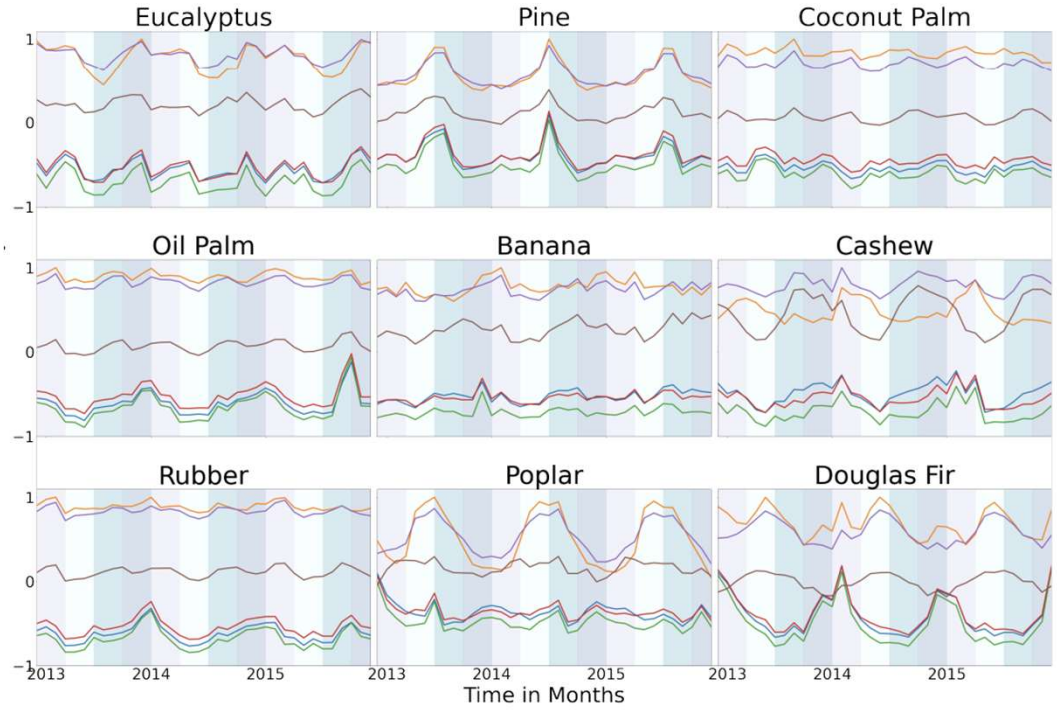

Pazos-Outon et al. (2024) (https://arxiv.org/abs/2406.18554)

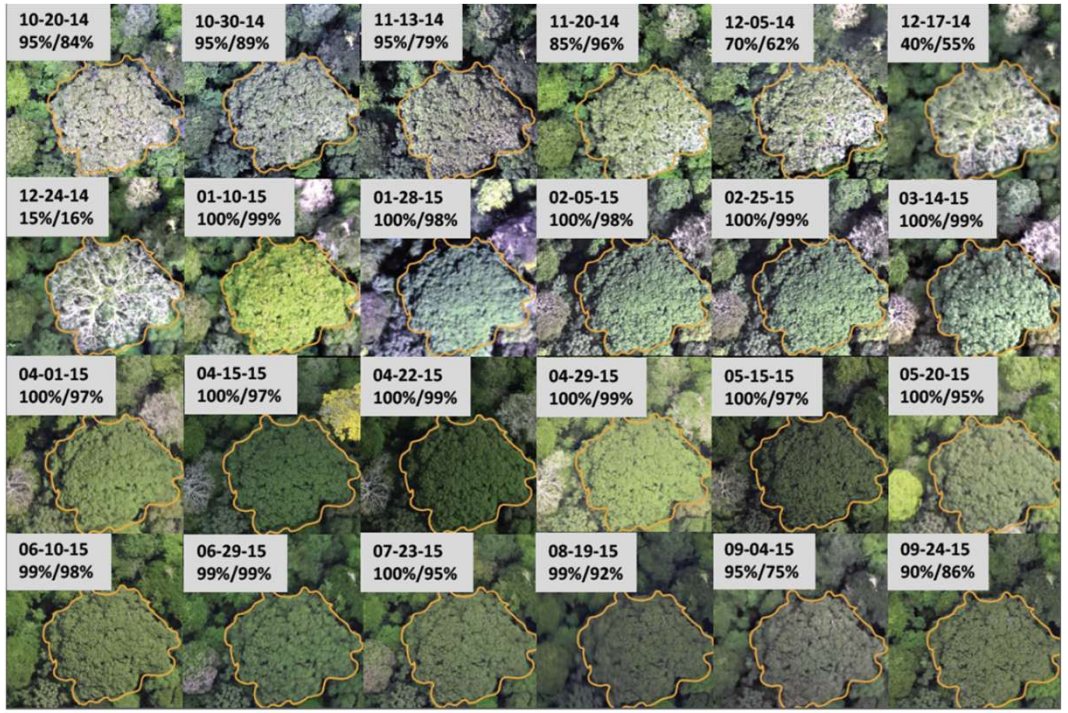

Park et al. (2019) (https://www.mdpi.com/2072-4292/11/13/1534)

Seasonal dynamics help identify land cover and detect change

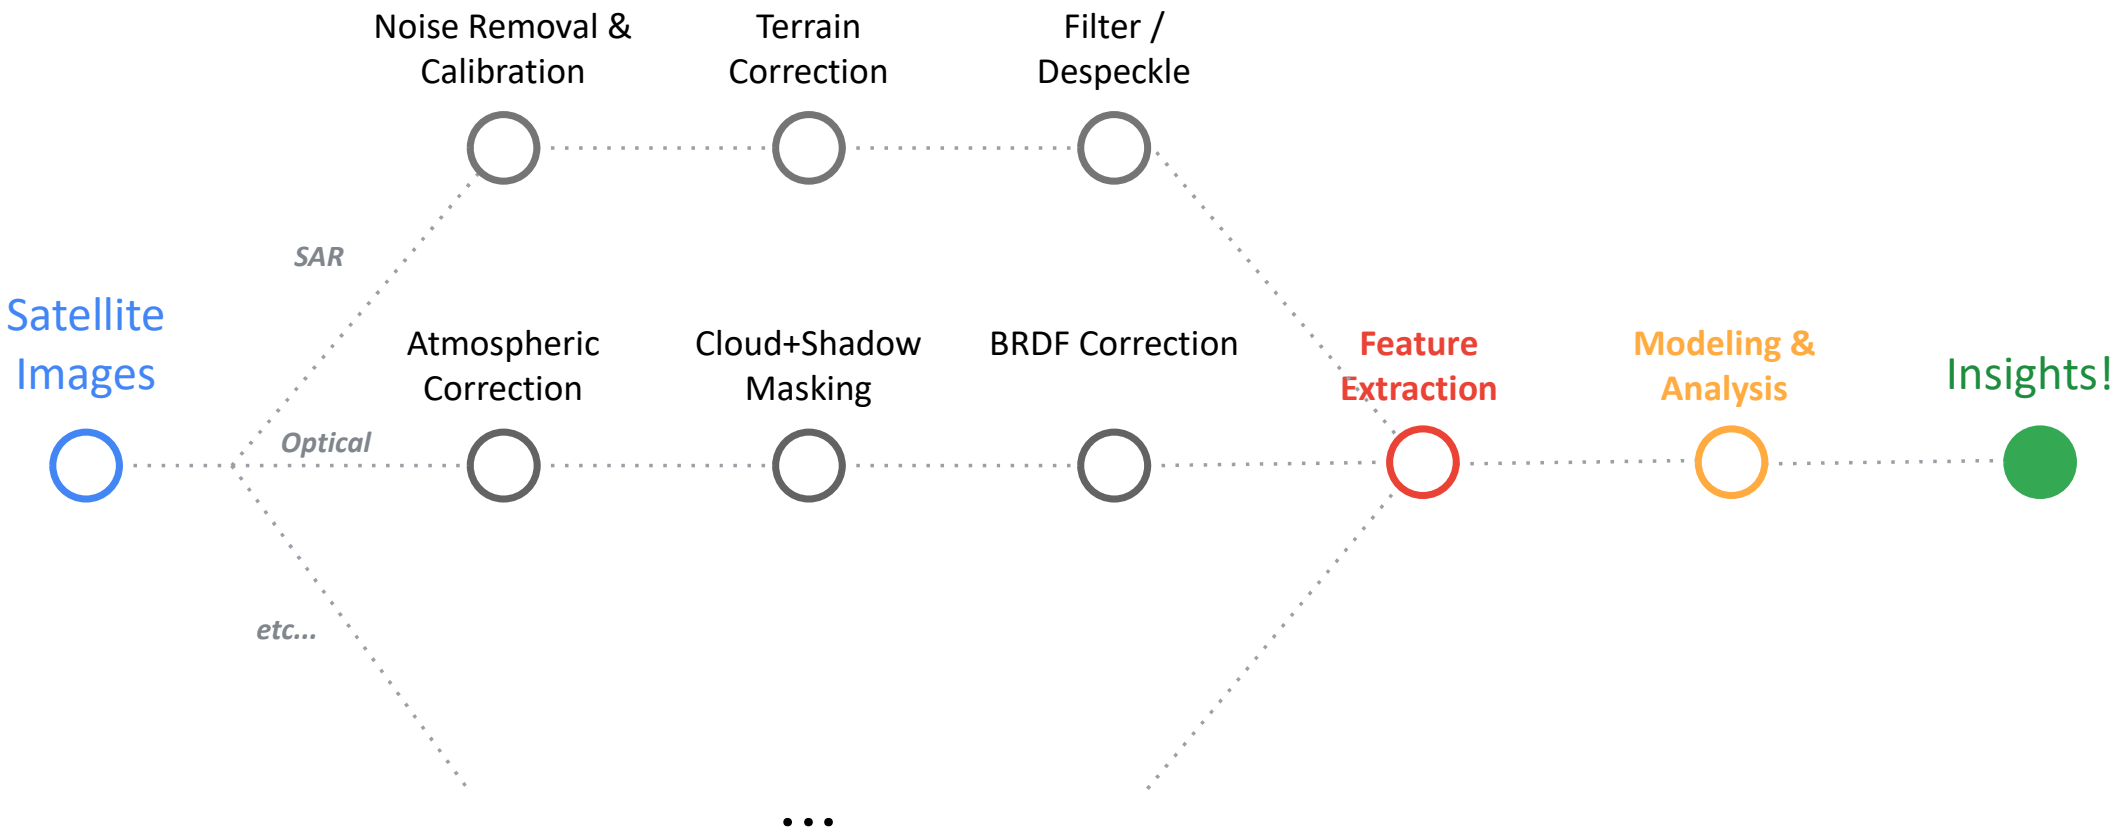

Image Credit: Google Alpha Earth

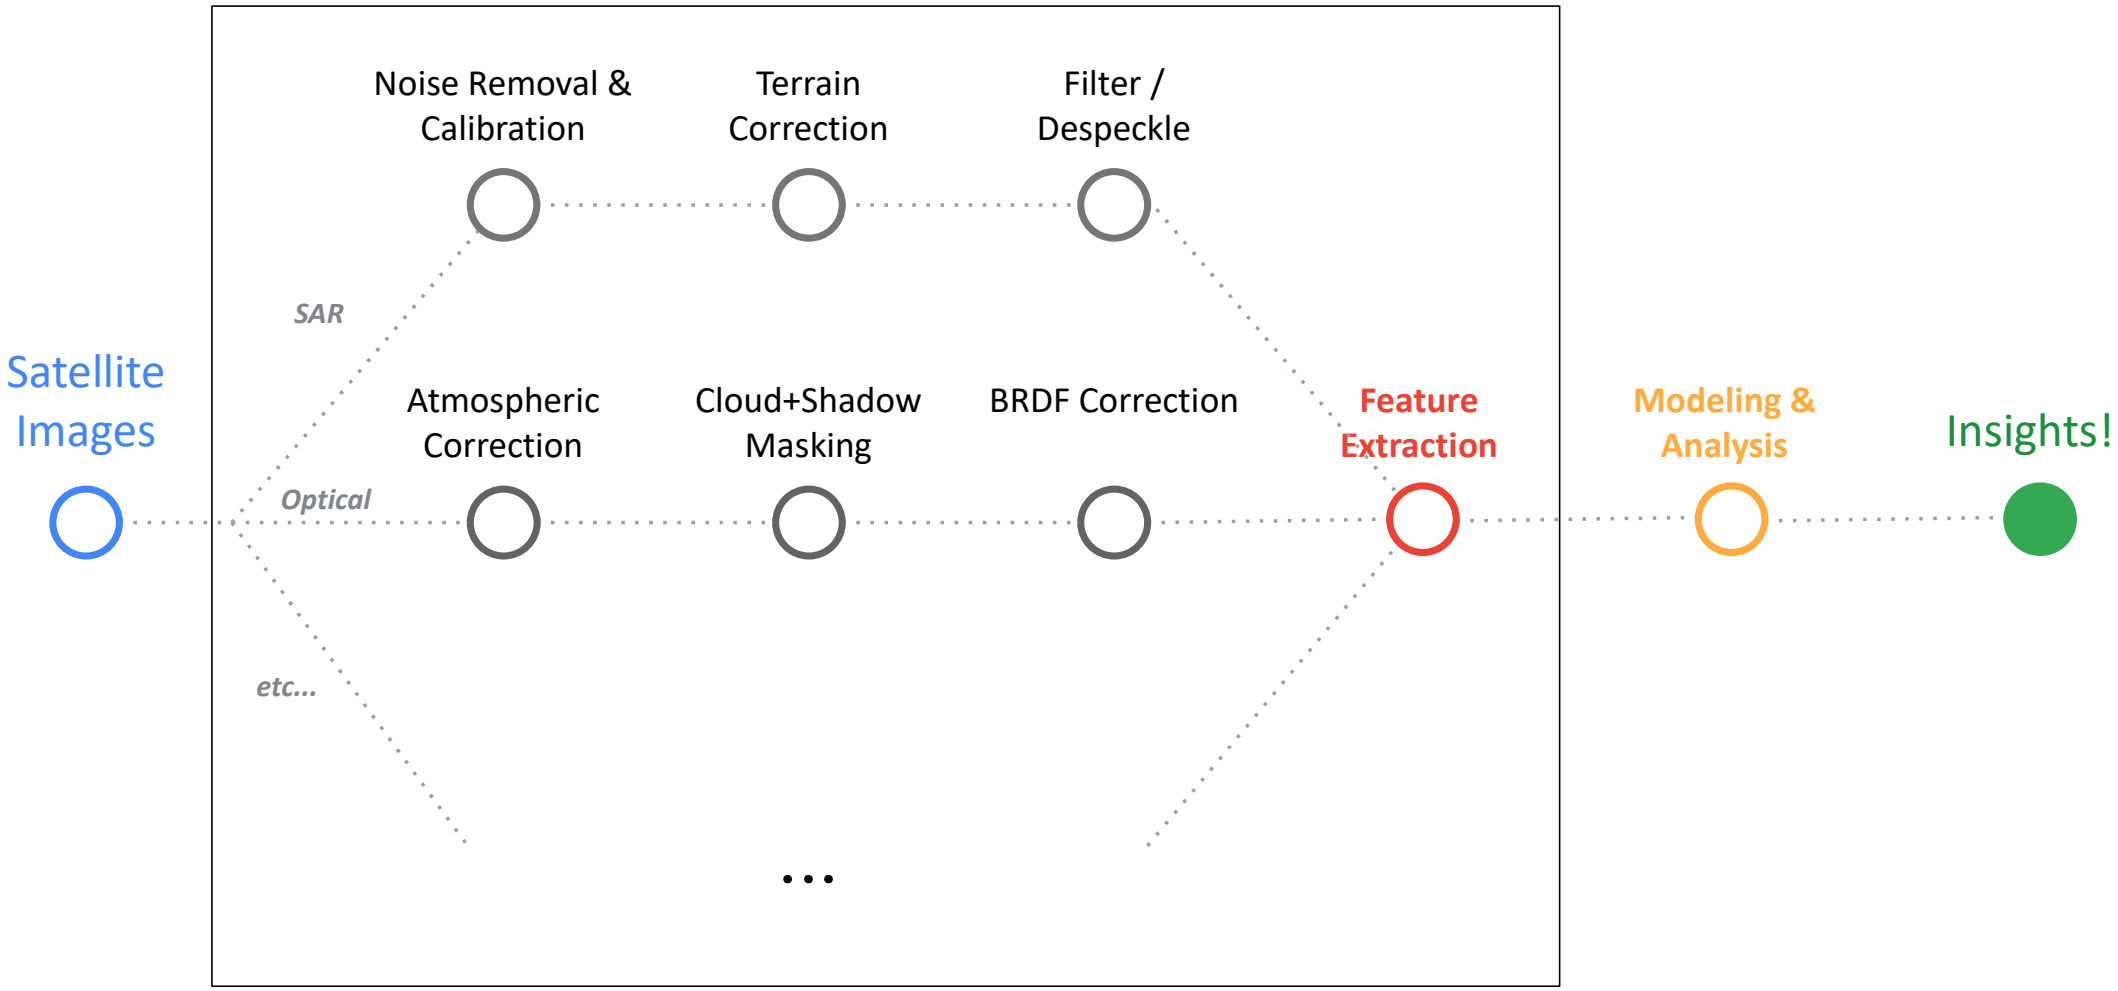

Image Credit: Google EFM

Satellite Images

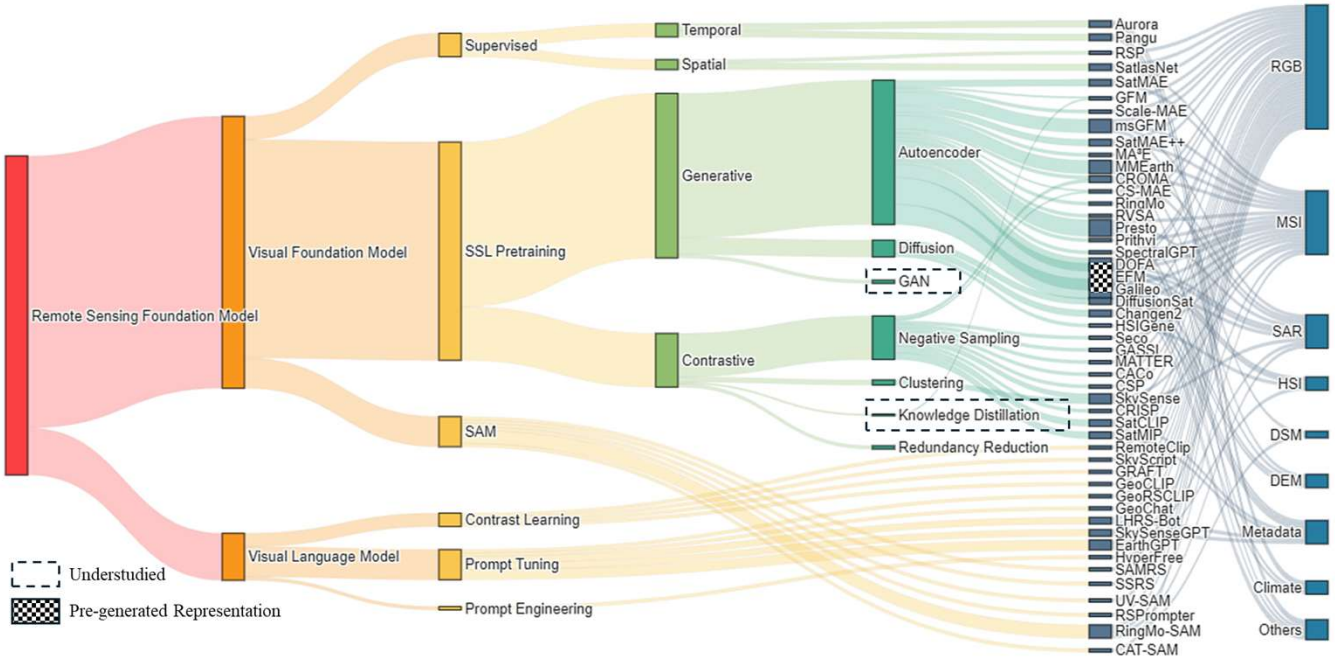

Modeling & Analysis

Insights!

Image Credit: Google EFM

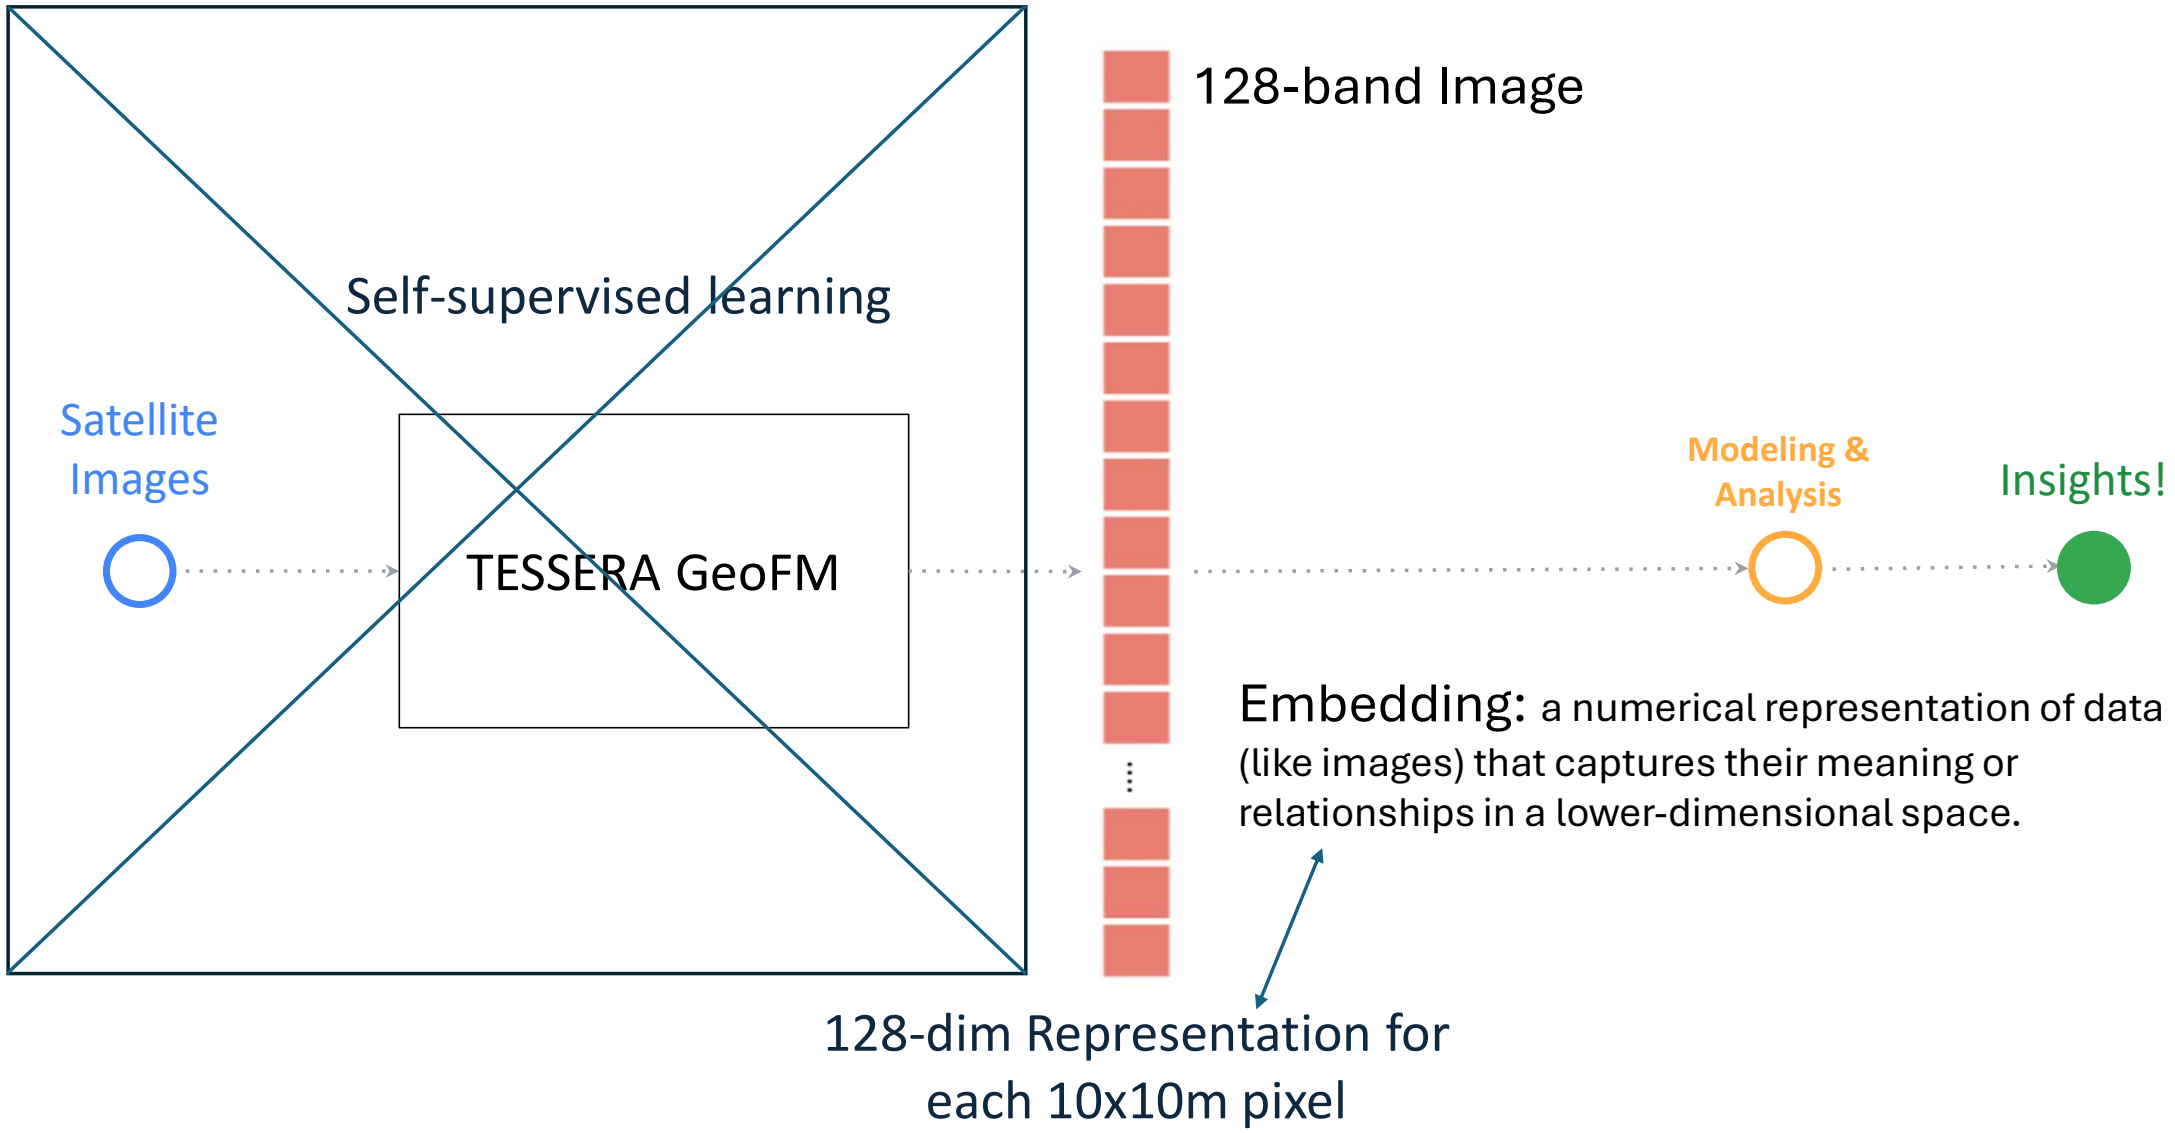
Self-supervised learning
Satellite Images
TESSERA GeoFM
128-band Image
Modeling & Analysis
Insights!
Embedding: a numerical representation of data (like images) that captures their meaning or relationships in a lower-dimensional space.
128-dim Representation for each 10x10m pixel

Self-supervised Learning

Masked Auto Encoding (MAE)

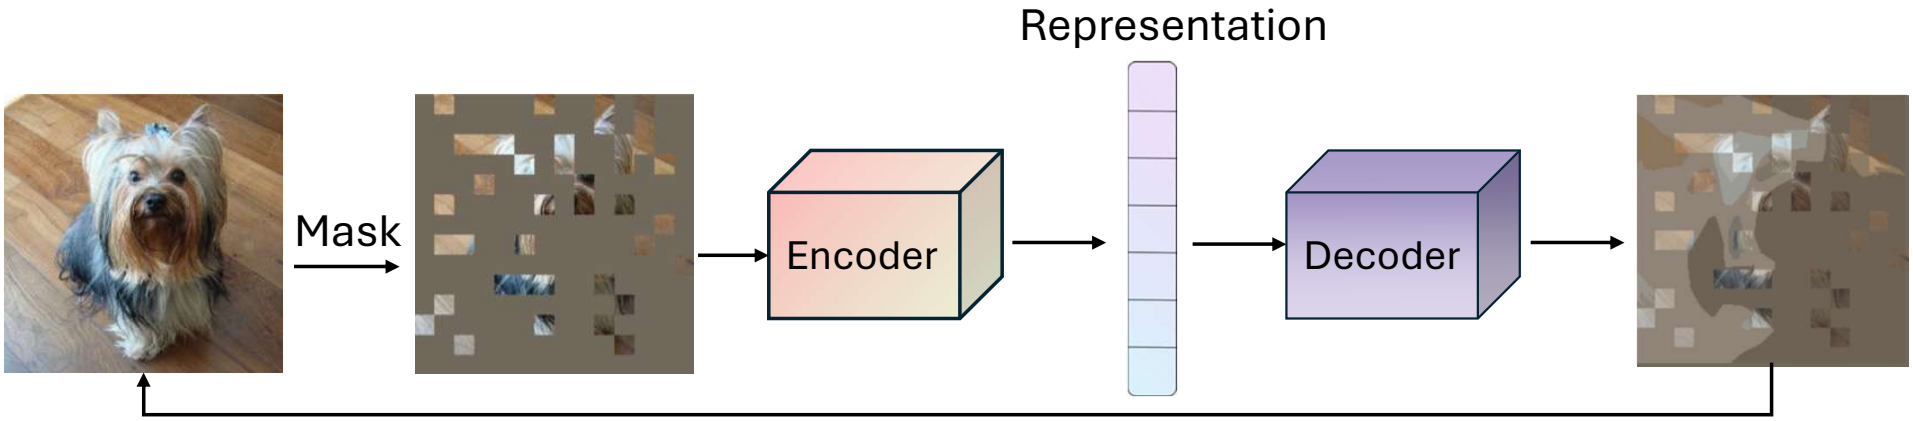

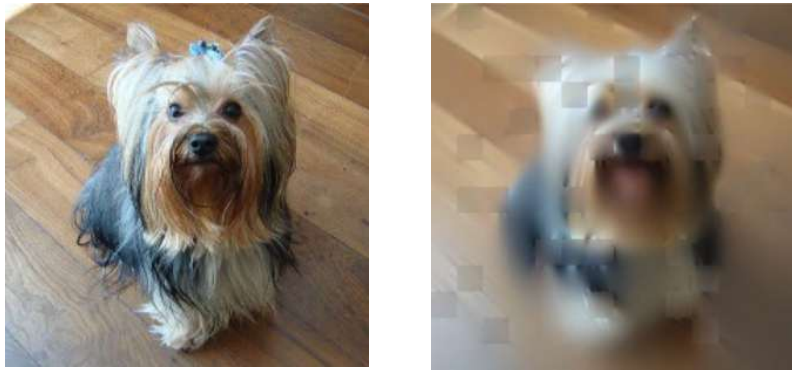

TESSERA Architecture

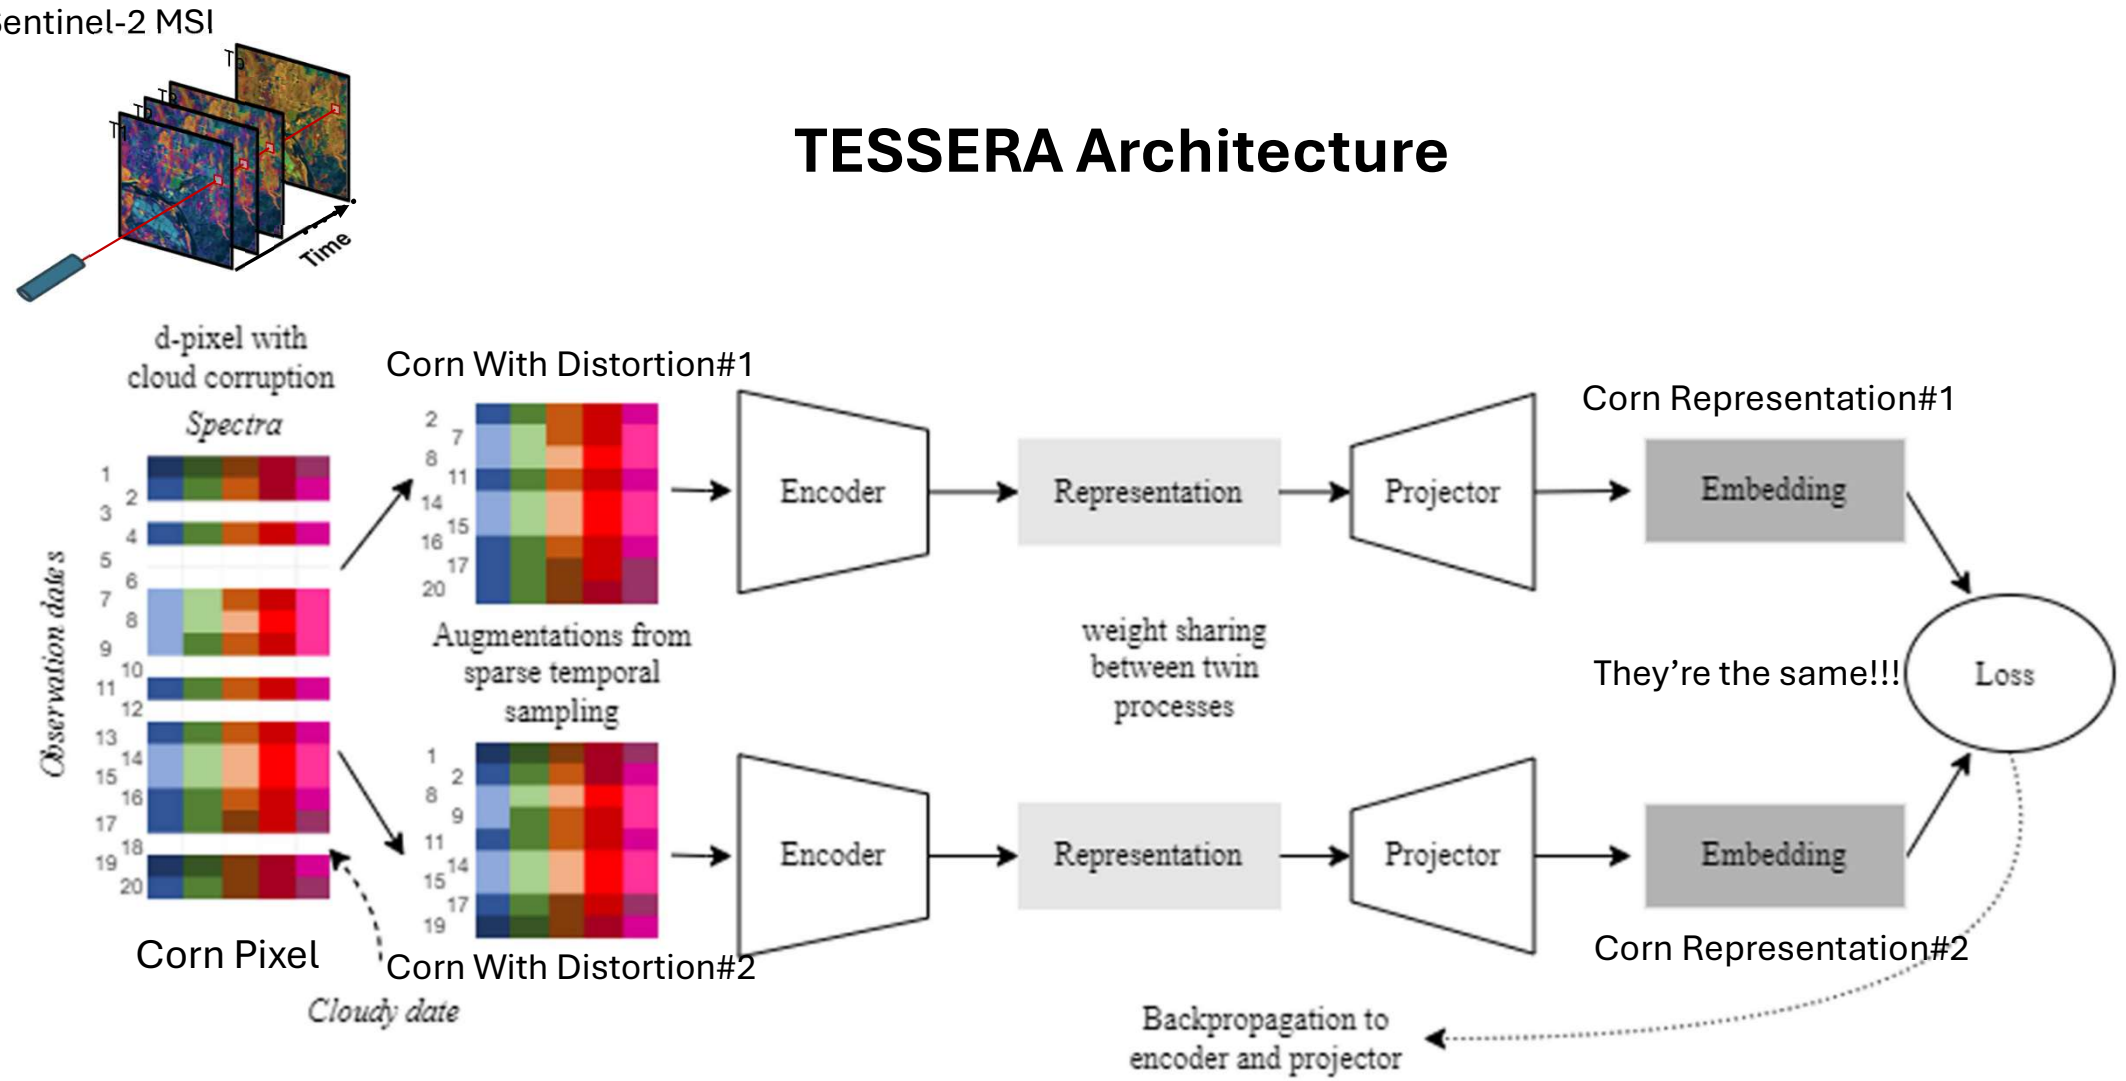

Lisaius et al. (2024) (https://ieeexplore.ieee.org/document/10592304)

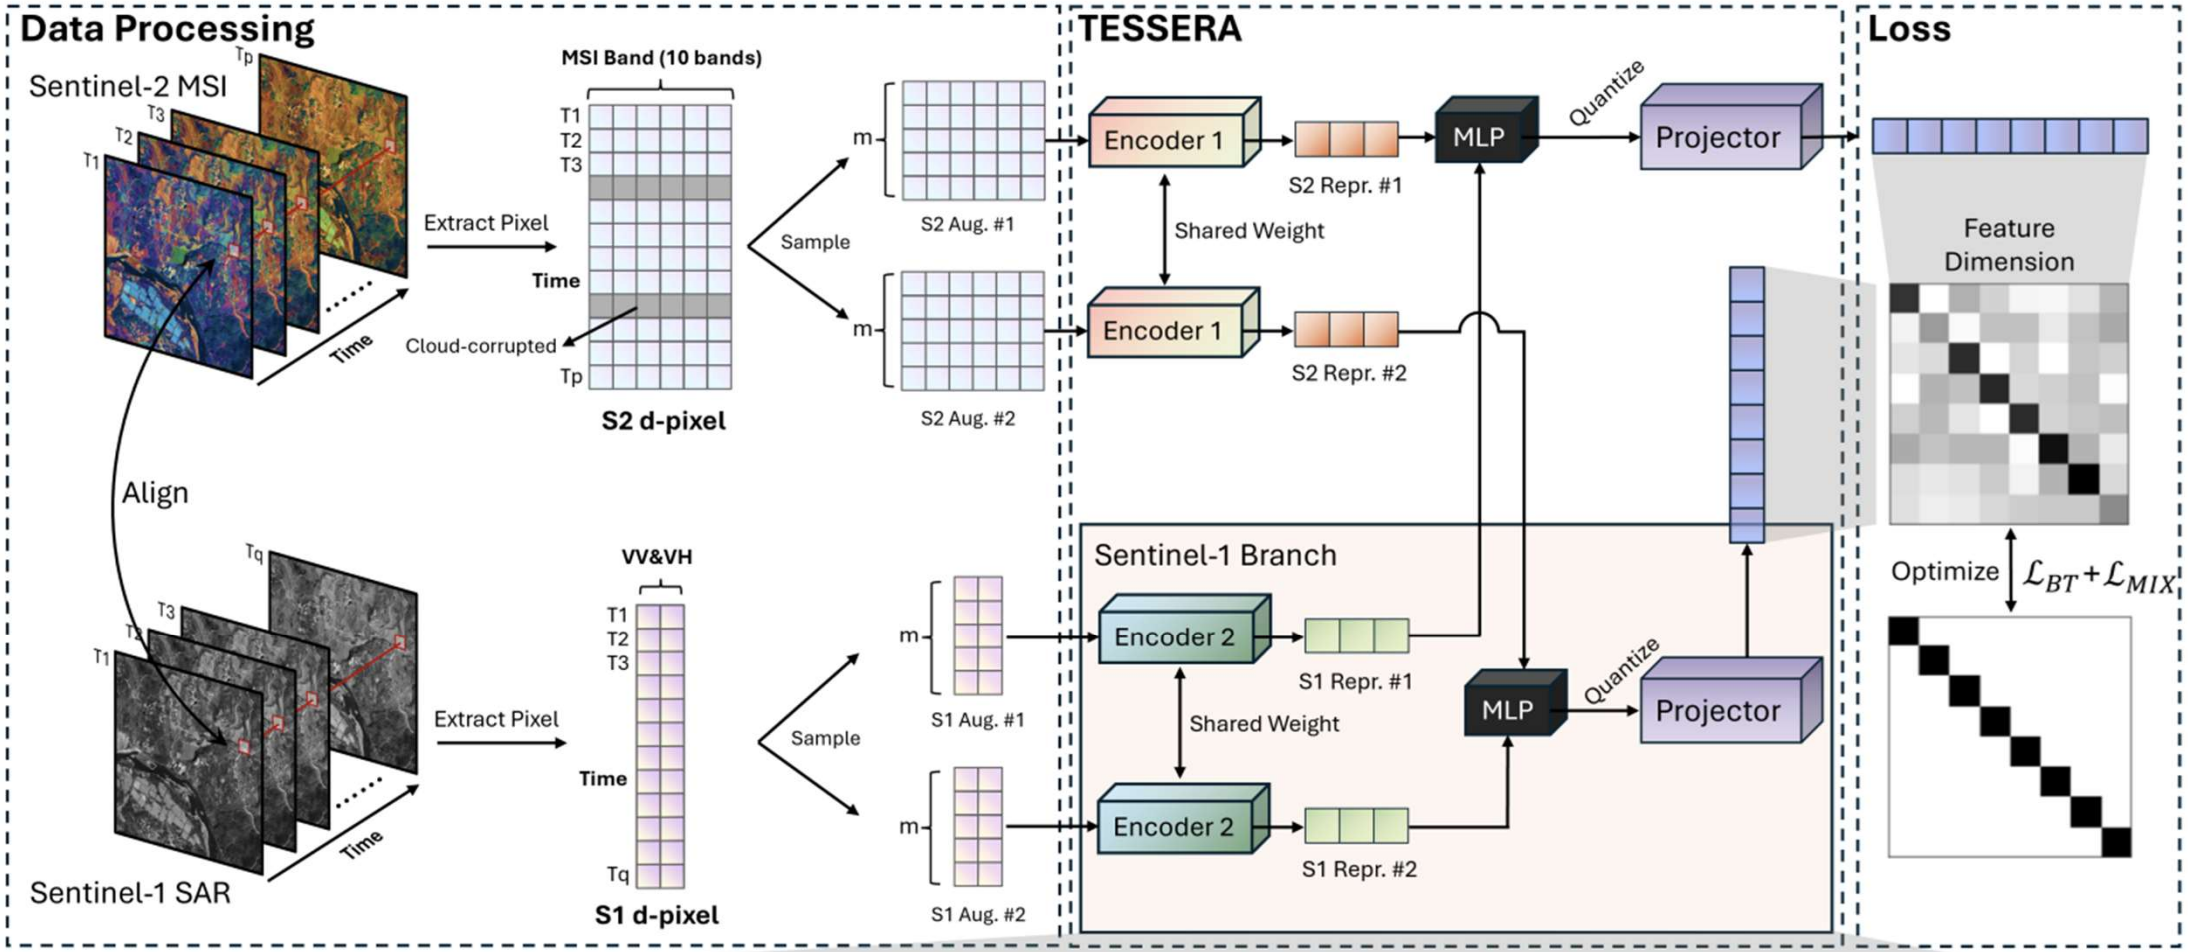
Data Processing
Sentinel-2 MSI
Tp
T3
T2
T1
Time
Extract Pixel
MSI Band (10 bands)
T1
T2
T3
Time
Cloud-corrupted
Tp
S2 d-pixel
Sample
m
S2 Aug. #1
S2 Aug. #2
Align
Sentinel-1 SAR
Tq
T3
T2
T1
Time
Extract Pixel
VV&VH
T1
T2
T3
Time
Tq
S1 d-pixel
Sample
m
S1 Aug. #1
S1 Aug. #2
TESSERA
Encoder 1
S2 Repr. #1
MLP
Quantize
Projector
Shared Weight
Encoder 1
S2 Repr. #2
Sentinel-1 Branch
Encoder 2
S1 Repr. #1
Shared Weight
MLP
Quantize
Projector
Encoder 2
S1 Repr. #2
Loss
Feature
Dimension
Optimize
$\mathcal{L}_{BT} + \mathcal{L}_{MIX}$

- Ready to use representations
- Cloud-free, encoding temporal spectral signal
- Sensor fusion: optical + radar data
- Globally available at 10m resolution, annually
- Open source

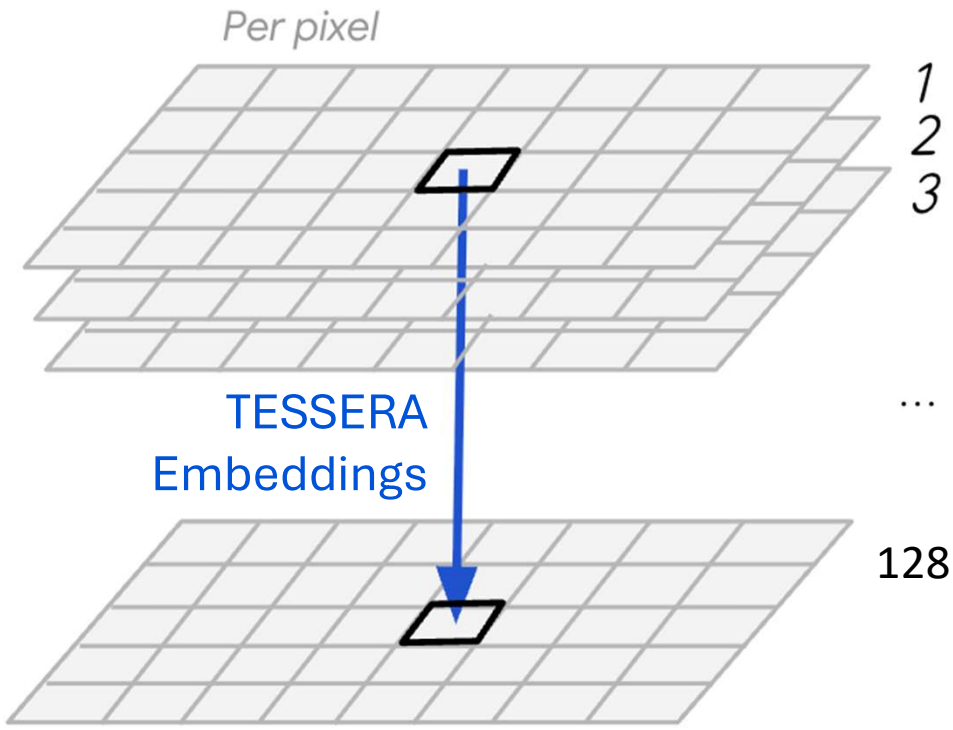

We use first 3 chennels out of 128 as R,G,B

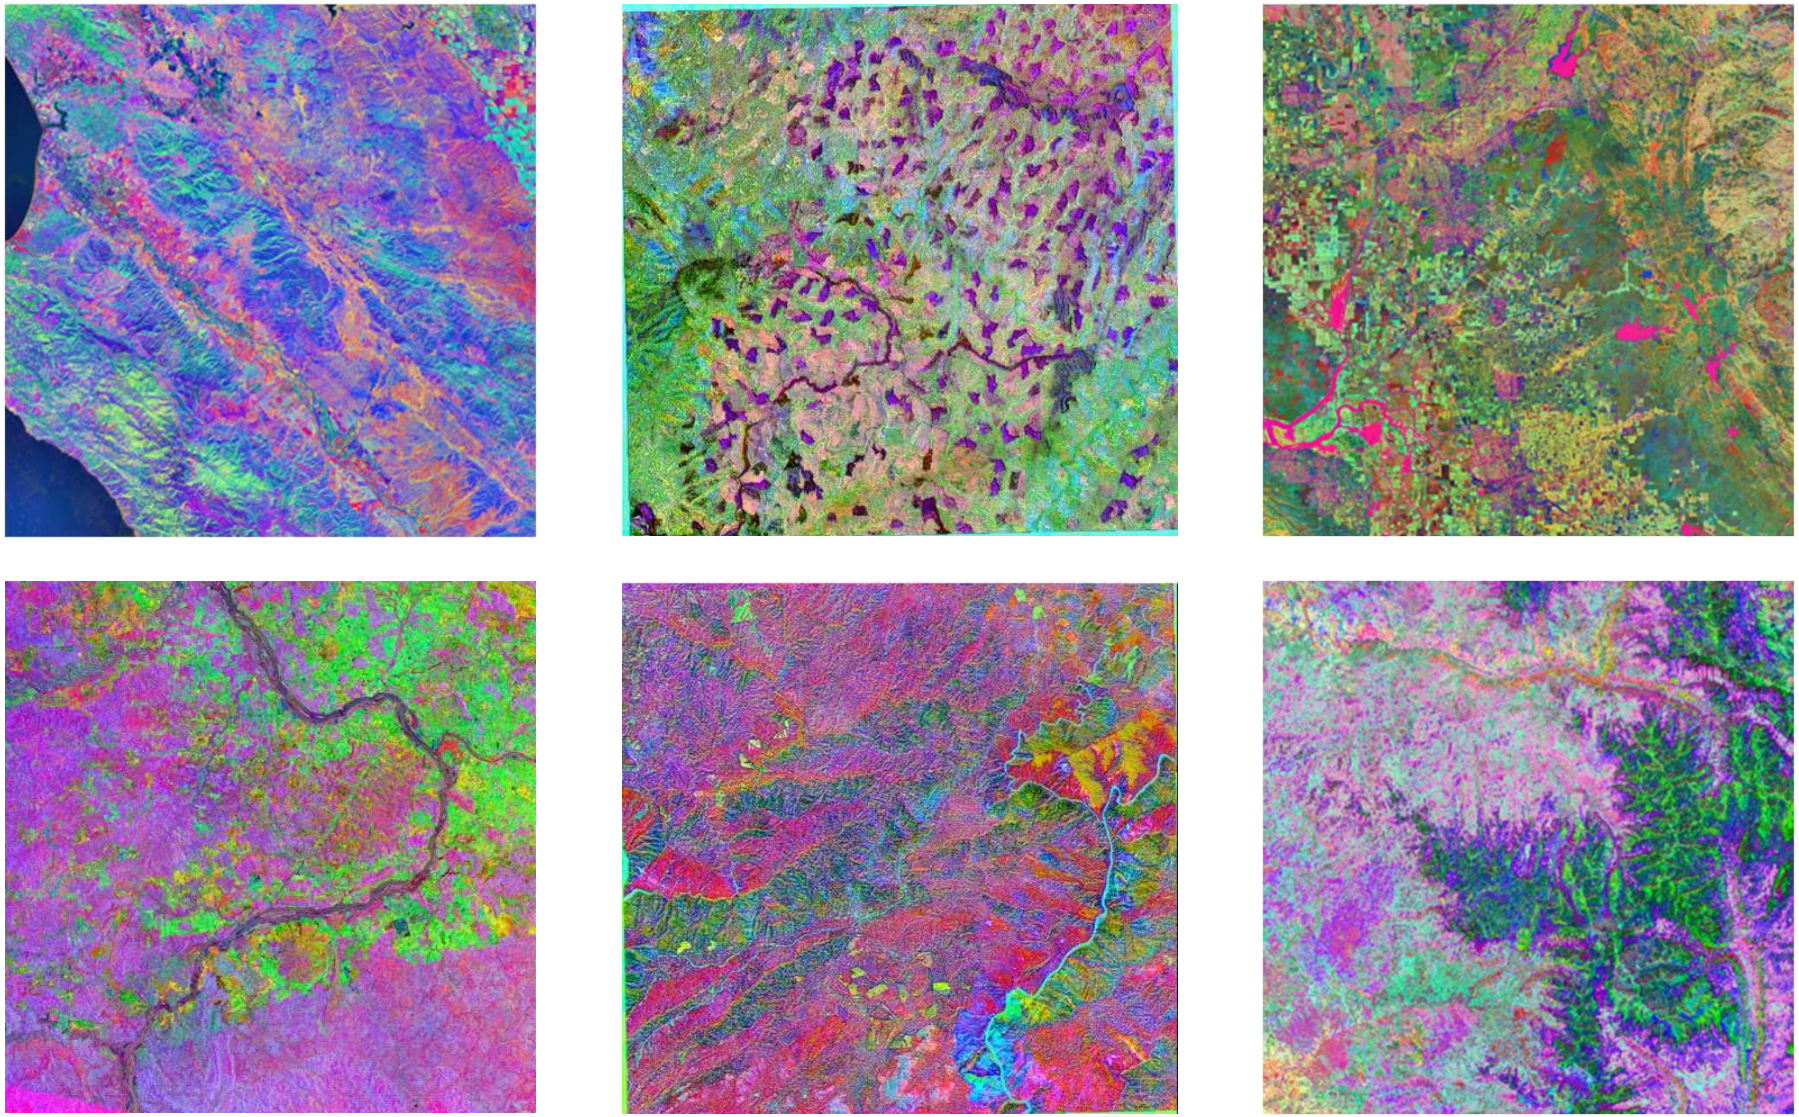

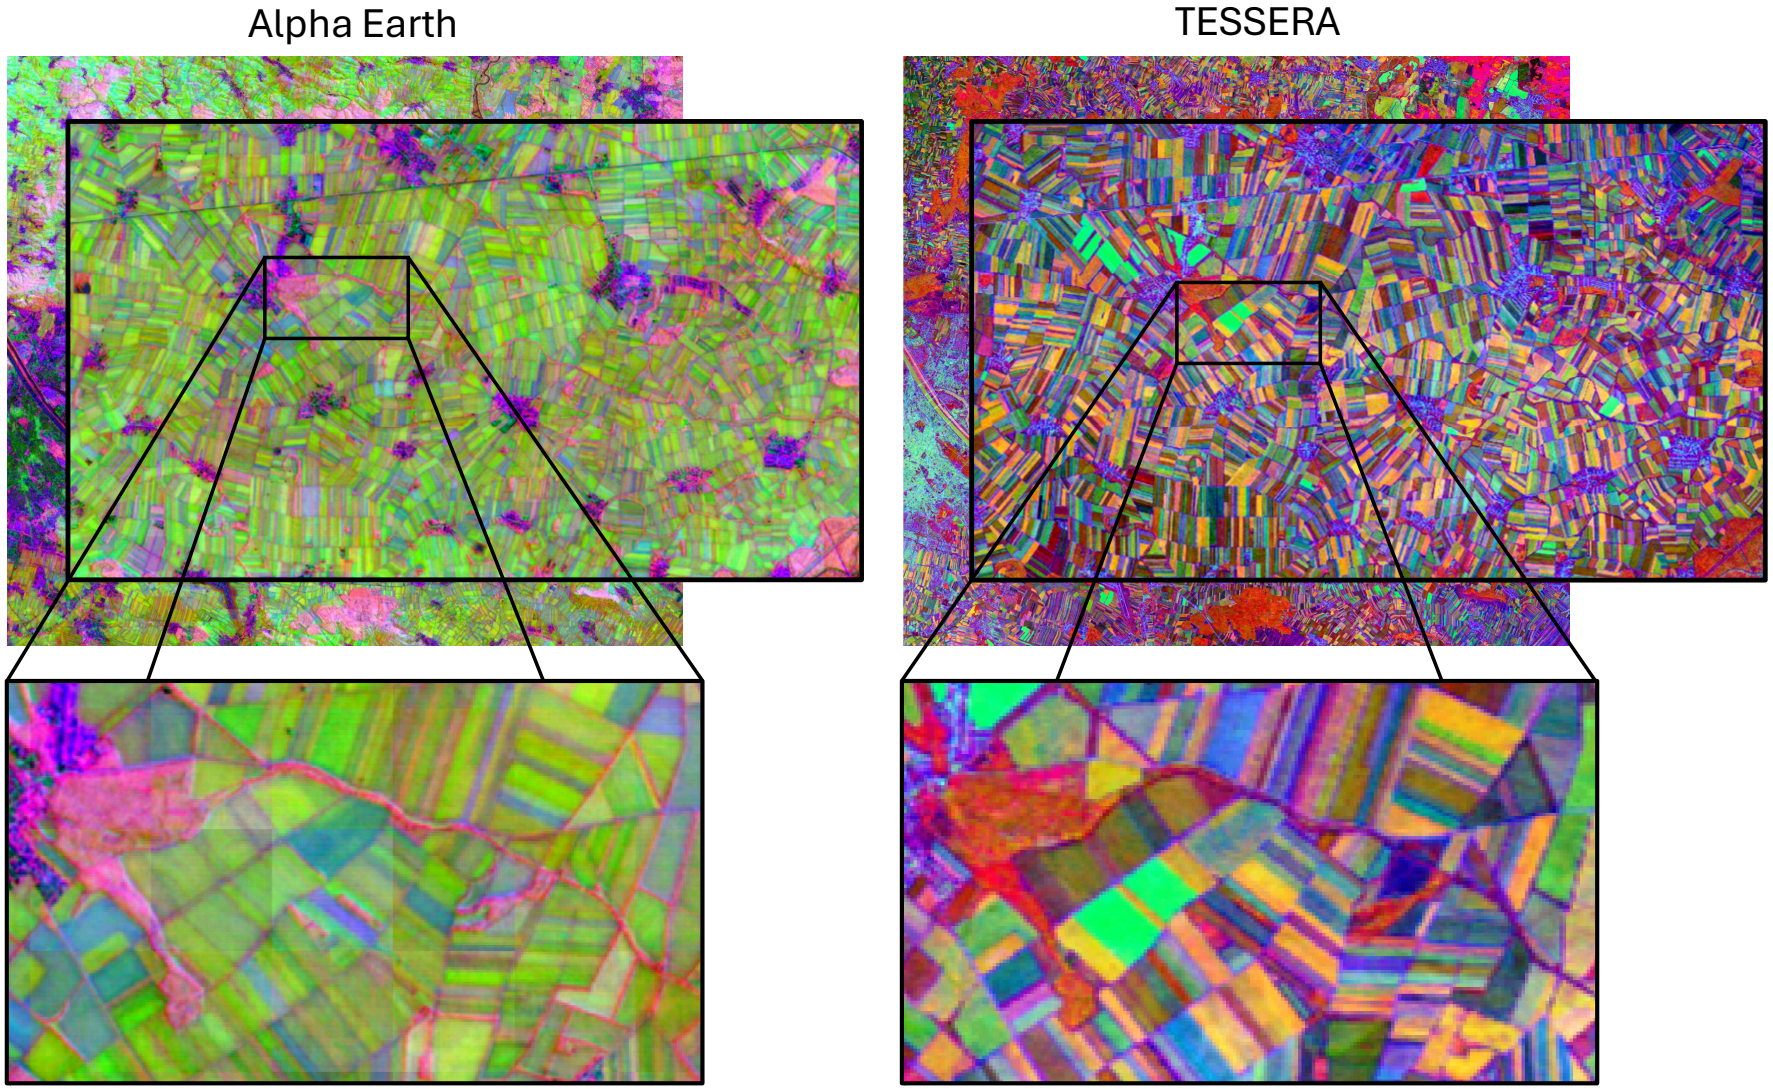
Alpha Earth
TESSERA

TESSERA vs AlphaEarth

	TESSERA	AlphaEarth
Training Data	S1, S2	S1, S2, Landsat, text, …
Coverage	Global 2017-2024*	Global 2017-2024
Embedding Dimension	128-d int8	64-d float64
Time Granularity	Annual	Annual
Resolution	10m	10m
Model Input	Pixel	Patch
Open-source	Yes	Partial
License	MIT	?

How well does it work?

Crop Classification

INVEKOS Austrian Crop Dataset

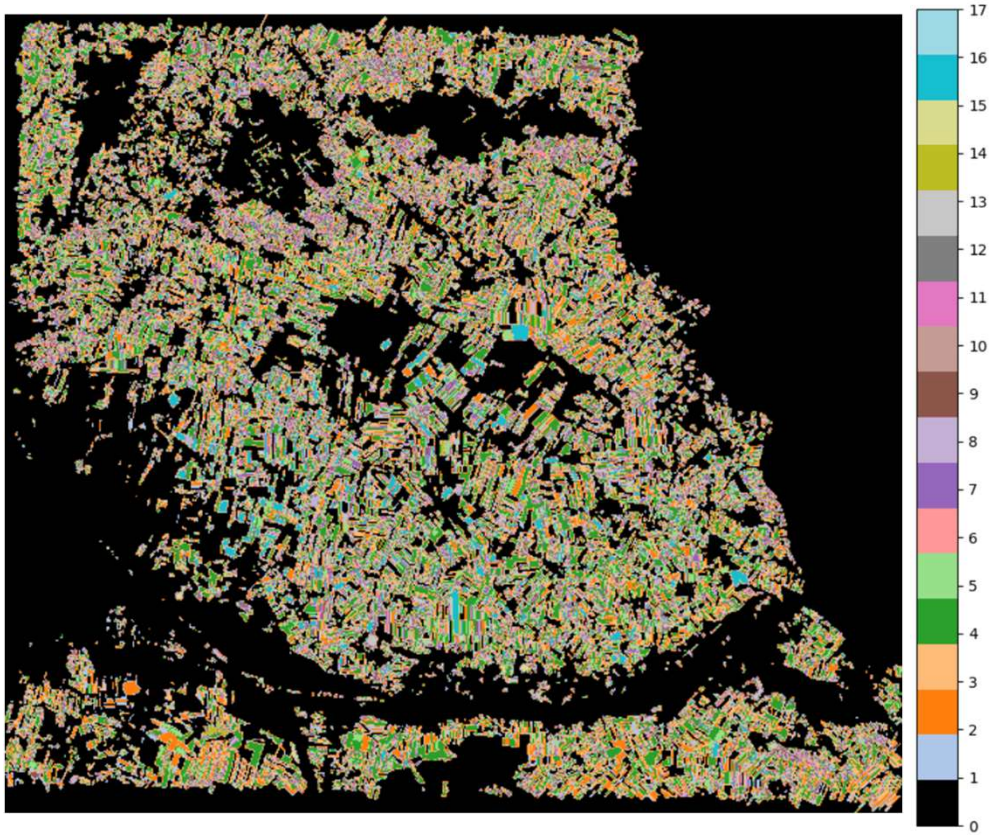

Name	Number of Samples	Percentage
Legume	2227	3.31%
Soy	5892	8.76%
Summer Grain	2475	3.68%
Winter Grain	24914	37.05%
Corn	6902	10.27%
Sunflower	207	0.31%
Mustard	1734	2.58%
Potato	2514	3.74%
Beet	1257	1.87%
Squash	2019	3.00%
Grapes	222	0.33%
Tree Fruit	347	0.52%
Cover Crop	1418	2.11%
Grass	2349	3.49%
Fallow	4484	6.67%
Other (Plants)	8220	12.23%
Other (Non Plants)	57	0.08%
Total	67238	100%

Performance

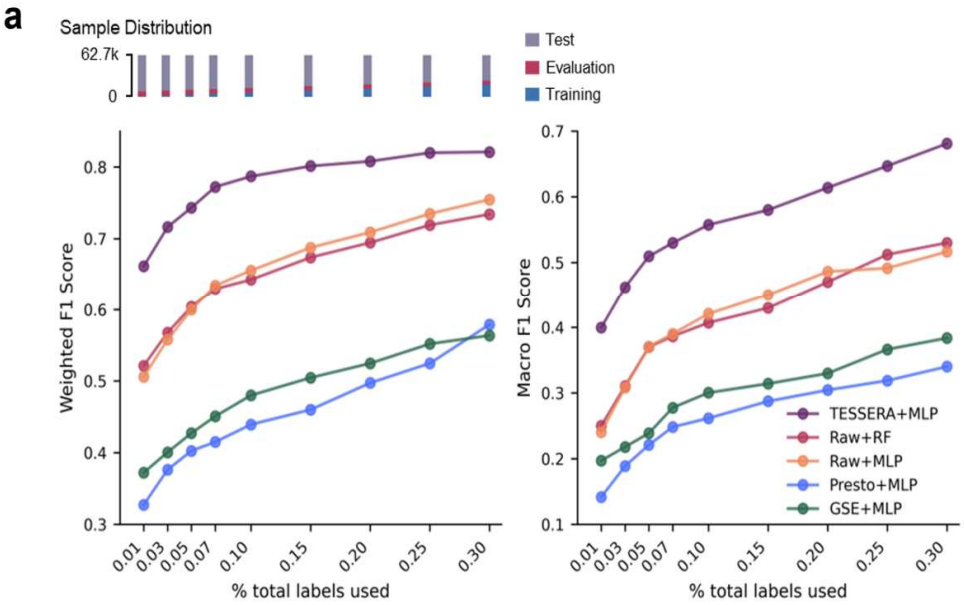

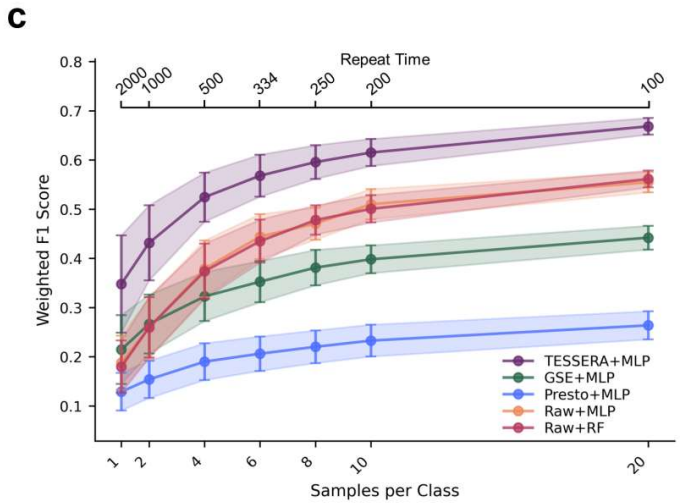

GSE is AlphaEarth

Presto is another GeoFM

Predicting Forest Height

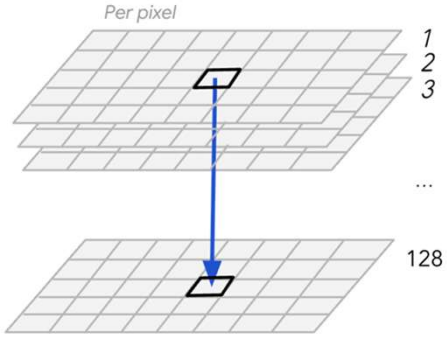

U-Net

Canopy Height from Airborne LiDAR

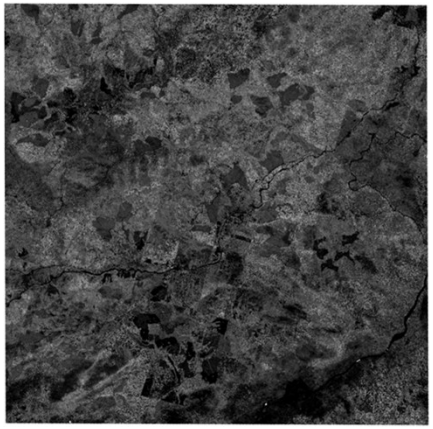

Danum Valley, Borneo
5 × 6 km

TESSERA Representations

Ground Truth

Prediction

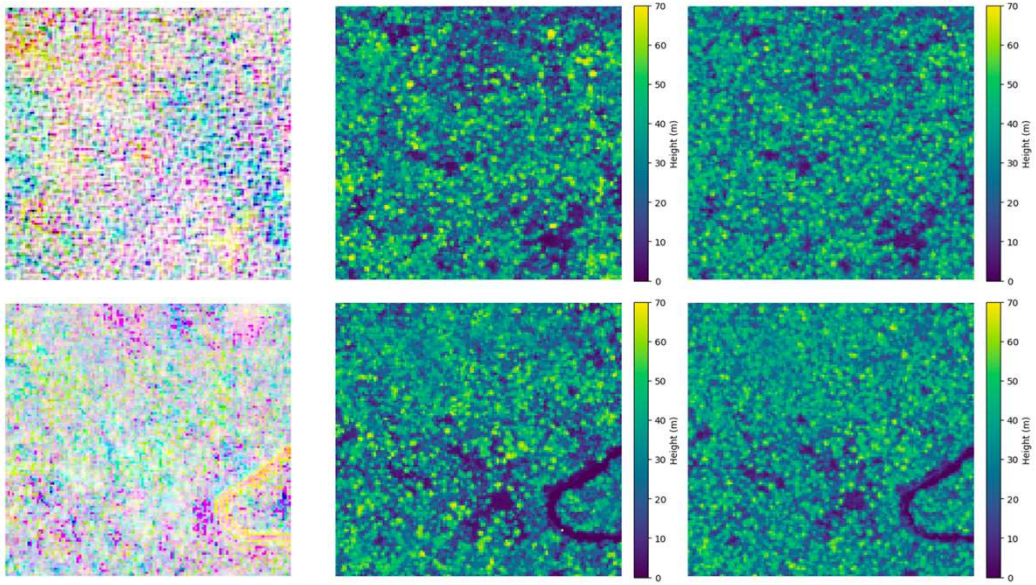

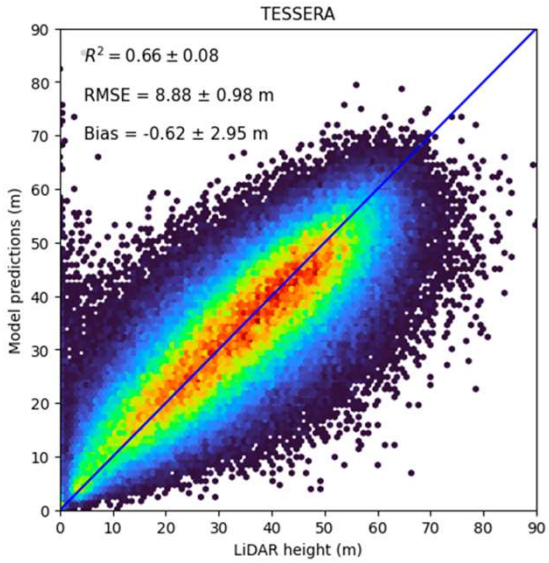

Predicting Forest Height

AlphaEarth

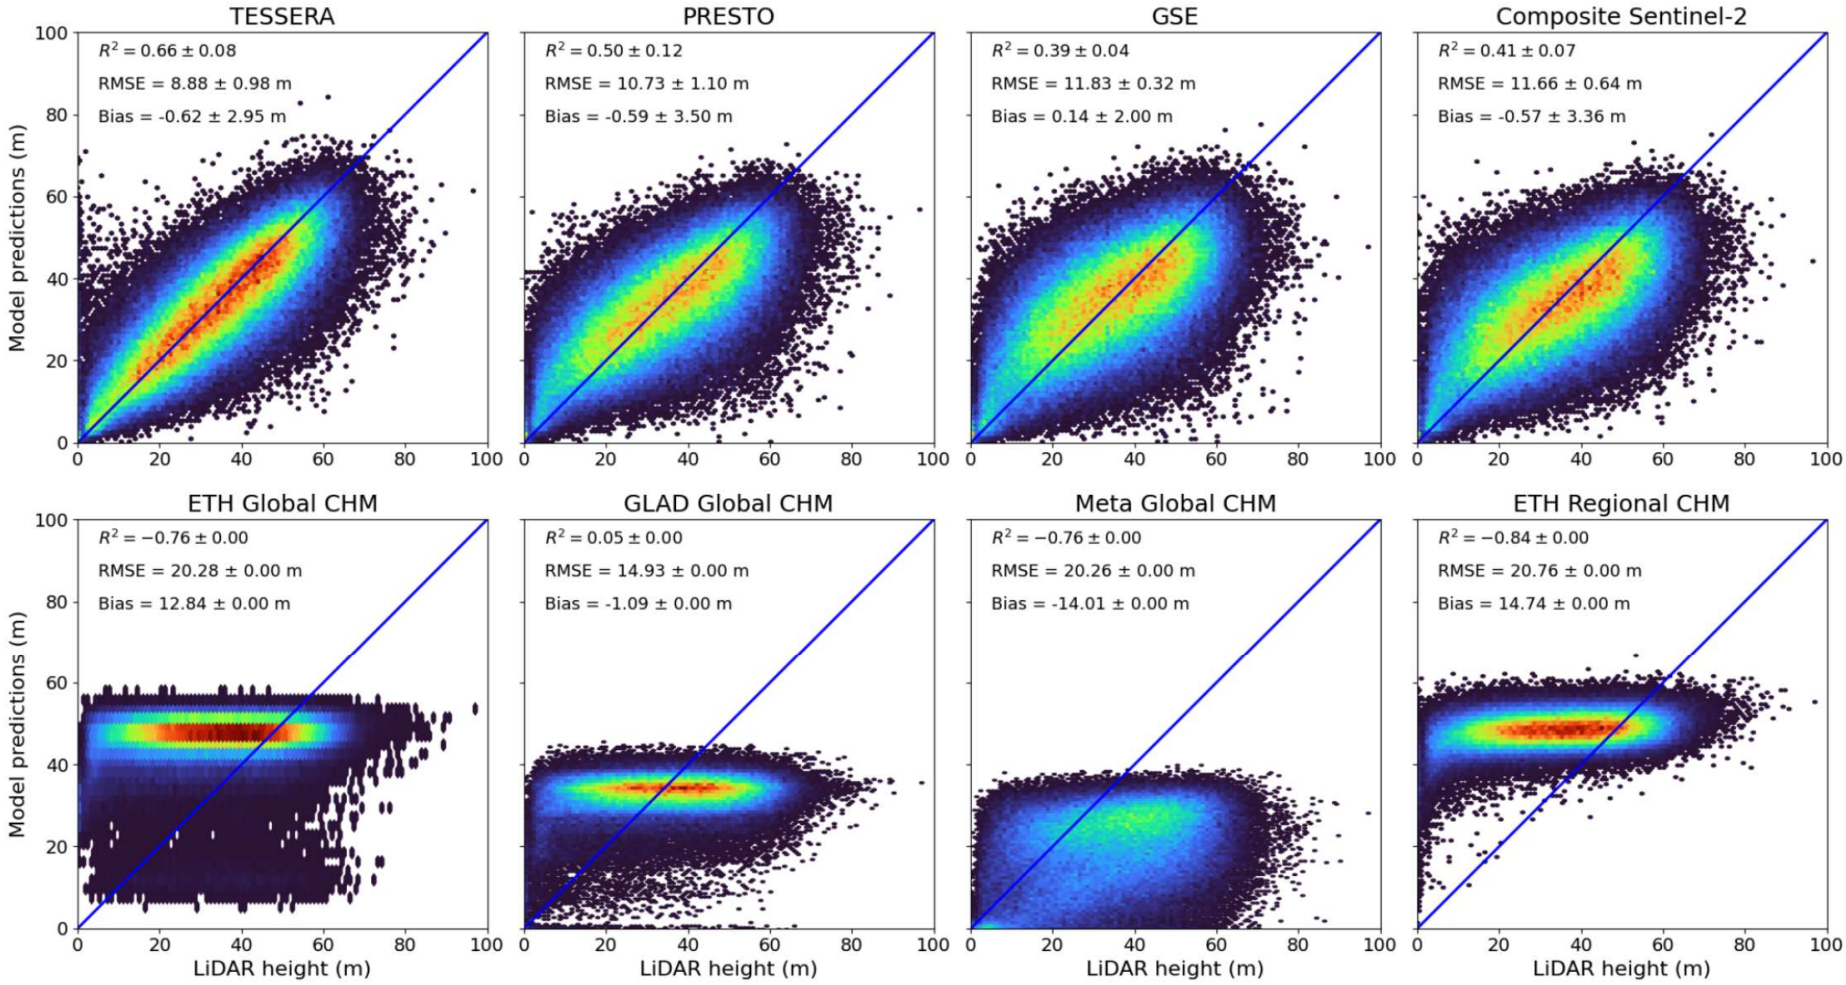

Detecting Disturbances: Wildfires in California

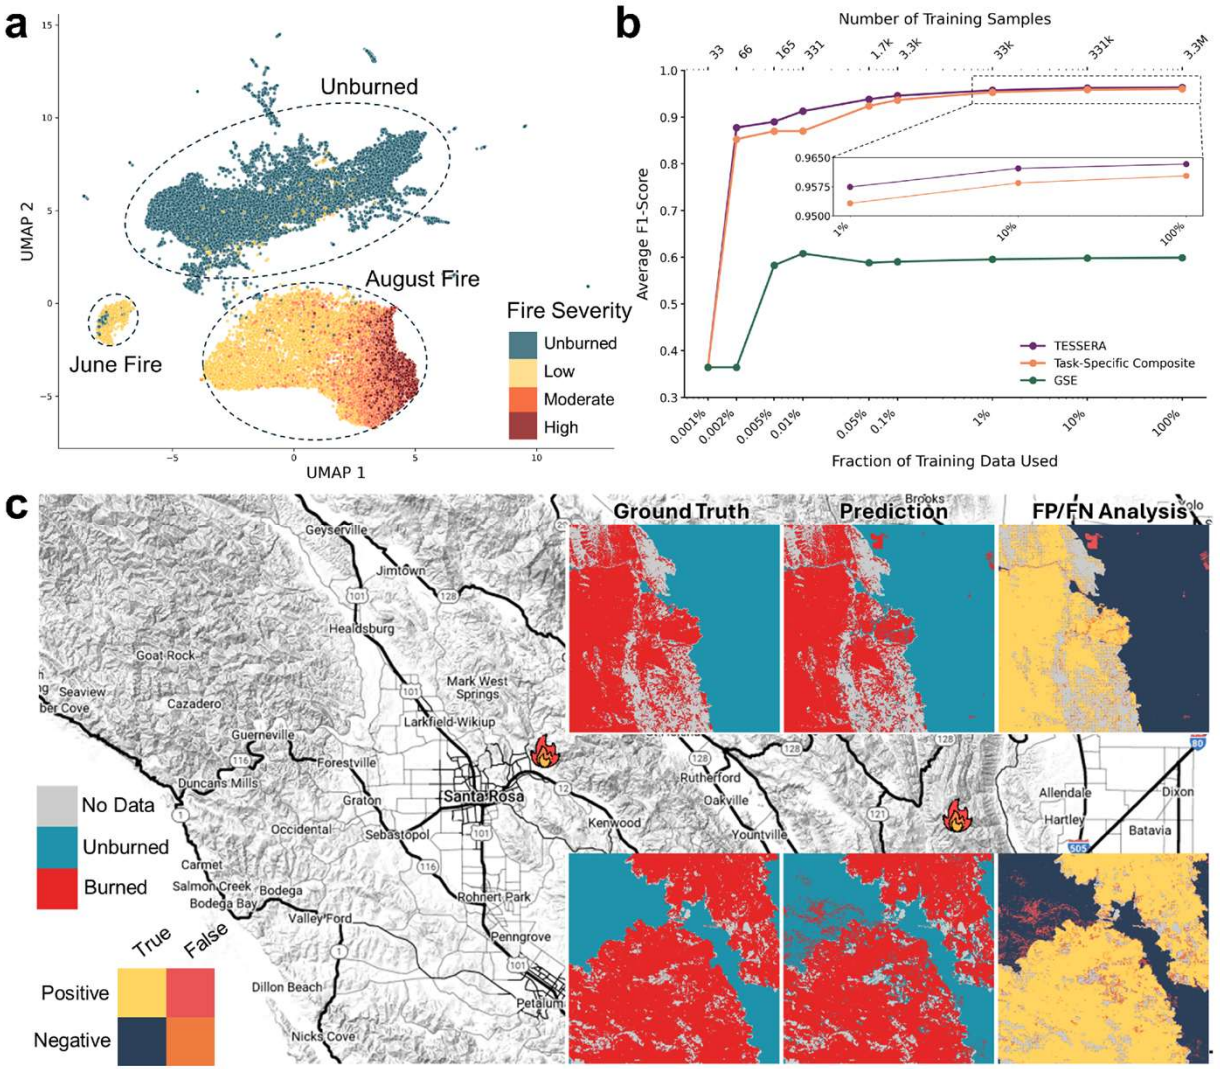

AGB in Finland (BIOMASSTRS)

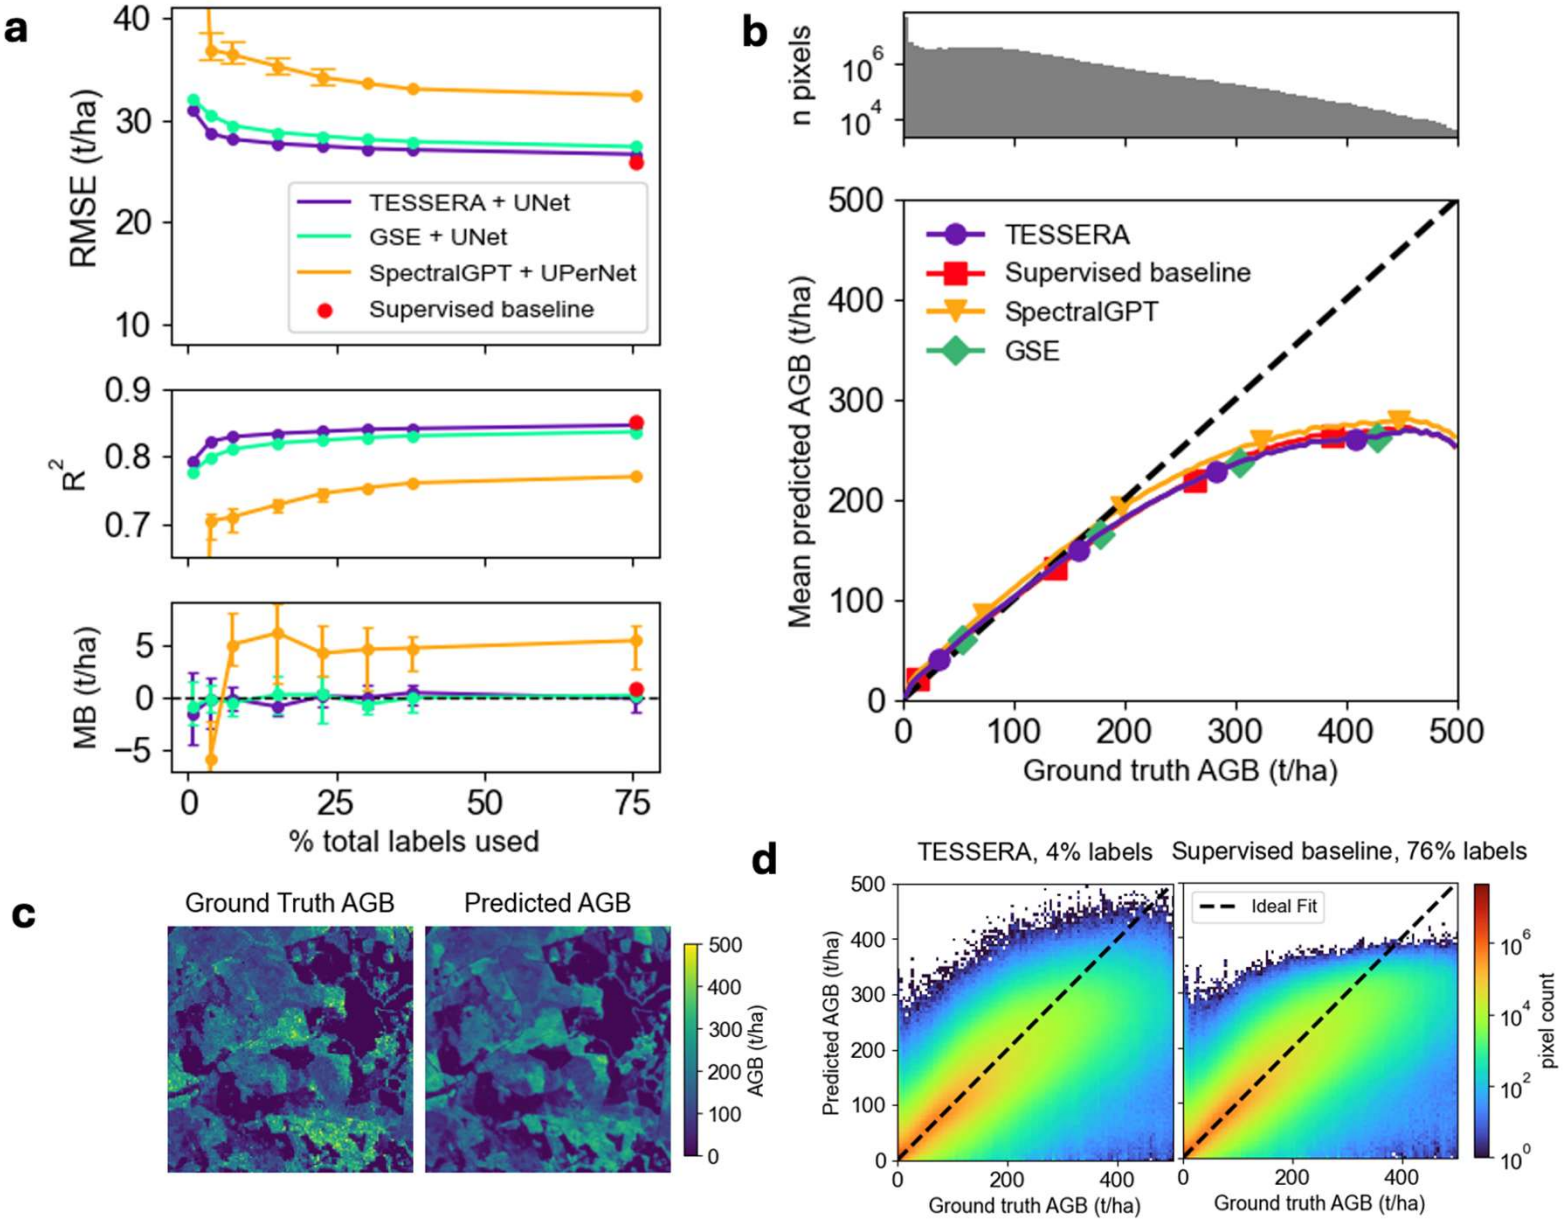

Land Use Change Detection

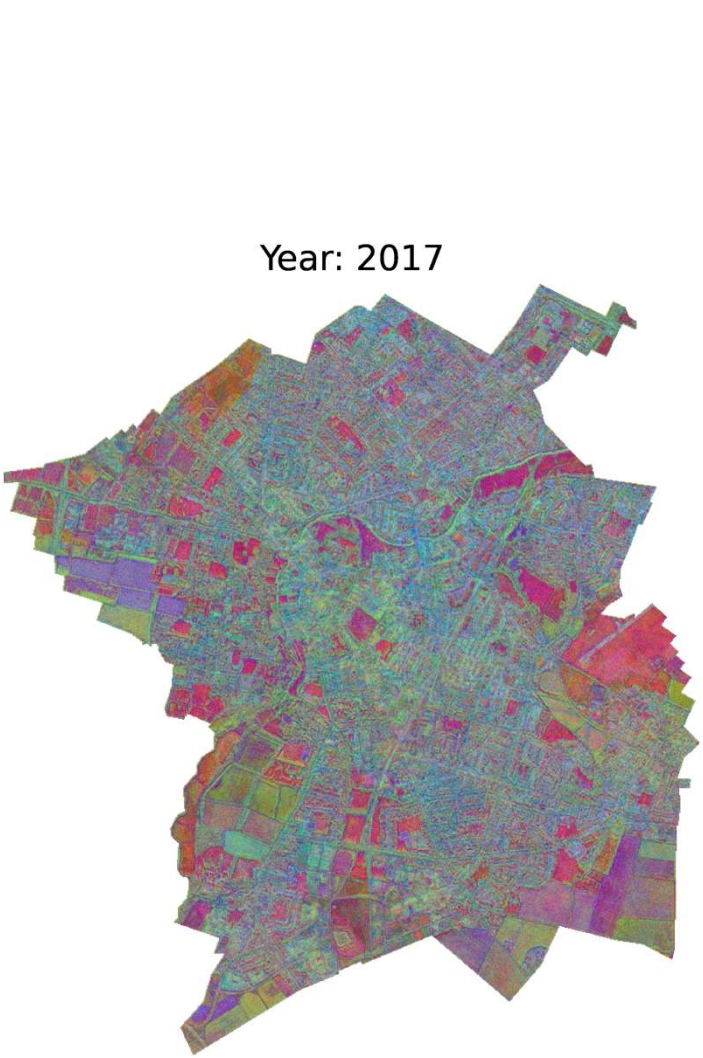

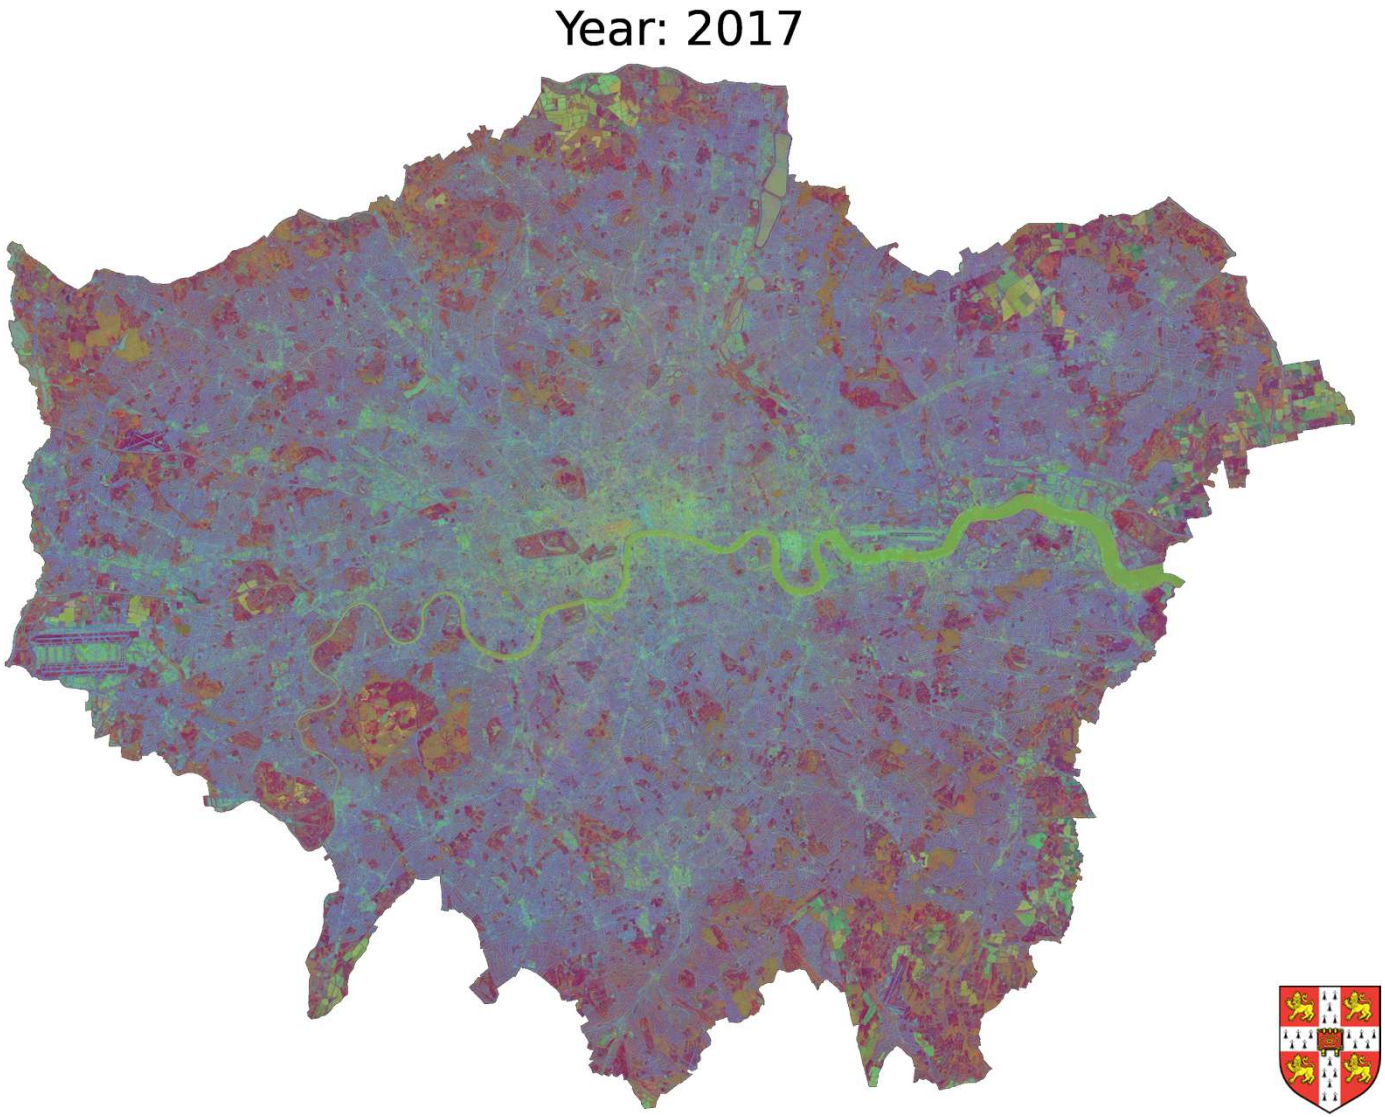

Land Use Change Detection

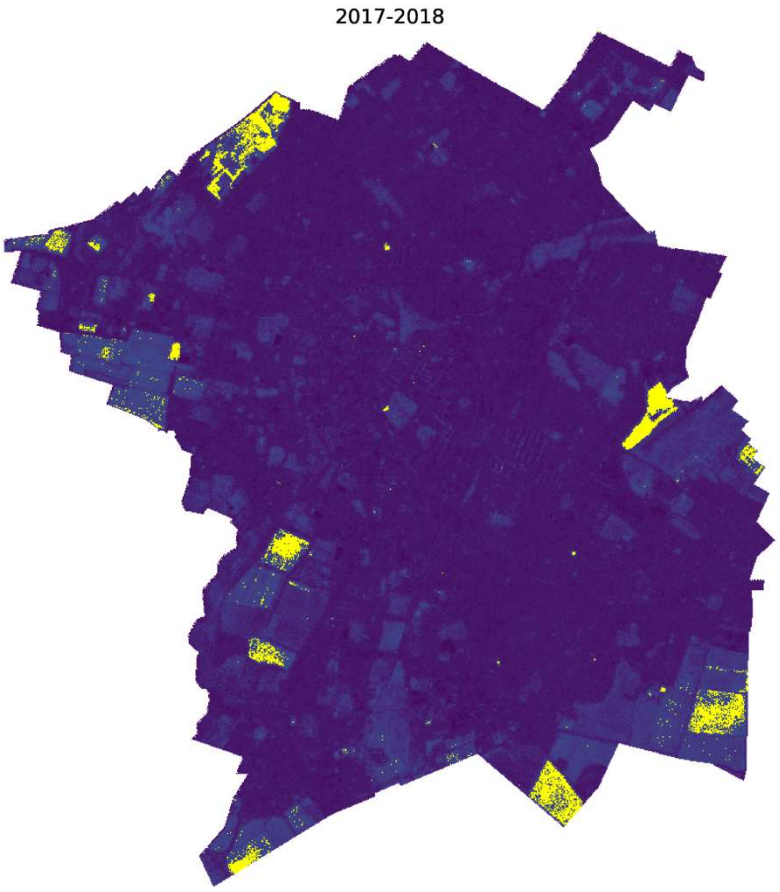

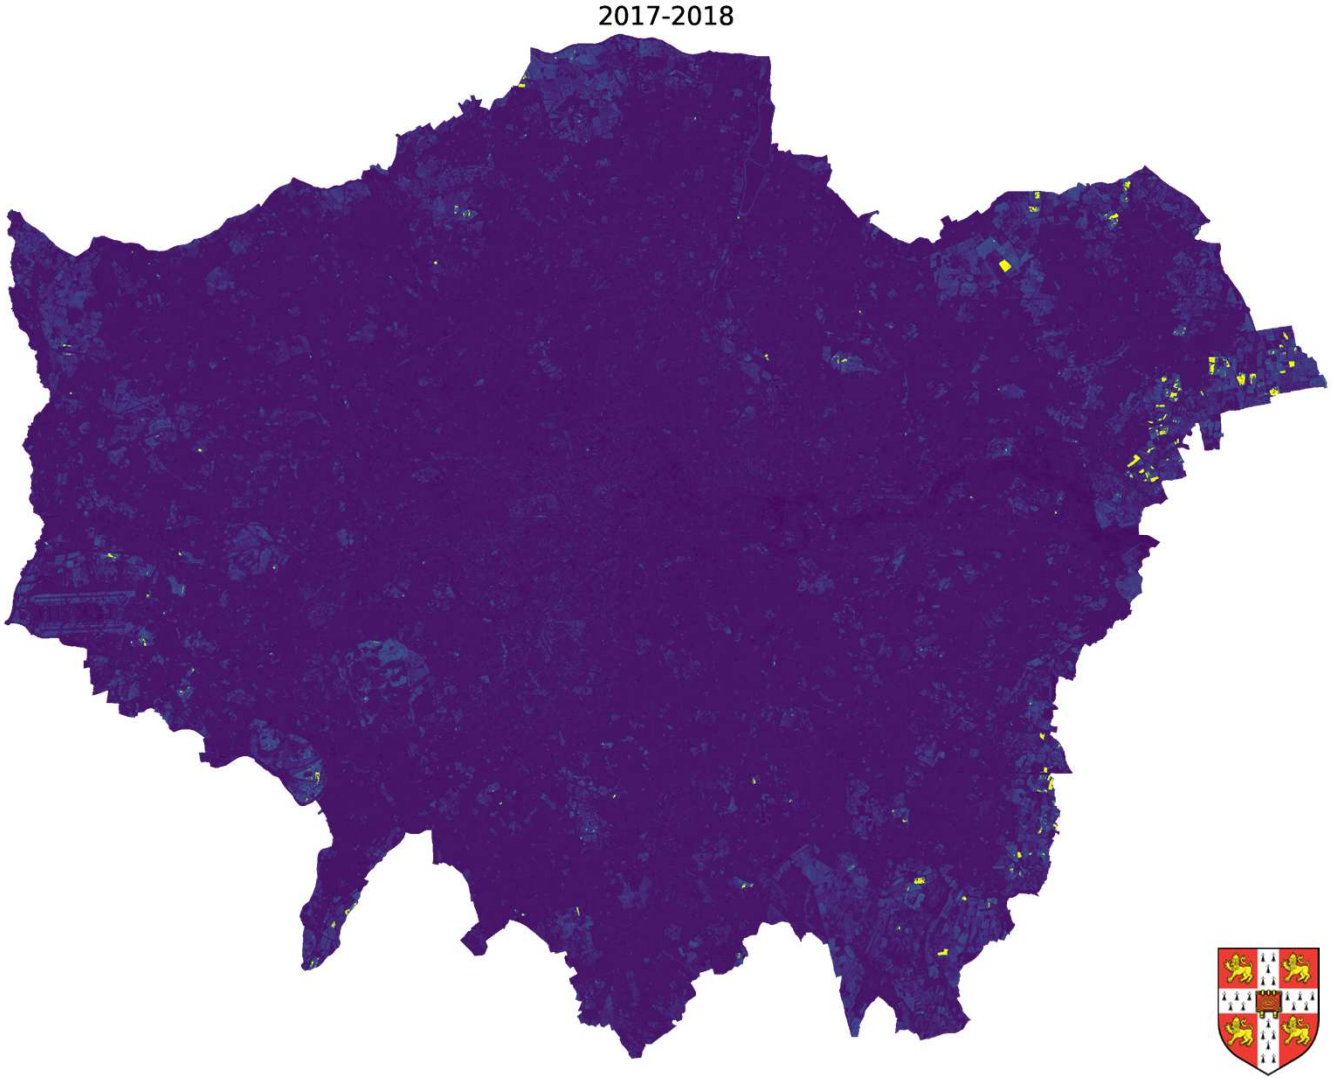

Some Systems Aspects

Data intensive task

- each tile is 150 GB
- we need to stream and process ~24PB
- 1.6M GPU hours + 11M GPU hours!

Need both raw compute and algorithmic design

Compute

- Dawn HPC
- AMD 8-GPU node
- 256-core 1TB RAM server
- ~300TB SSD

We are still CPU- and GPU- hungary!

TESSERA Tool Kit

Tessera

Temporal Embeddings of Surface Spectra for Earth Representation and Analysis (TESSERA)

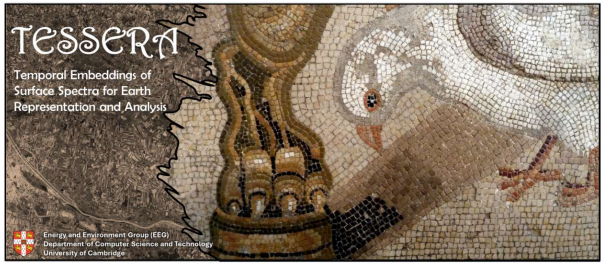

View Our Paper | Report Bug | Request Feature | Podcast

PyPI version v0.6.0 License MIT

Table of Contents

Fully open-sourced repo
- Train the model
- Evaluate the model
- Create embeddings **yourself**

Geotessera

GeoTessera

Python library for accessing and working with Tessera geospatial foundation model embeddings.

Overview

GeoTessera provides access to geospatial embeddings from the Tessera foundation model, which processes Sentinel-1 and Sentinel-2 satellite imagery to generate 128-channel representation maps at 10m resolution. These embeddings compress a full year of temporal-spectral features into dense representations optimized for downstream geospatial analysis tasks. Read more details about the model.

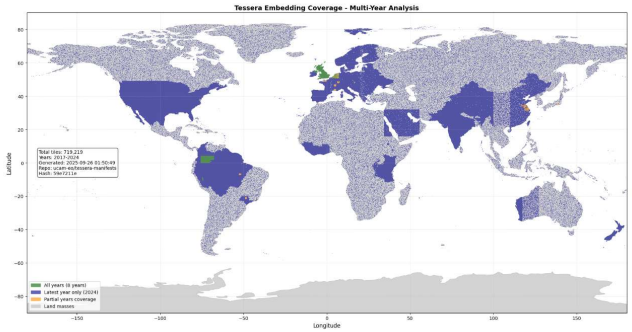

User Request (lat/lon bbox)
↓
Registry Lookup (find available tiles)
↓
Download Files (via Pooch with caching)
├— embedding.npy (quantized)
└— embedding_scales.npy
↓
Dequantization (multiply arrays)
↓
Output Format
├— NumPy arrays → Direct analysis
└— GeoTIFF → GIS integration

Interactive Map

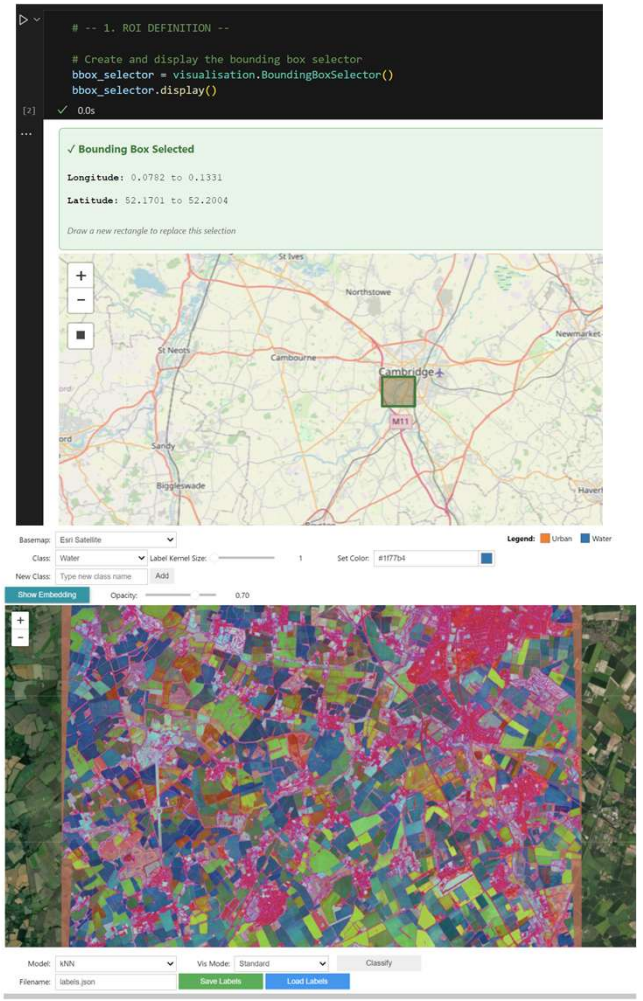

Track Farm PV Panels in the UK

2h to track the farm solar panels in the UK from 2017-2024

1. Go check the data availability map in Geotessera repo

1min

2. Download all the embeddings in the UK from 2017-2024

60min

3. Create some PV panel labels! (I only did this for 2024)

15min

4. Ask Claude Code to give me a pipeline for pixel-wise PV panel detection!

10min

5. Train the XGBoost model

1min

6. Use the model to detect PV panels in the UK from 2017-2024

40min

1+60+15+10+1+40
= 118 min

Does the result look good?

Global dataset from Li etc.

TESSERA

Note: The above two outputs are from 2020, but we only use 2024 embeddings/labels for training!

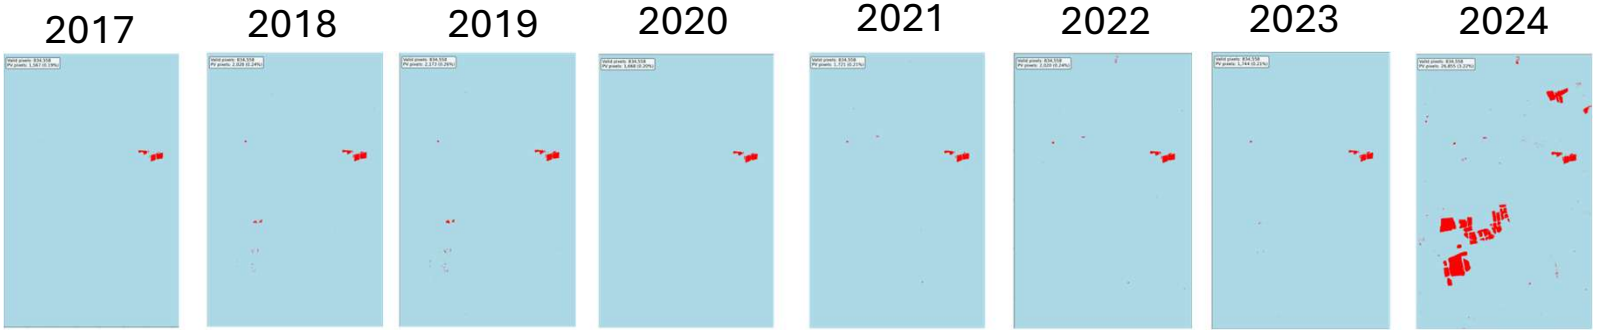

Interdisciplinary team

Computer Science

- AI
- Systems
- Computer vision

Frank Feng
Jovana Knezevic
Robin Young
Sadiq Jaffer
Andrew Blake
S. Keshav
Anil Madhavapeddy

Remote sensing

- Satellite data
- Radiative transfer models

Maddy Lisaius
Frank Feng
Clement Atzberger
Markus Immitzer

Agriculture/Ecology

- Forest processes
- Crop behaviour

Maddy Lisaius
Clement Atberger
David Coomes
James Ball

TESSERA

We are looking for collaborators!

Read our paper on arXiv
(https://arxiv.org/abs/2506.20380)

Our Python Library
(https://github.com/ucam-eo)