TESSERA

Zhengpeng(Frank) Feng

May 20, 2025

with C. Atzberger, M. Lisaius, S. Keshav,, A. Blake, D.A. Coomes, S. Jaffer, J. Knezevic, A. Madhavapeddy, R. Young

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

We'd like to

Monitor
- Land use
- Biodiversity

Assess status of
- Crops
- Soils
- Forests

Quantify
- Forest degradation
- Deforestation
- Carbon stock

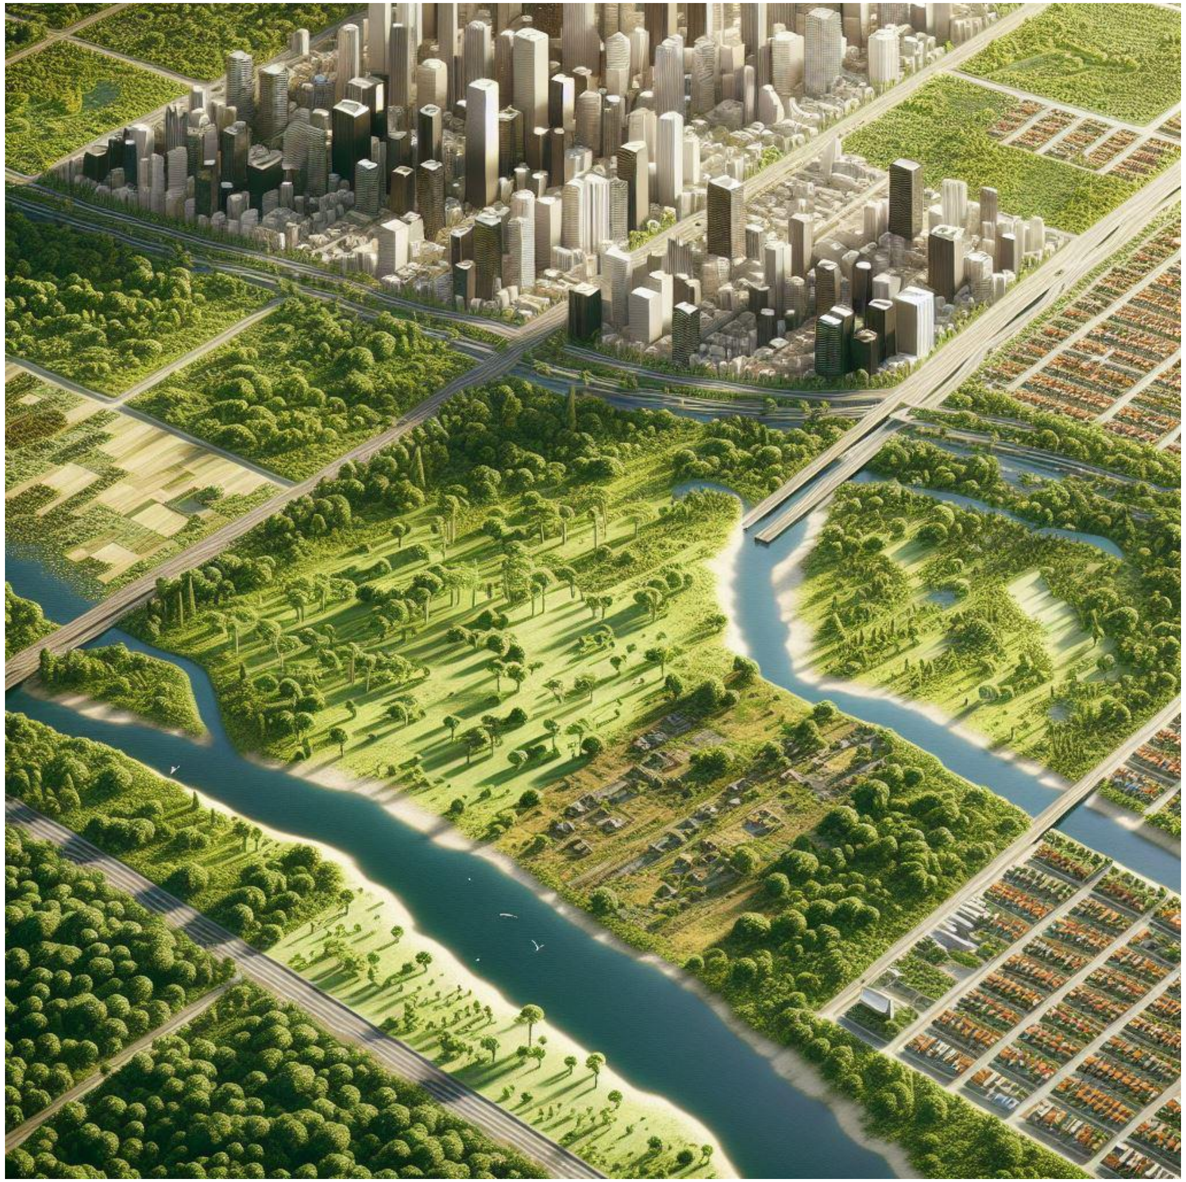

zf281@cam.ac.uk

Lots of public satellite data!

10
million
Landsat
TM / ETM+ / OLI

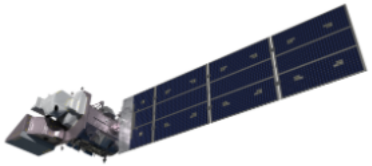

31
million
Sentinel-2
L1C

2
million
Sentinel-1
GRD

Problems

- Clouds
  - >50% in an image often cloud-covered; many models only accept <10%
  - Clouds block surface info, adding noise
- Varying time gaps
  - Irregular intervals hurt temporal learning
  - Large gaps = abrupt, hard-to-learn changes
- Changes in lighting
  - Lighting shifts mimic real land changes
  - Models may misclassify due to brightness

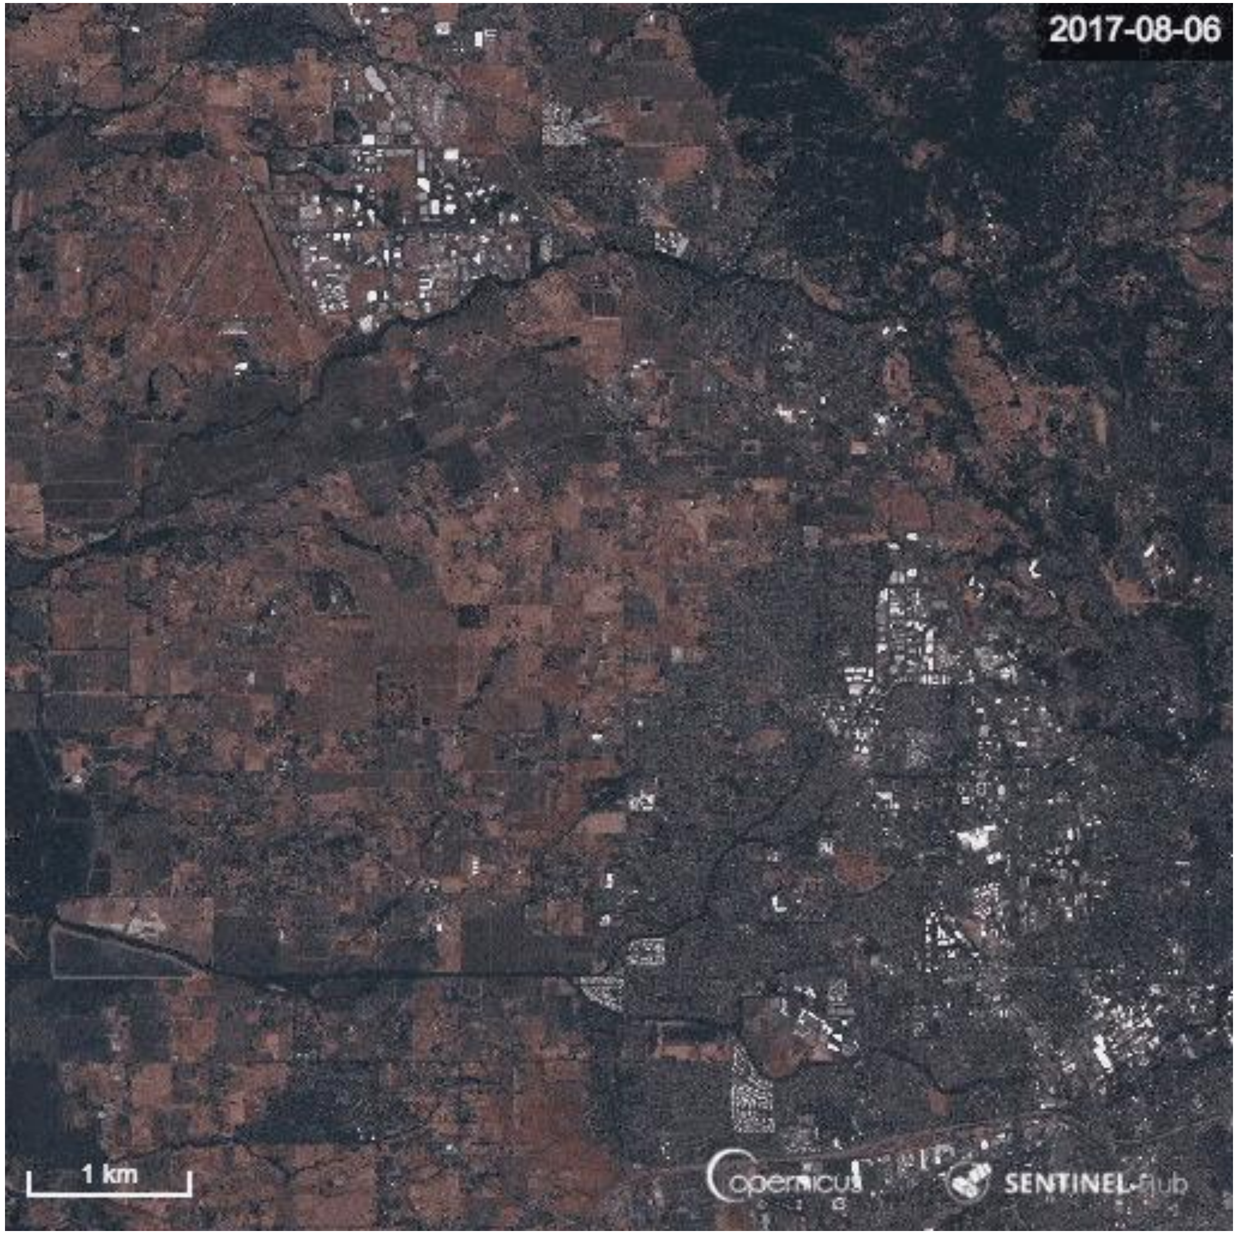

Current practice: Compositing

Multi-timestep Cloud-corrupted → **Single-/Few- timestep Cloud-free**

Cloud corrupted
Sentinel-2 MSI

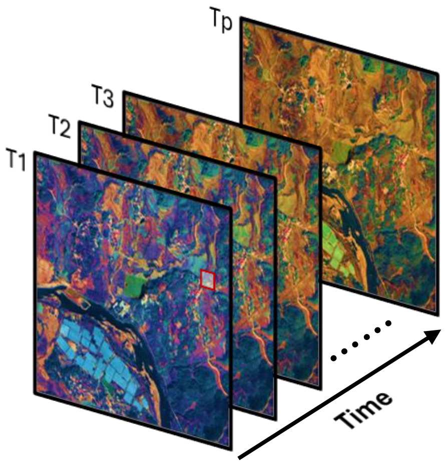

https://www.esa.int/ESA_Multimedia/Images/2018/05/Cloud-free_Europe_from_Sentinel-3

zf281@cam.ac.uk

Problem with Compositing

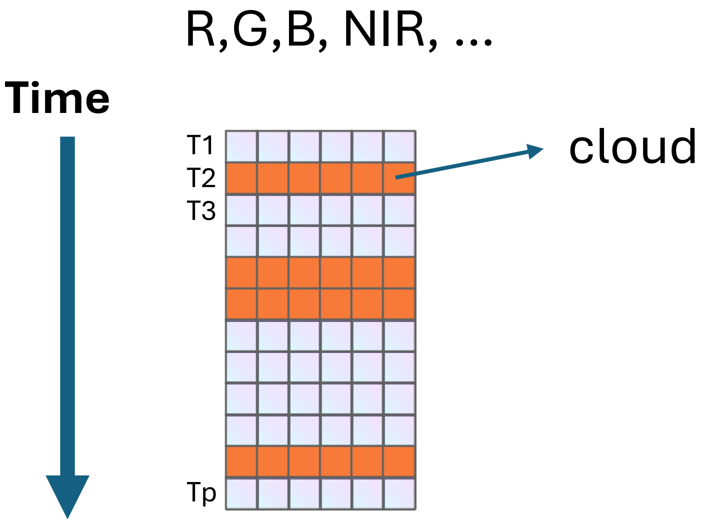

Compositing hides the temporal signal!

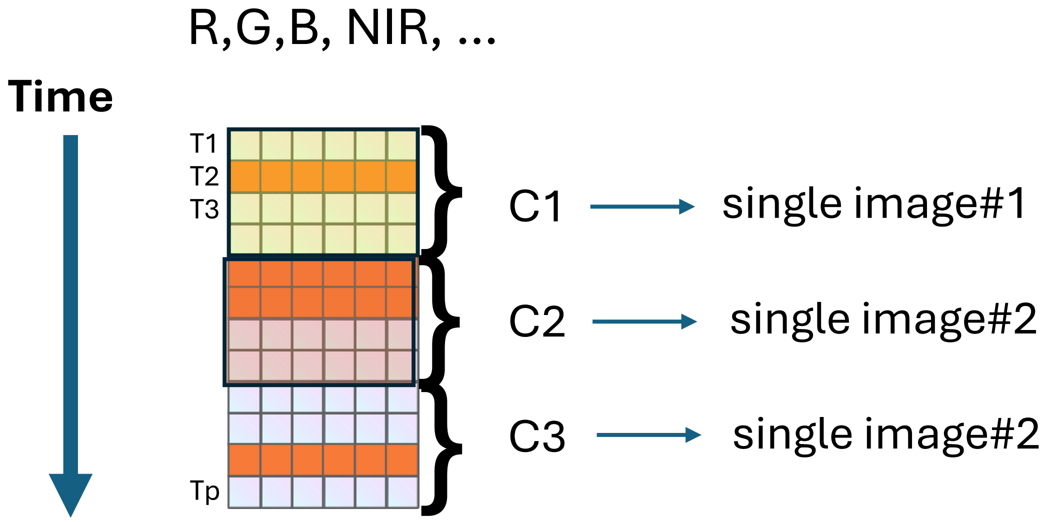

zf281@cam.ac.uk

We use AI to produce a
128-dimension cloud-free "image"
encoding the temporal evolution of each spectral band

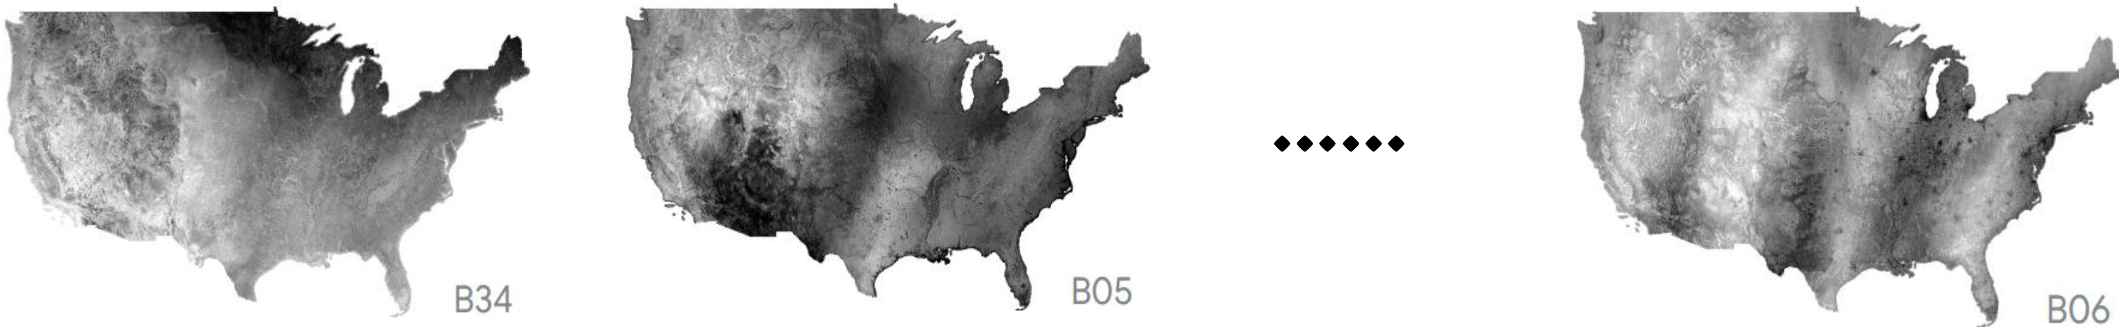

Image from Google

How?

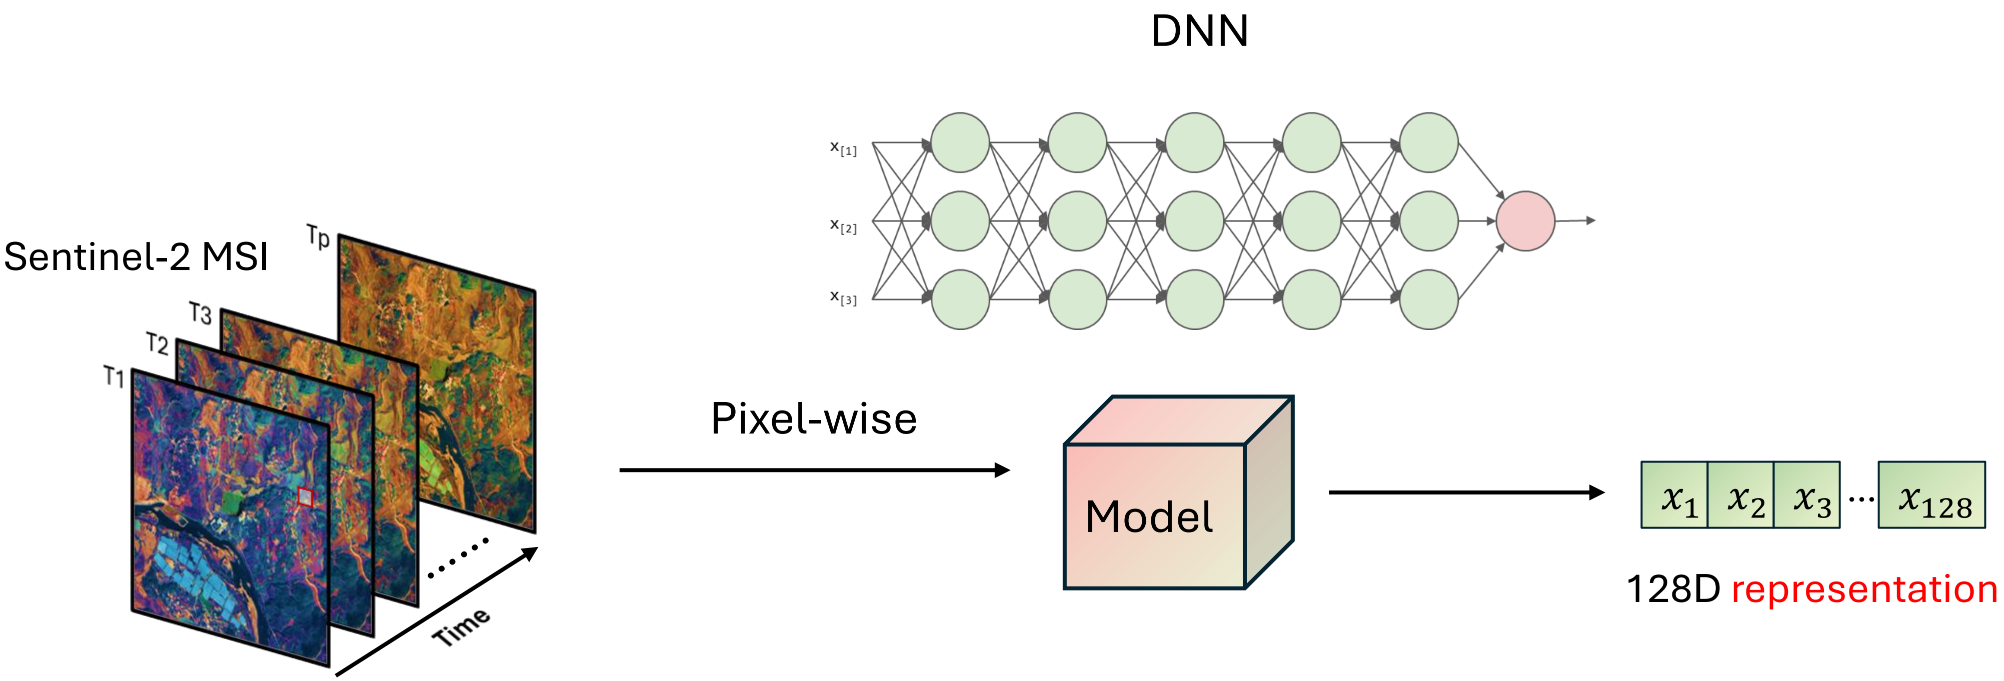

zf281@cam.ac.uk

	Supervised Learning	Self-supervised Learning (SSL)
Data	(x,y) where x is data, y is label	Only x!
Goal	Learn a function to map x -> y	Learn some underlying hidden structure of the data (**Representation**)
Tasks	• Land Classification • AGB Regression • Object Detection	Land Classification / AGB Estimation / Object Detection
Performance	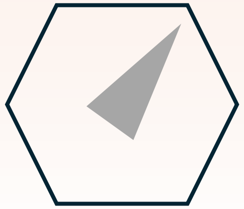	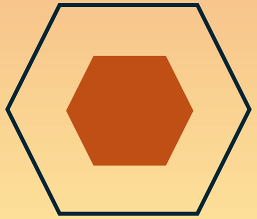 Finetune 

Land Classification

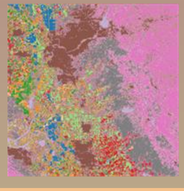

AGB Estimation

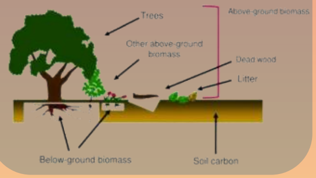

Object Detection

zf281@cam.ac.uk

SSL example

Masked Auto Encoding (MAE)

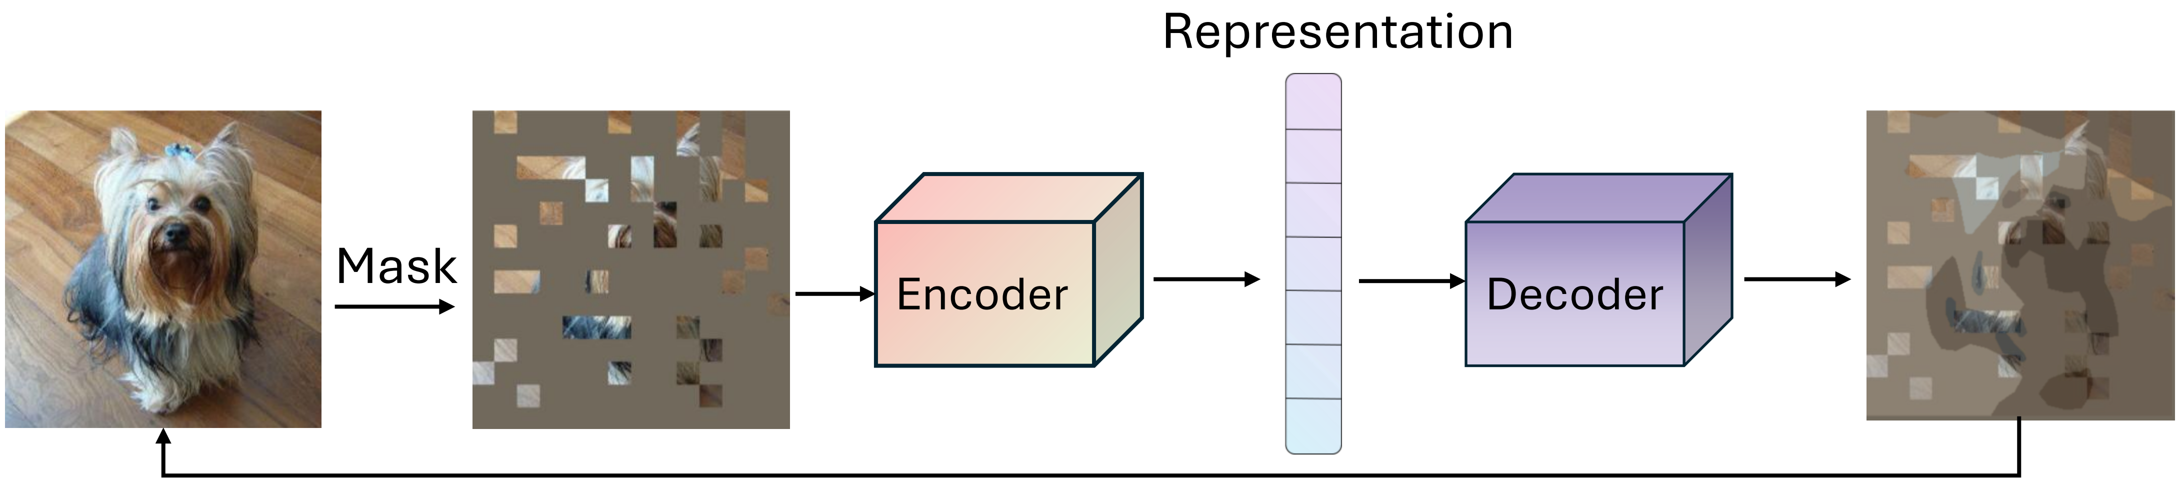

Are they similar?

zf281@cam.ac.uk

Barlow Twins SSL*

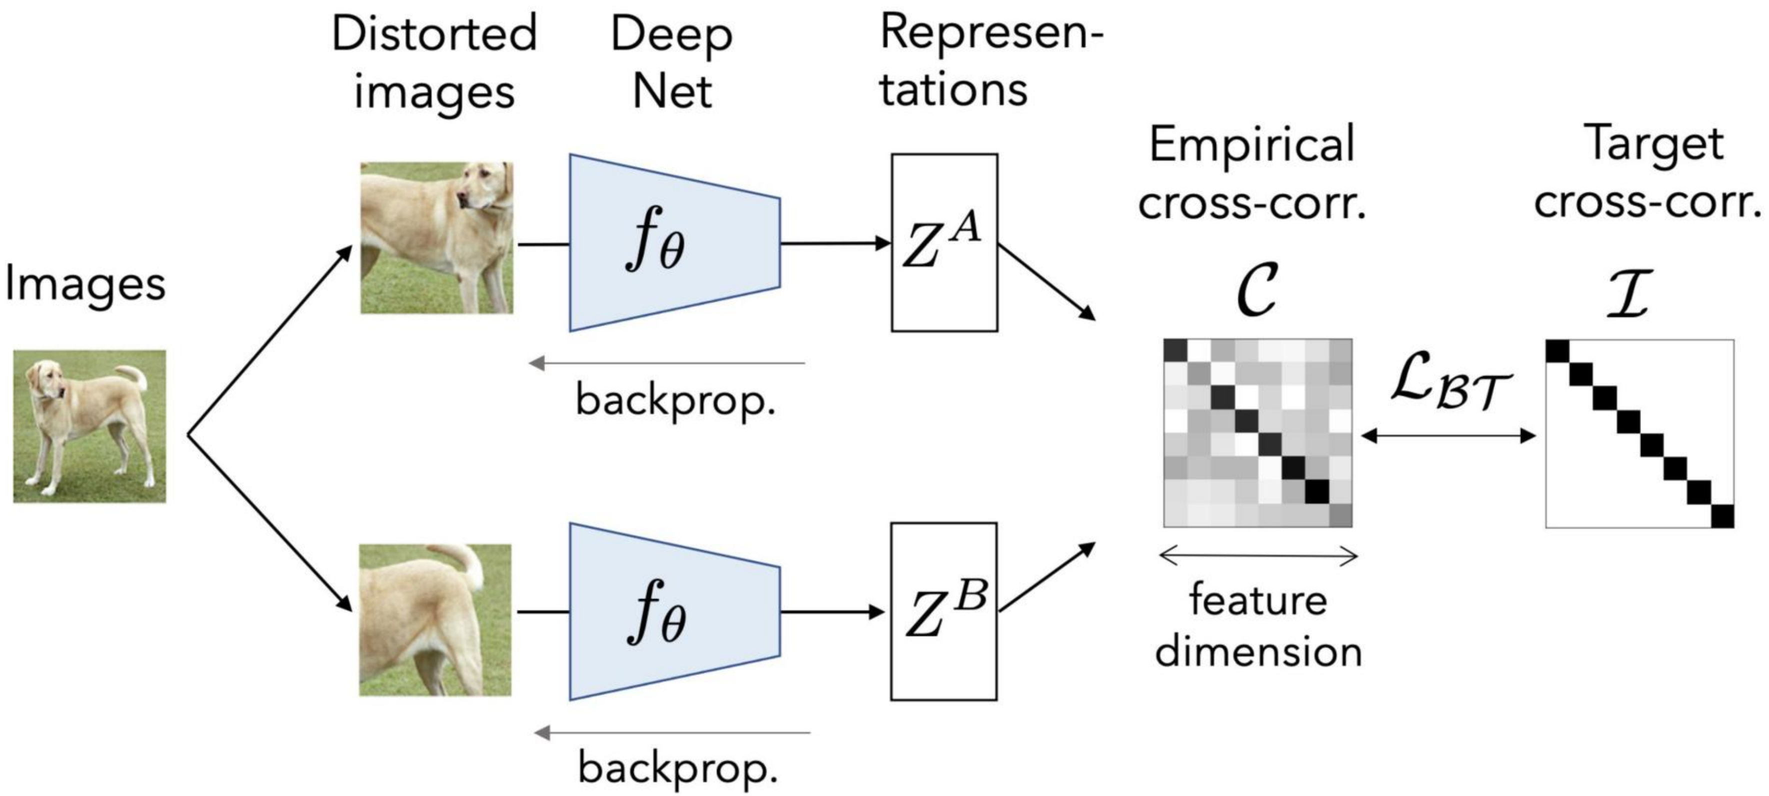

$$\mathcal{L}_{\mathcal{BT}} \triangleq \underbrace{\sum_i (1-\mathcal{C}_{ii})^2}_{\text{invariance term}} + \lambda \underbrace{\sum_i \sum_{j \neq i} {\mathcal{C}_{ij}}^2}_{\text{redundancy reduction term}}$$

$$\mathcal{C}_{ij} \triangleq \frac{\sum_b z^A_{b,i} z^B_{b,j}}{\sqrt{\sum_b \left(z^A_{b,i}\right)^2} \sqrt{\sum_b \left(z^B_{b,j}\right)^2}}$$

* simplified

Cross-correlation Matrix

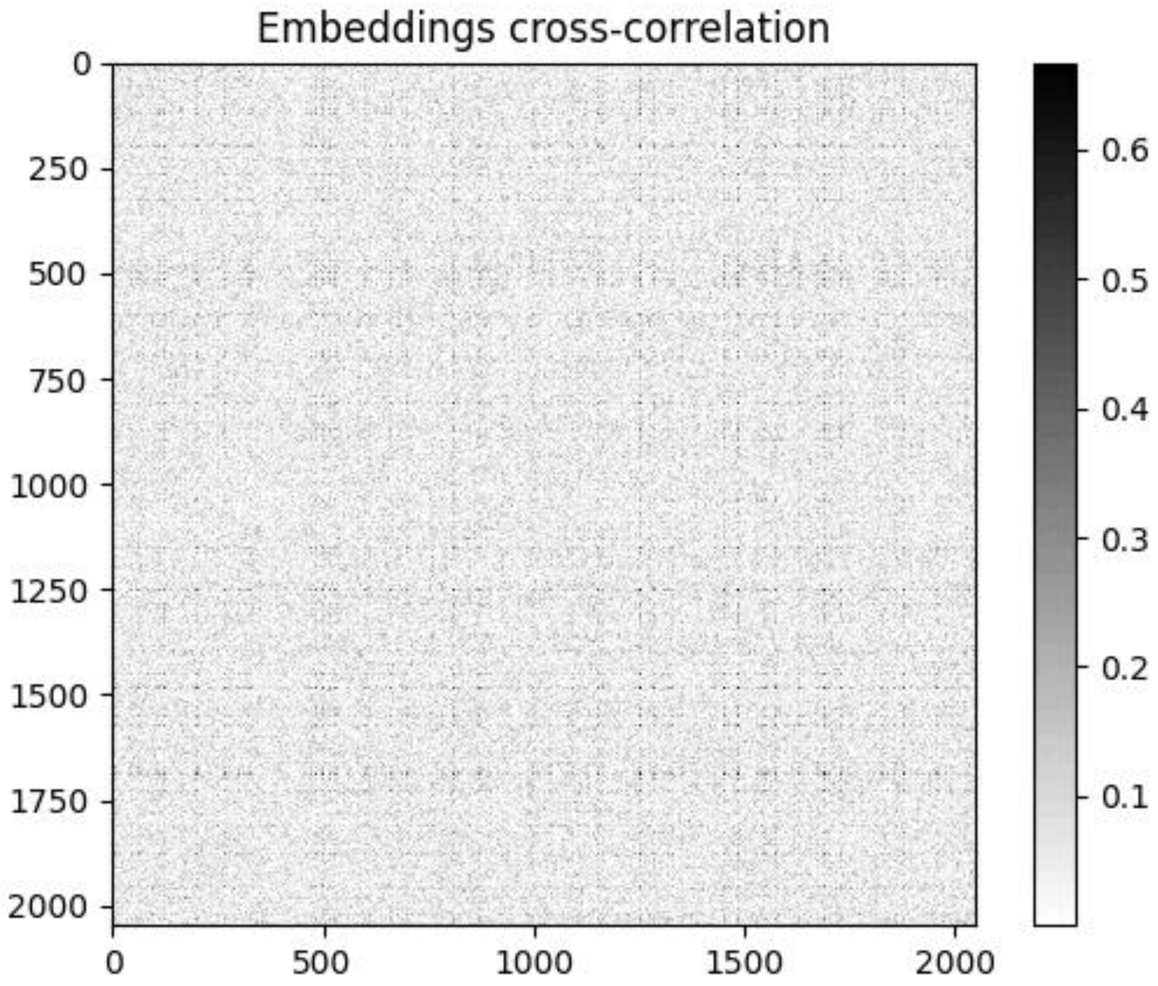

zf281@cam.ac.uk

Barlow Twins for Remote Sensing

Create two augmentations (sub-samples) from the same sample

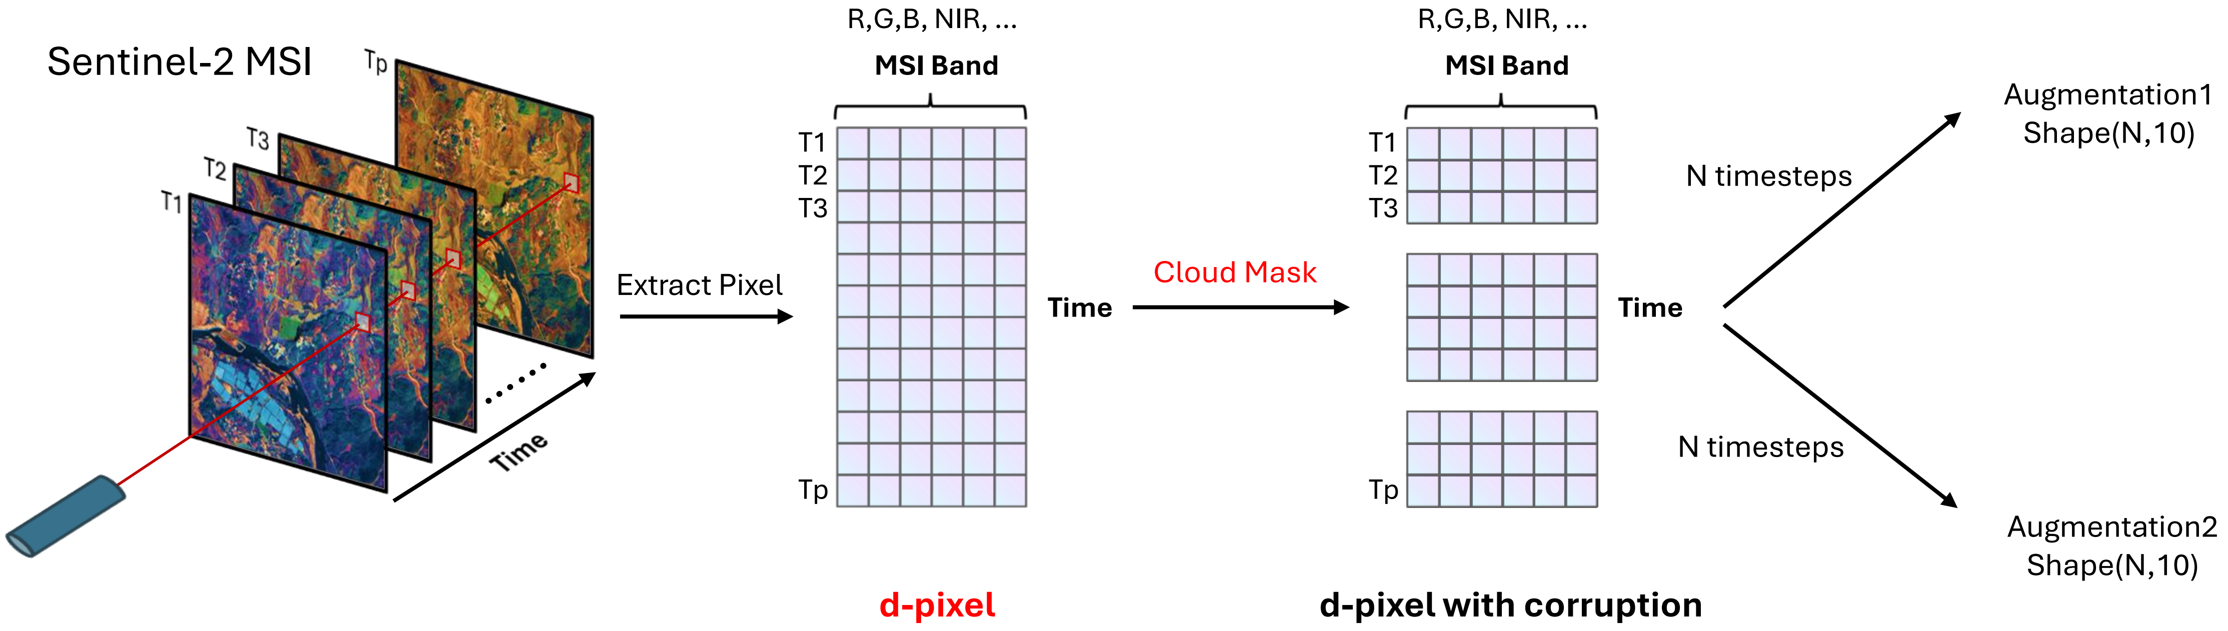

Lisaius M C, Blake A, Keshav S, et al. Using Barlow twins to create representations from cloud-corrupted remote sensing time series[J]. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 2024.

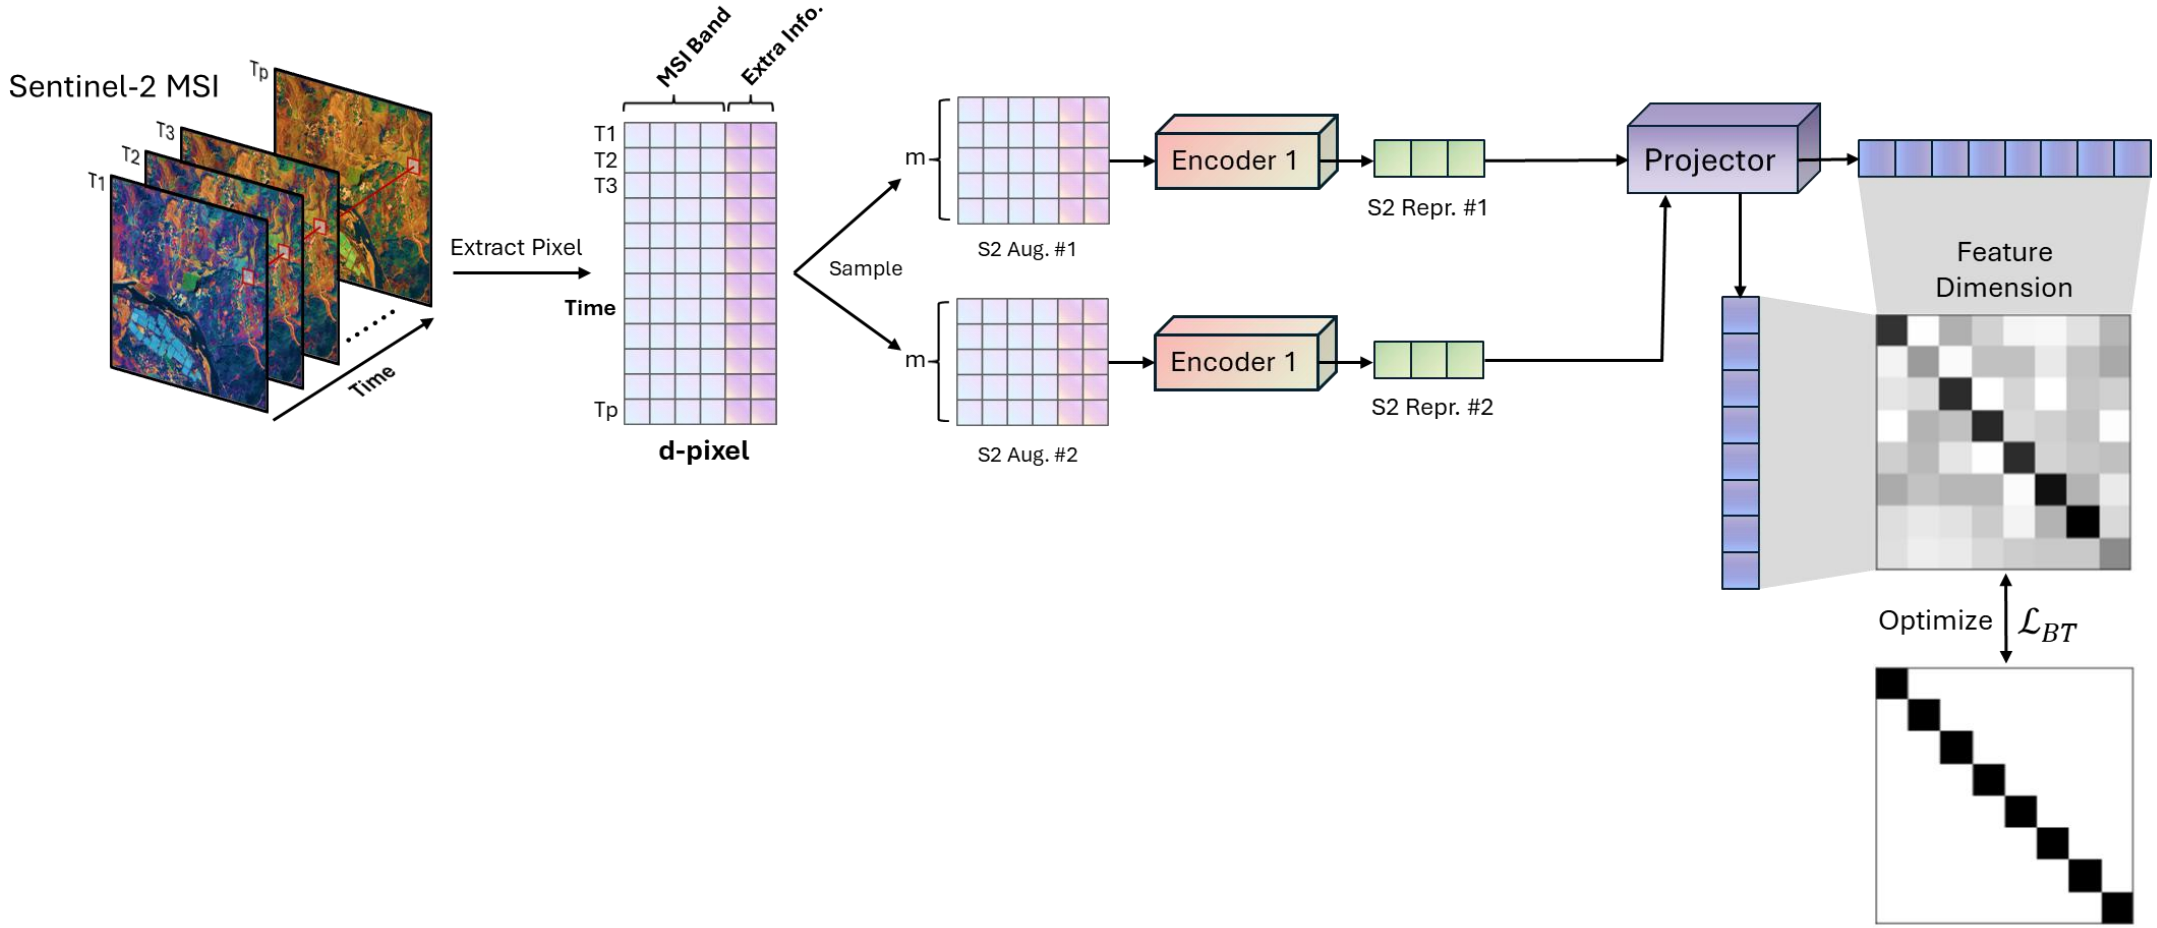

zf281@cam.ac.uk

Add Sentinel 1

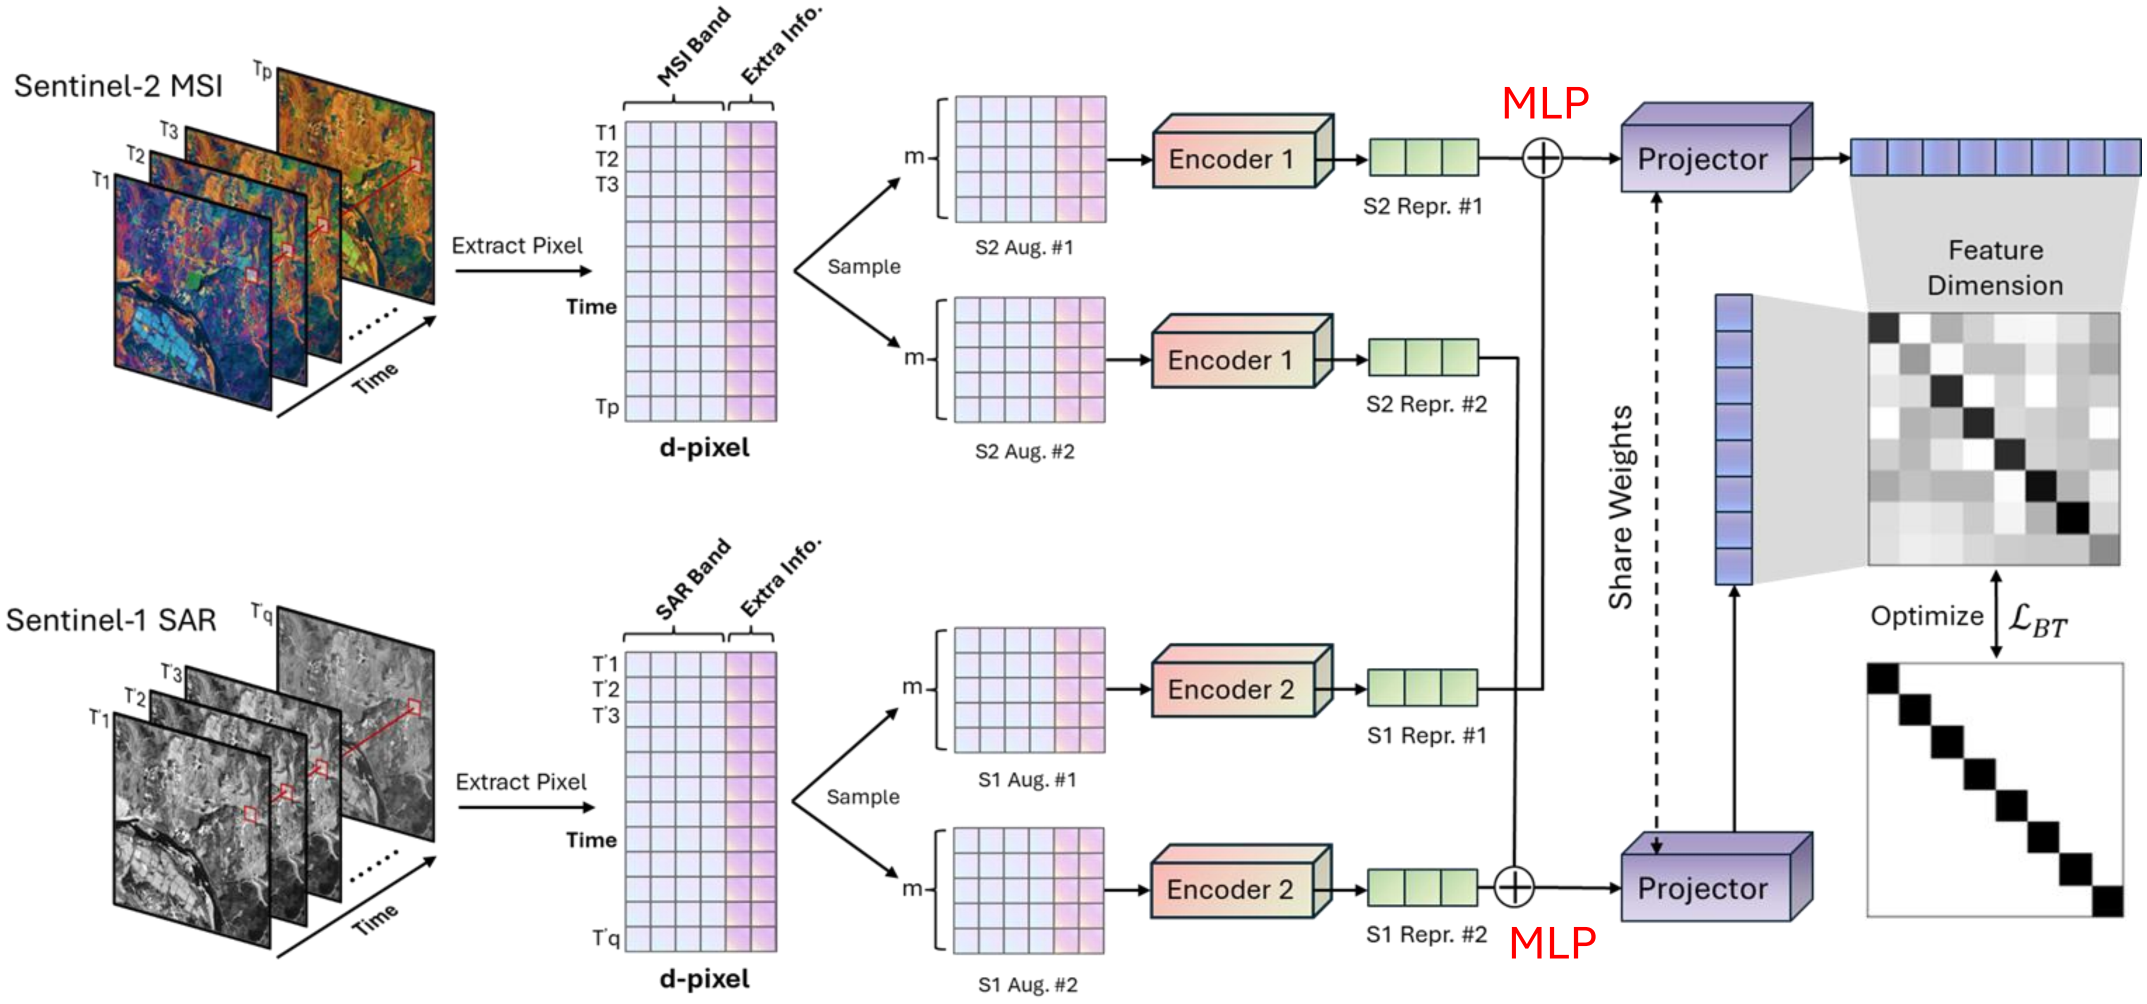

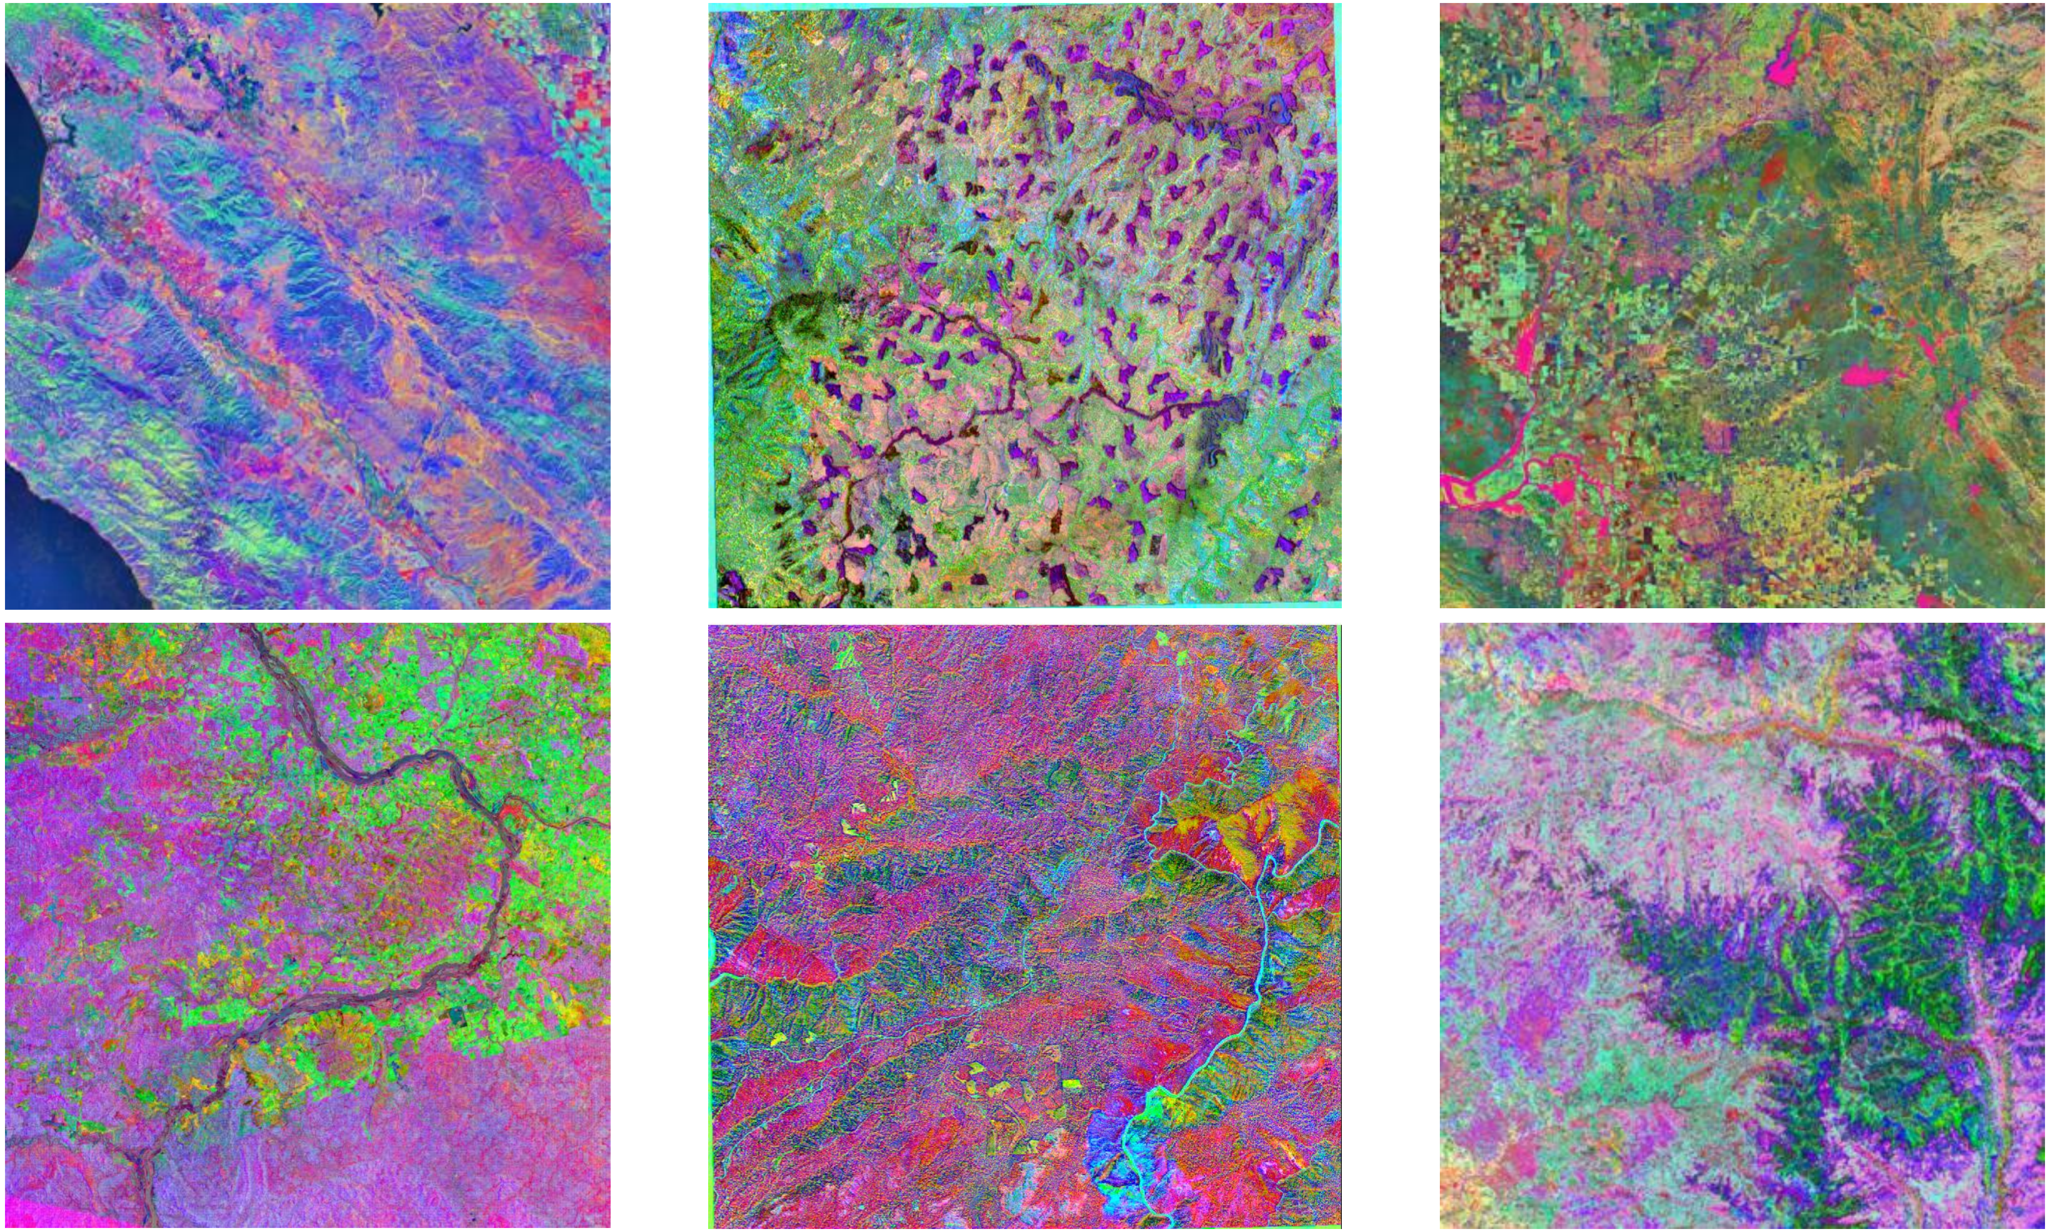

zf281@cam.ac.uk

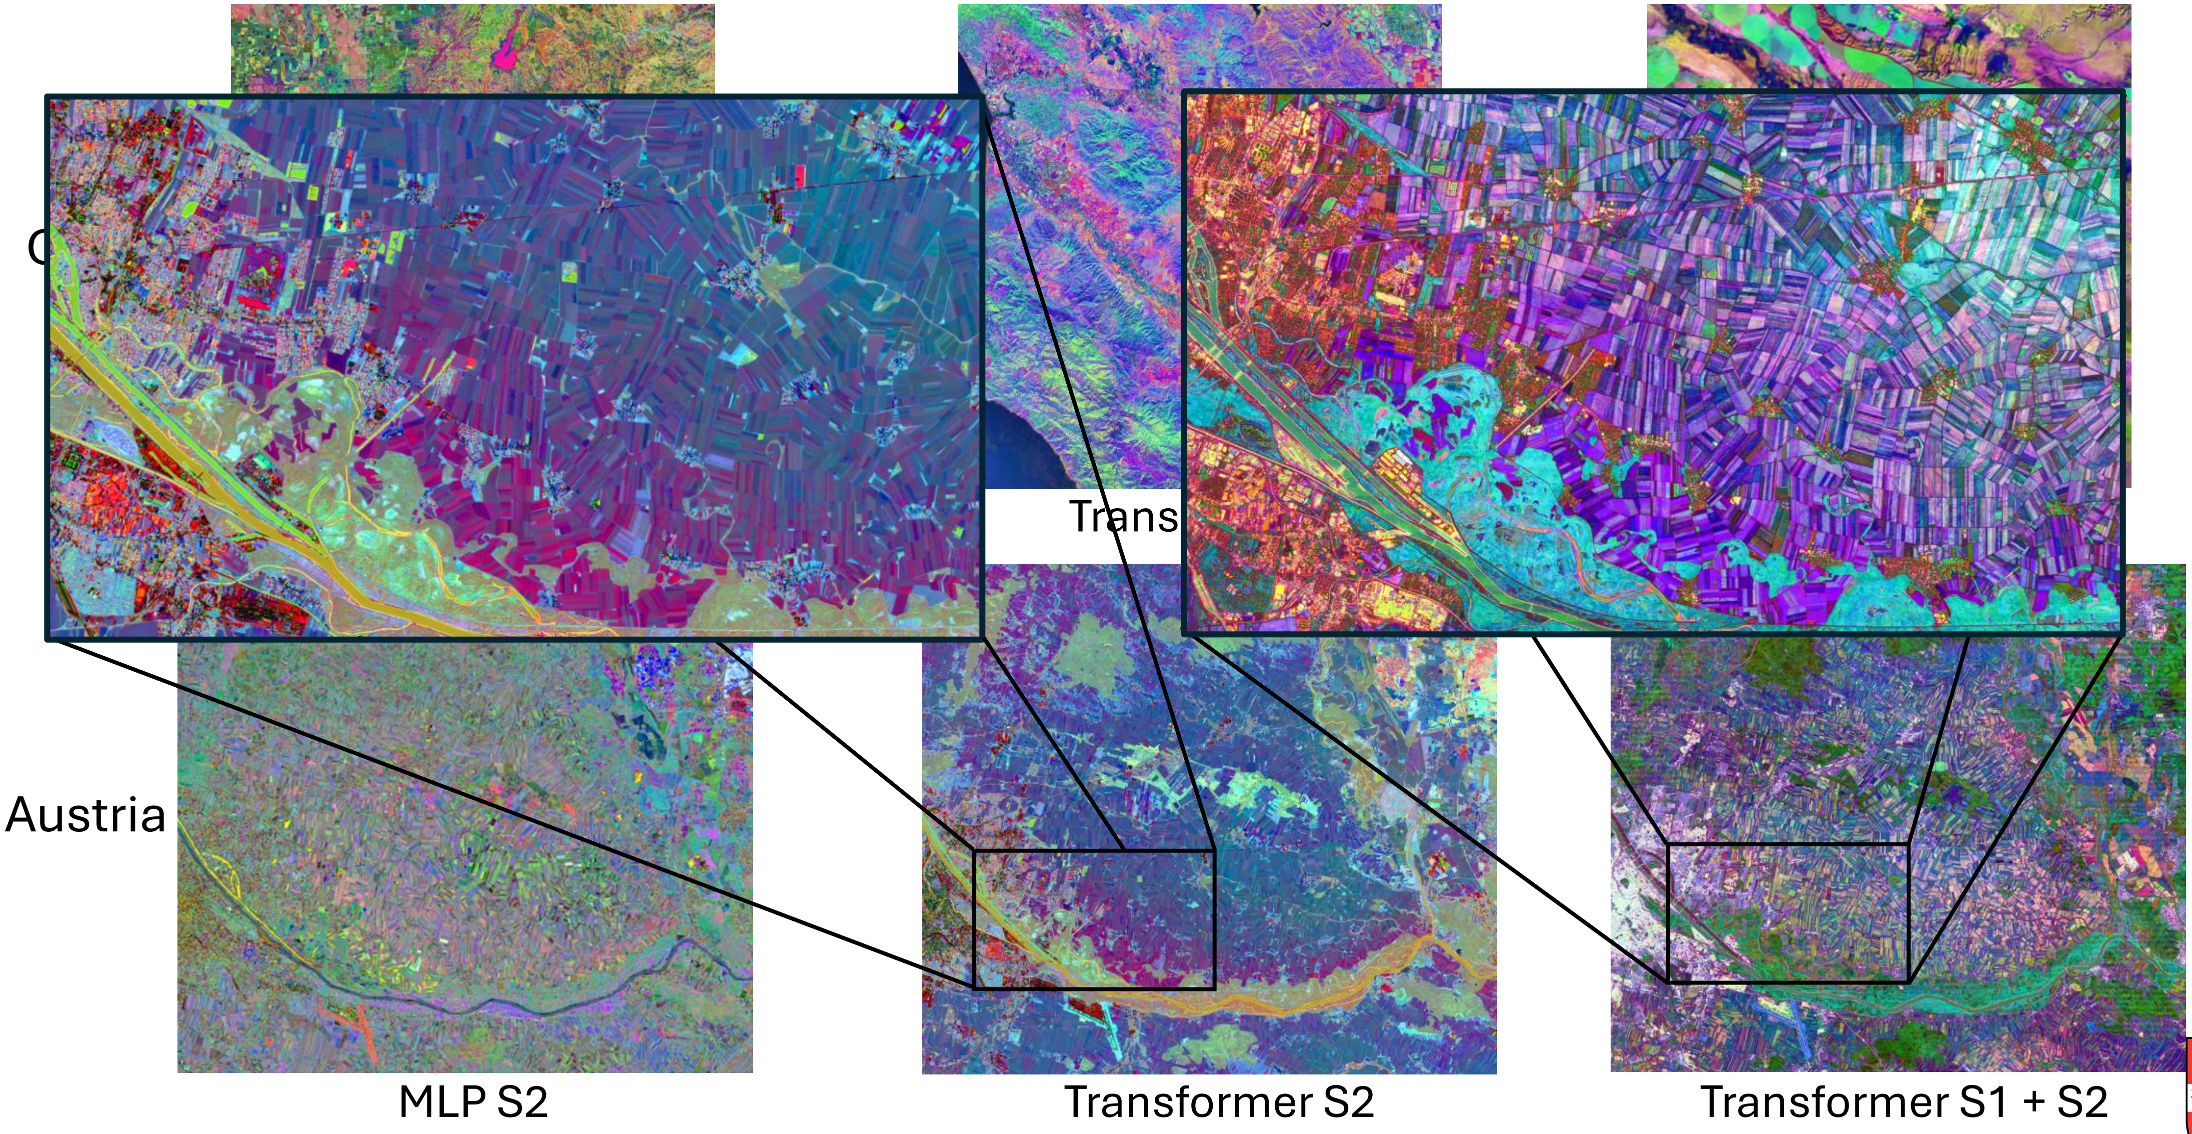
Austria
MLP S2
Transformer S2
Transformer S1 + S2

Using representations in downstream tasks

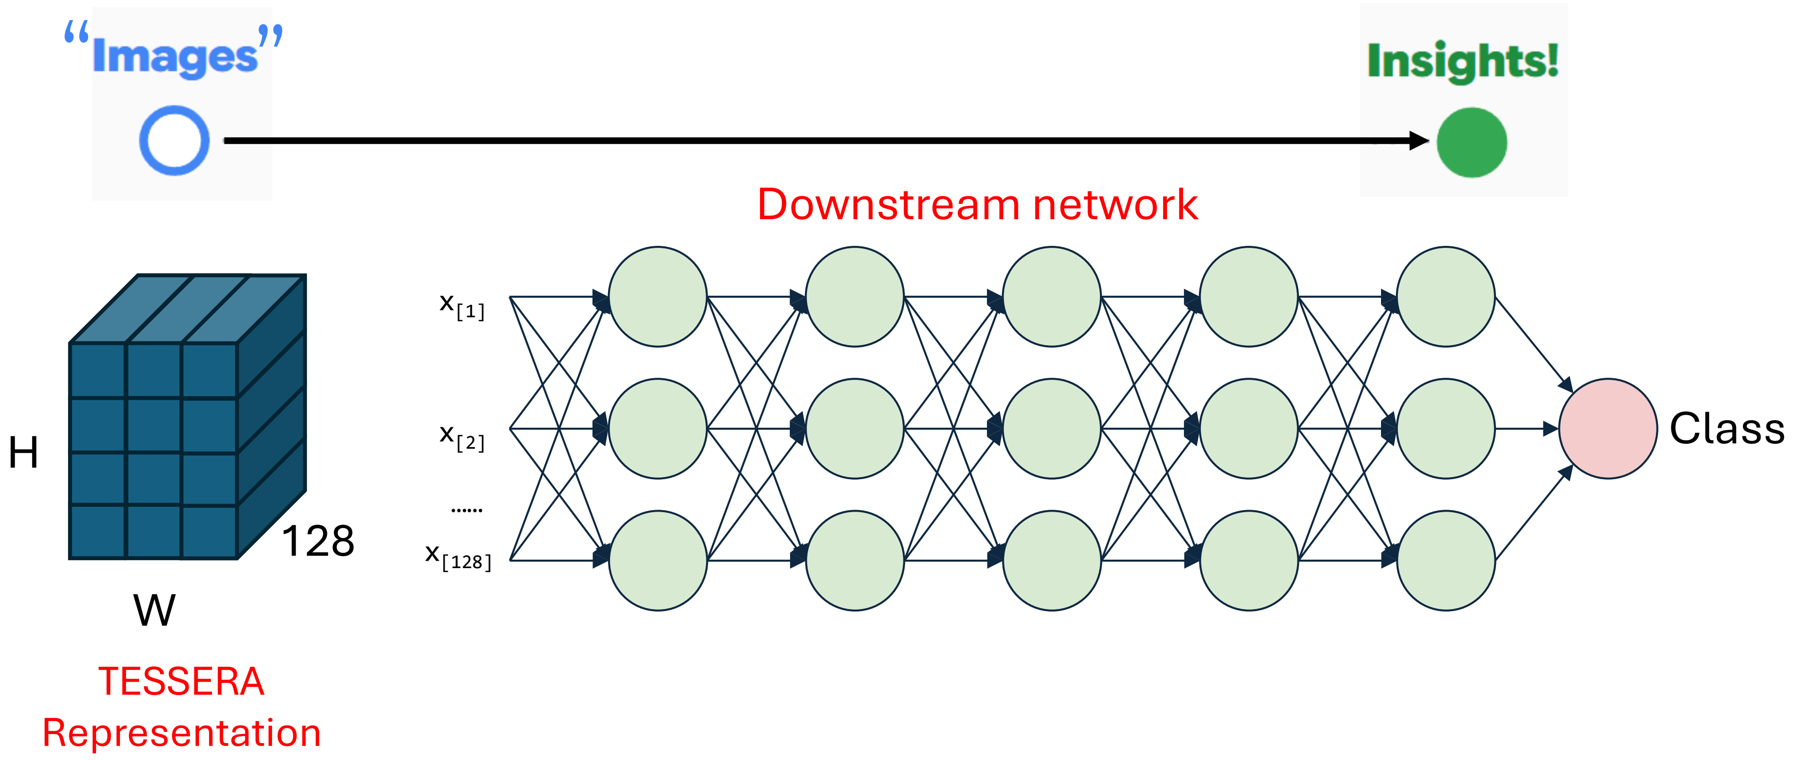

zf281@cam.ac.uk

Data used to train the TESSERA Model

```
s2_s1_global_project
├ data------------------------------------------ Folder to store raw S2 tiles
│ ├ 19FEF--------------------------------------- MGRS number
│ │ └ red
│ │   └ S2A_19FEF_20220101_0_L2A.tiff
├ data_processed-------------------------------- Folder to store processed S2 and S1 data
│ ├ 19FEF
│ │ ├ band_mean.npy
│ │ ├ band_std.npy
│ │ ├ bands.npy
│ │ ├ doys.npy
│ │ ├ masks.npy
│ │ ├ sar_ascending.npy
│ │ ├ sar_ascending_doy.npy
│ │ ├ sar_descending.npy
│ │ └ sar_descending_doy.npy
├ data_sar ------------------------------------- Folder to store raw S1 tiles
├ data_sar_processed --------------------------- Folder to store calibrated S1 tiff files
├ all_land_tiles.txt---------------------------- Text file including all the MGRS numbers (13672 in total)
├ Cargo.toml-----------┐
├ process_tile.rs------┤
├ Dockerfile-----------┼------------------------ S2 processing program by Robin; I made minor changes
├ s2_process_tile------┘
├ download_s2_tile.py -------------------------- Download S2 tiles
├ s1_sar_helper.py ----------------------------- S1 SAR geographic alignment, downsampling, stacking, etc.
└ orchestrate.py ------------------------------- Orchestrate all the above files in parallel
```

Trained on 3000+ MGRS tiles across the globe

zf281@cam.ac.uk

Some systems aspects

Data intensive task
- each tile is 150 GB

Need both raw compute and algorithmic design

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

Algorithms
- Per-tile d-pixel creation time : 2w -> 30 min
- Typical ROI inferencing: 10h -> 30s

Crop Classification

Classification Task

17 classes

Austrian Crop Dataset

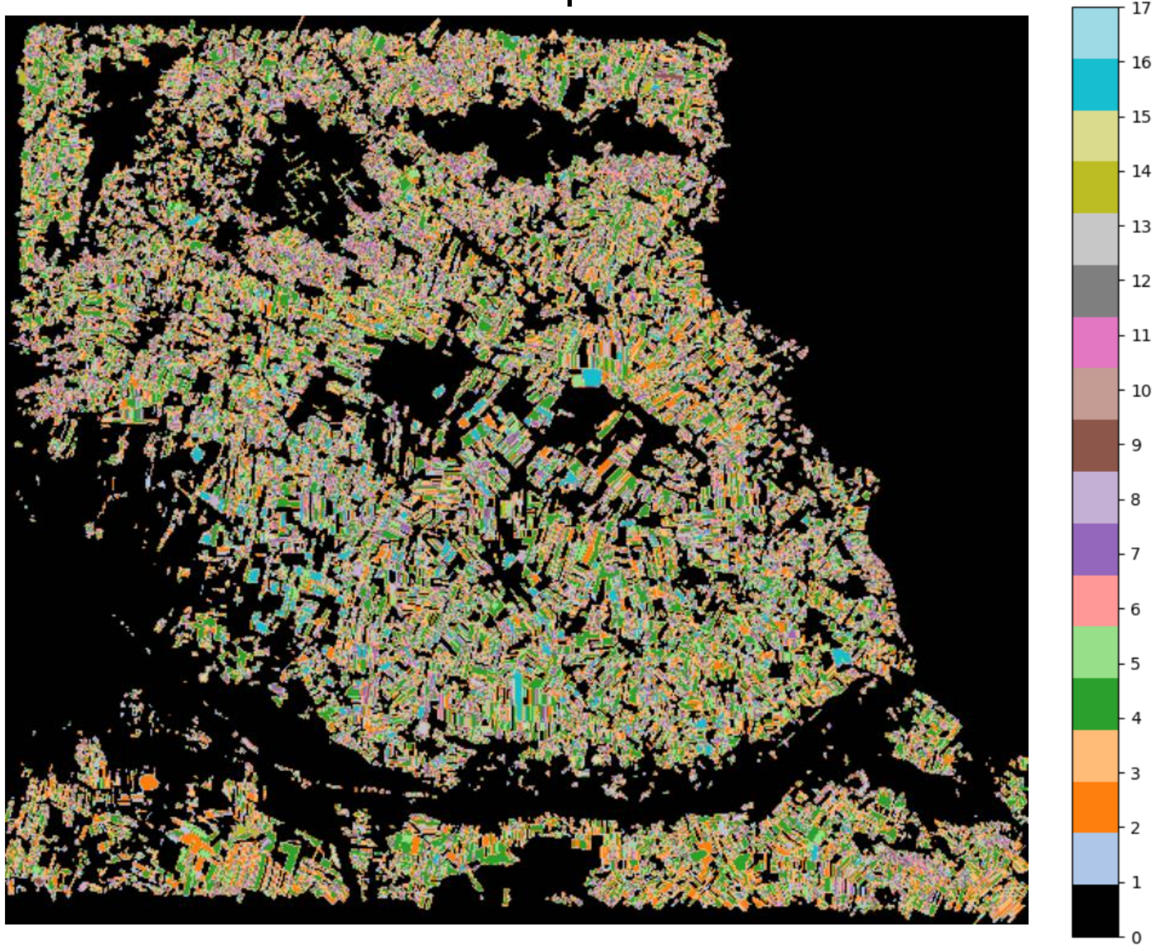

Name	Number of Samples	Percentage
Legume	2227	3.31%
Soy	5892	8.76%
Summer Grain	2475	3.68%
Winter Grain	24914	37.05%
Corn	6902	10.27%
Sunflower	207	0.31%
Mustard	1734	2.58%
Potato	2514	3.74%
Beet	1257	1.87%
Squash	2019	3.00%
Grapes	222	0.33%
Tree Fruit	347	0.52%
Cover Crop	1418	2.11%
Grass	2349	3.49%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

zf281@cam.ac.uk

Crop Classification

Preliminary Results

Encoder	Region	Data	Acc	F1
RF	Austria	S2	74.39	73.10
MLP	Austria	S2	75.21	73.58
Transformer	Austria	S2	78.43	77.61
Transformer+	**Austria**	**S2**	**79.48**	**78.77**

TESSERA outperforms RF

Model	Region	Data	Acc	F1
Transformer	California	S2	76.02	75.24
Transformer	**Austria**	**S2**	**78.43**	**77.61**

Representation can generalize

Crop Classification

What if we only have very limited samples with ground truth labels (close to real-world scenarios)

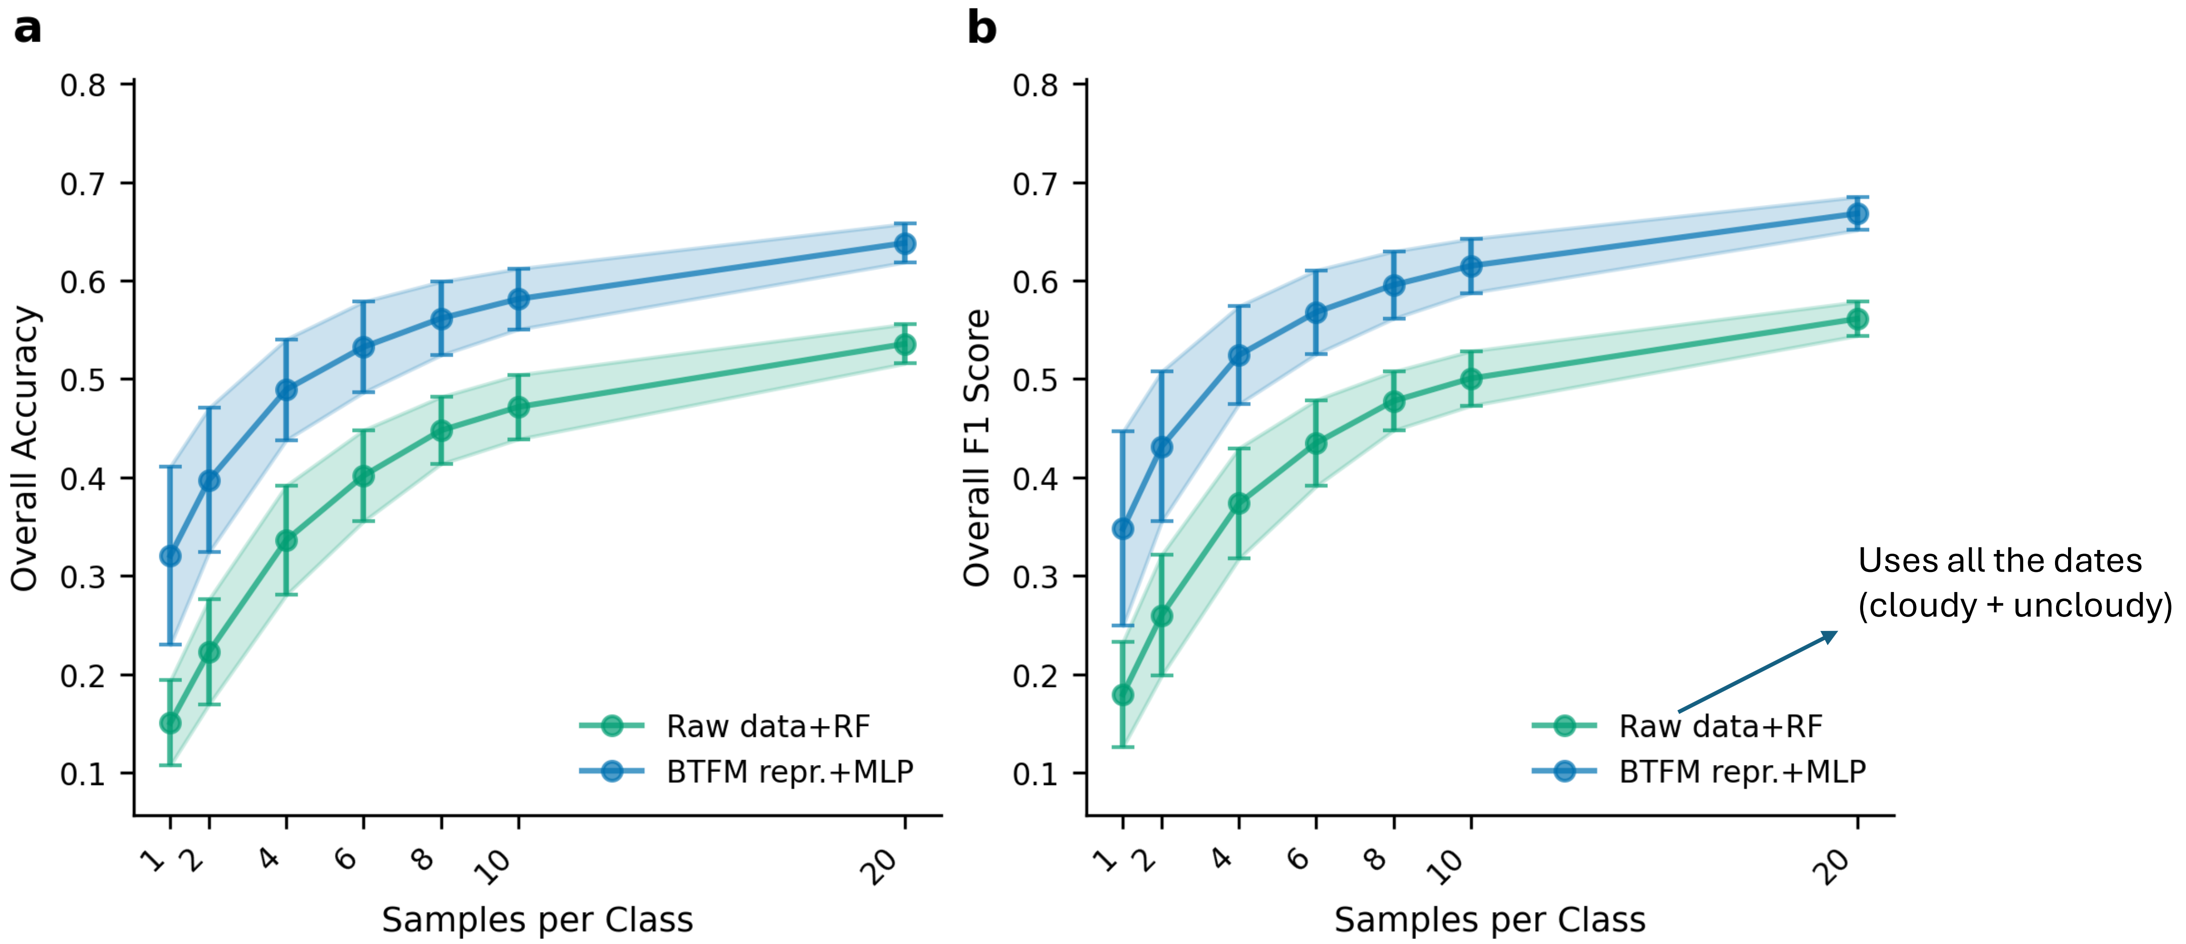

Competition: Google Embedding Field Model

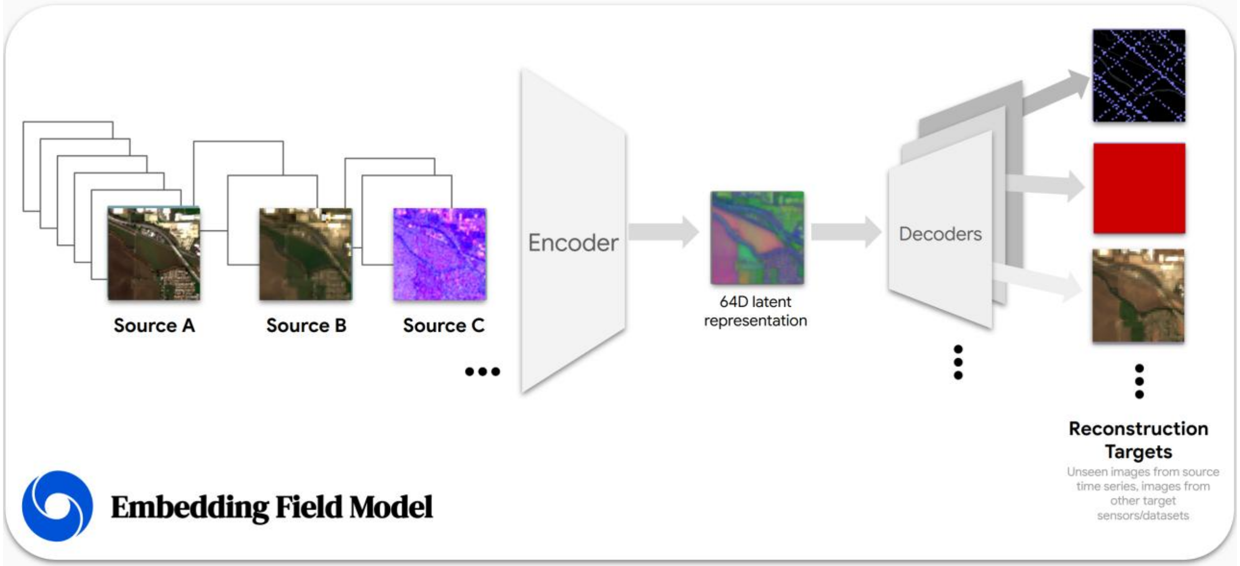

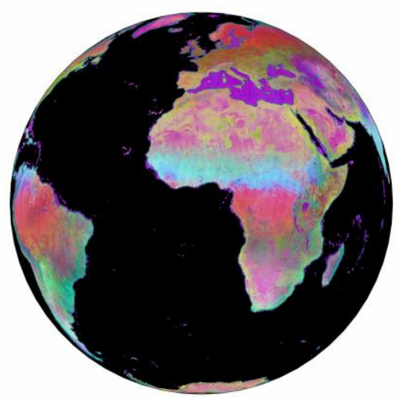

Table 1: EFM v2 data sources

Type	Dataset	Product	Bands	Resolution (m)	Usage
Optical	Sentinel-2	L1C	B2 (Blue), B3 (Green), B4 (Red), B8 (NIR), B11 (SWIR)	10, 20, 60	input, target
Optical, Thermal	Landsat-8, Landsat-9	L1C	B2 (Blue), B3 (Green), B4 (Red), B5 (NIR), B6 (SWIR), B8 (Panchromatic), B10 (Thermal)	15, 30, 100	input, target
C-band SAR	Sentinel-1A, Sentinel-1B	GRD	VV, VH, HH, HV, angle	10	input, target
L-band SAR	ALOS PALSAR ScanSAR	Level 2.2	HH, HV, lin	25	target
Elevation	Copernicus DEM	GLO-30	DEM (elevation)	30	target
LiDAR	GEDI	L2A	Relative height metrics (rh*)	25	target
Climate	ERA5-Land	Monthly aggregates	total precipitation (sum, min, max), air temperature 2m (and min, max), dew-point temperature 2m (and min, max), surface pressure (and min, max)	11132	target
Gravity fields	GRACE	Monthly mass grids	equivalent liquid water thickness	11132	target (@50%)
Land cover	National Land Cover Database	NLCD 2019, 2021	landcover	N/A	target (@50%)
Text	Wikipedia	geocoded articles	text embeddings	N/A	target
Text	GBIF	Research-grade obs	text embeddings (class, genus, and species)	N/A	target

zf281@cam.ac.uk

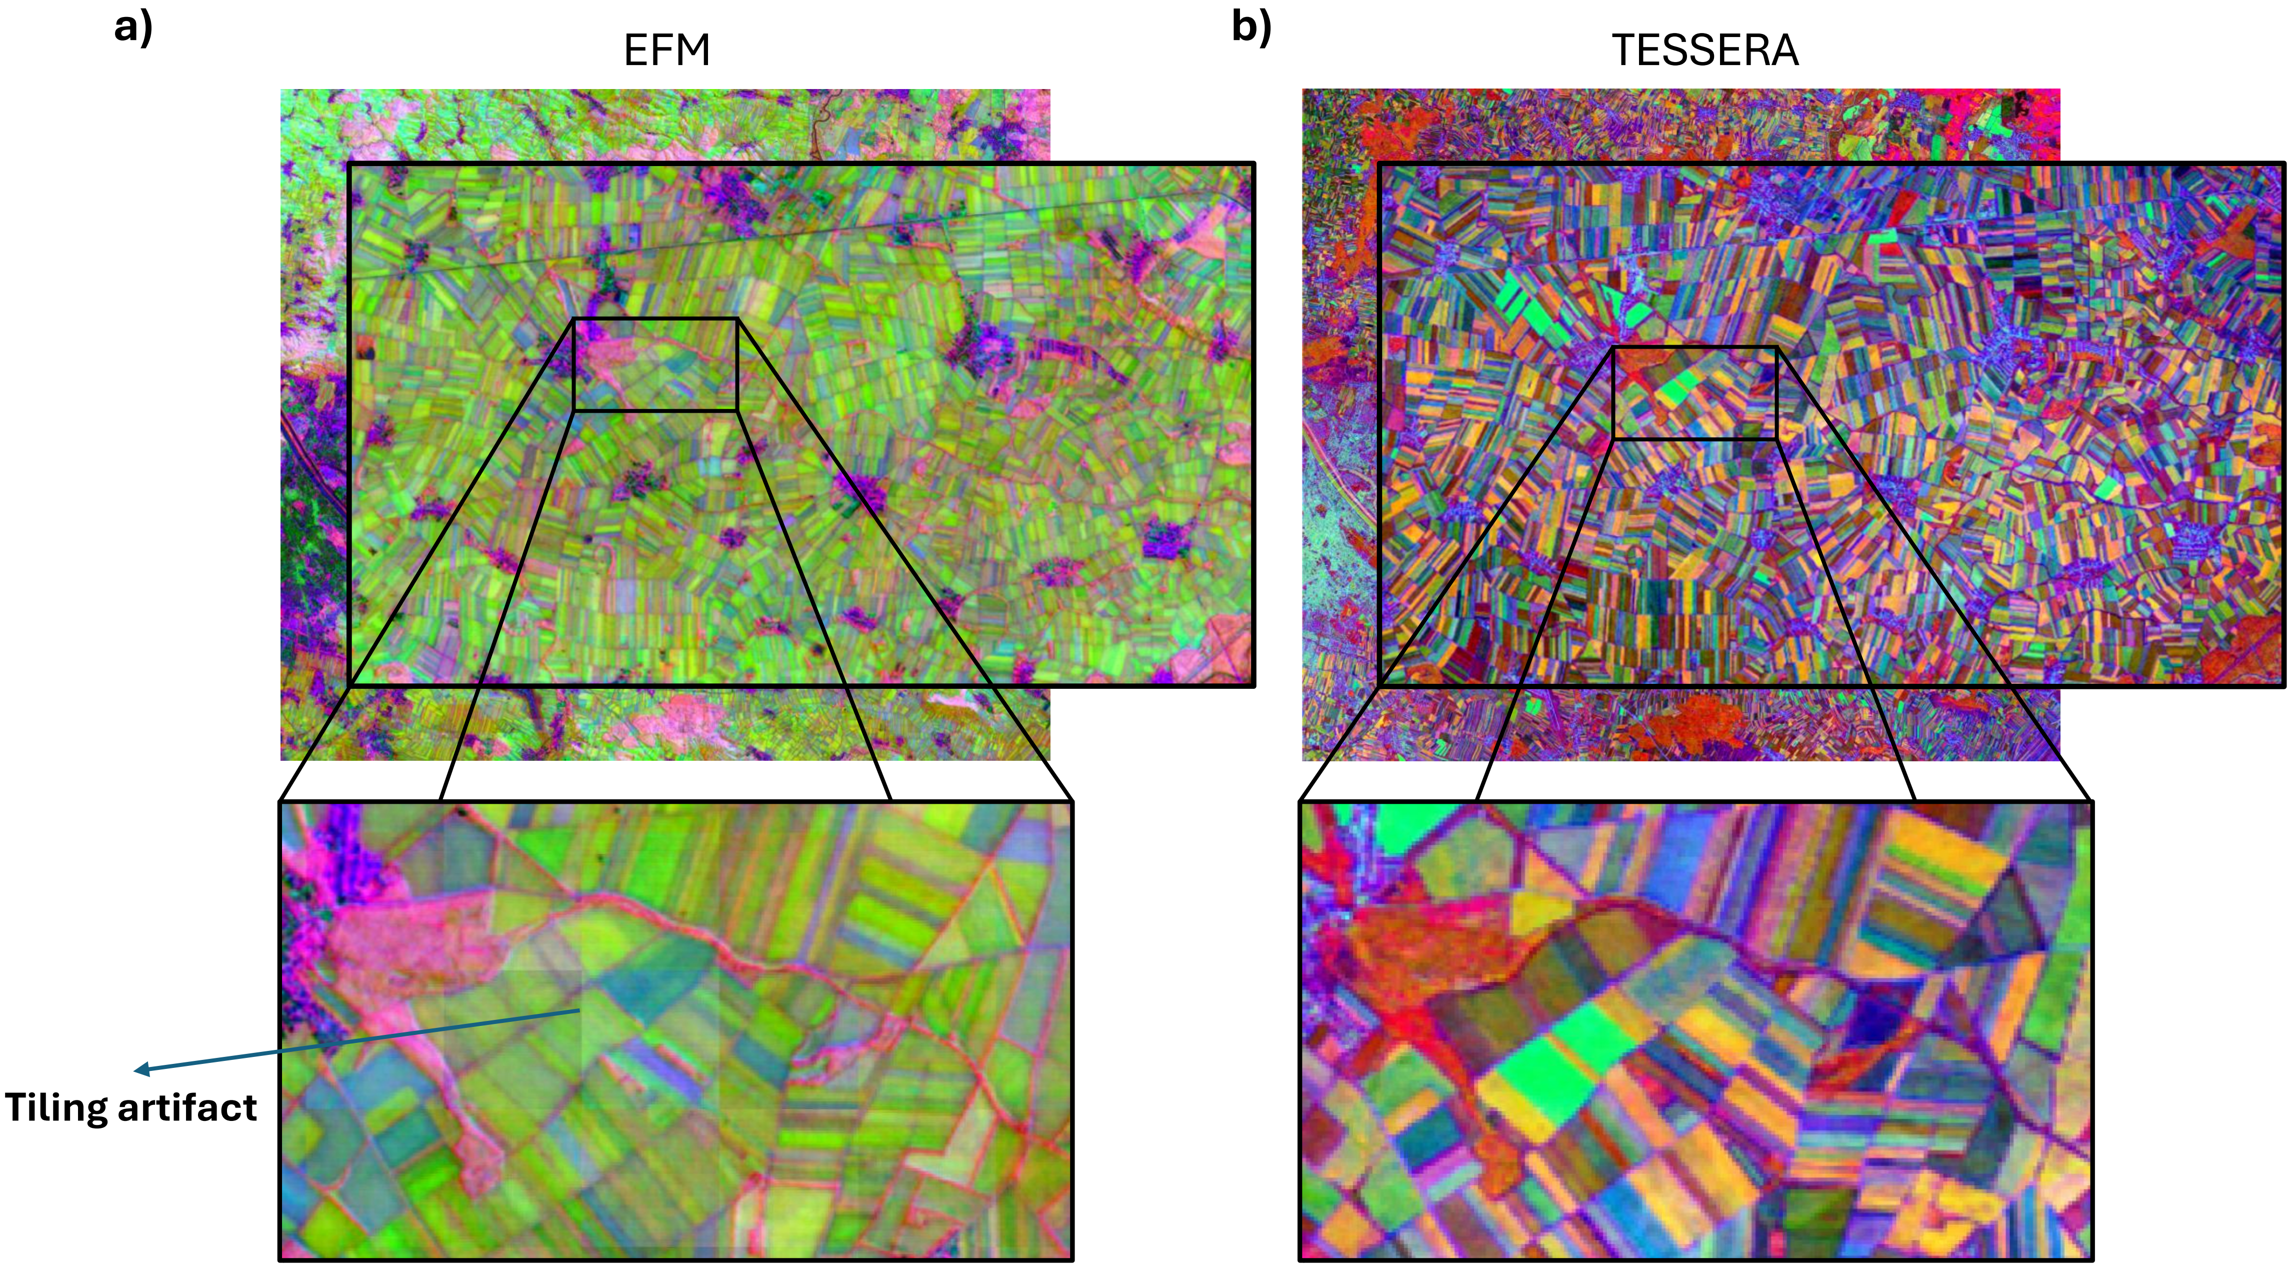
a)
EFM
Tiling artifact
b)
TESSERA

Crop Classification

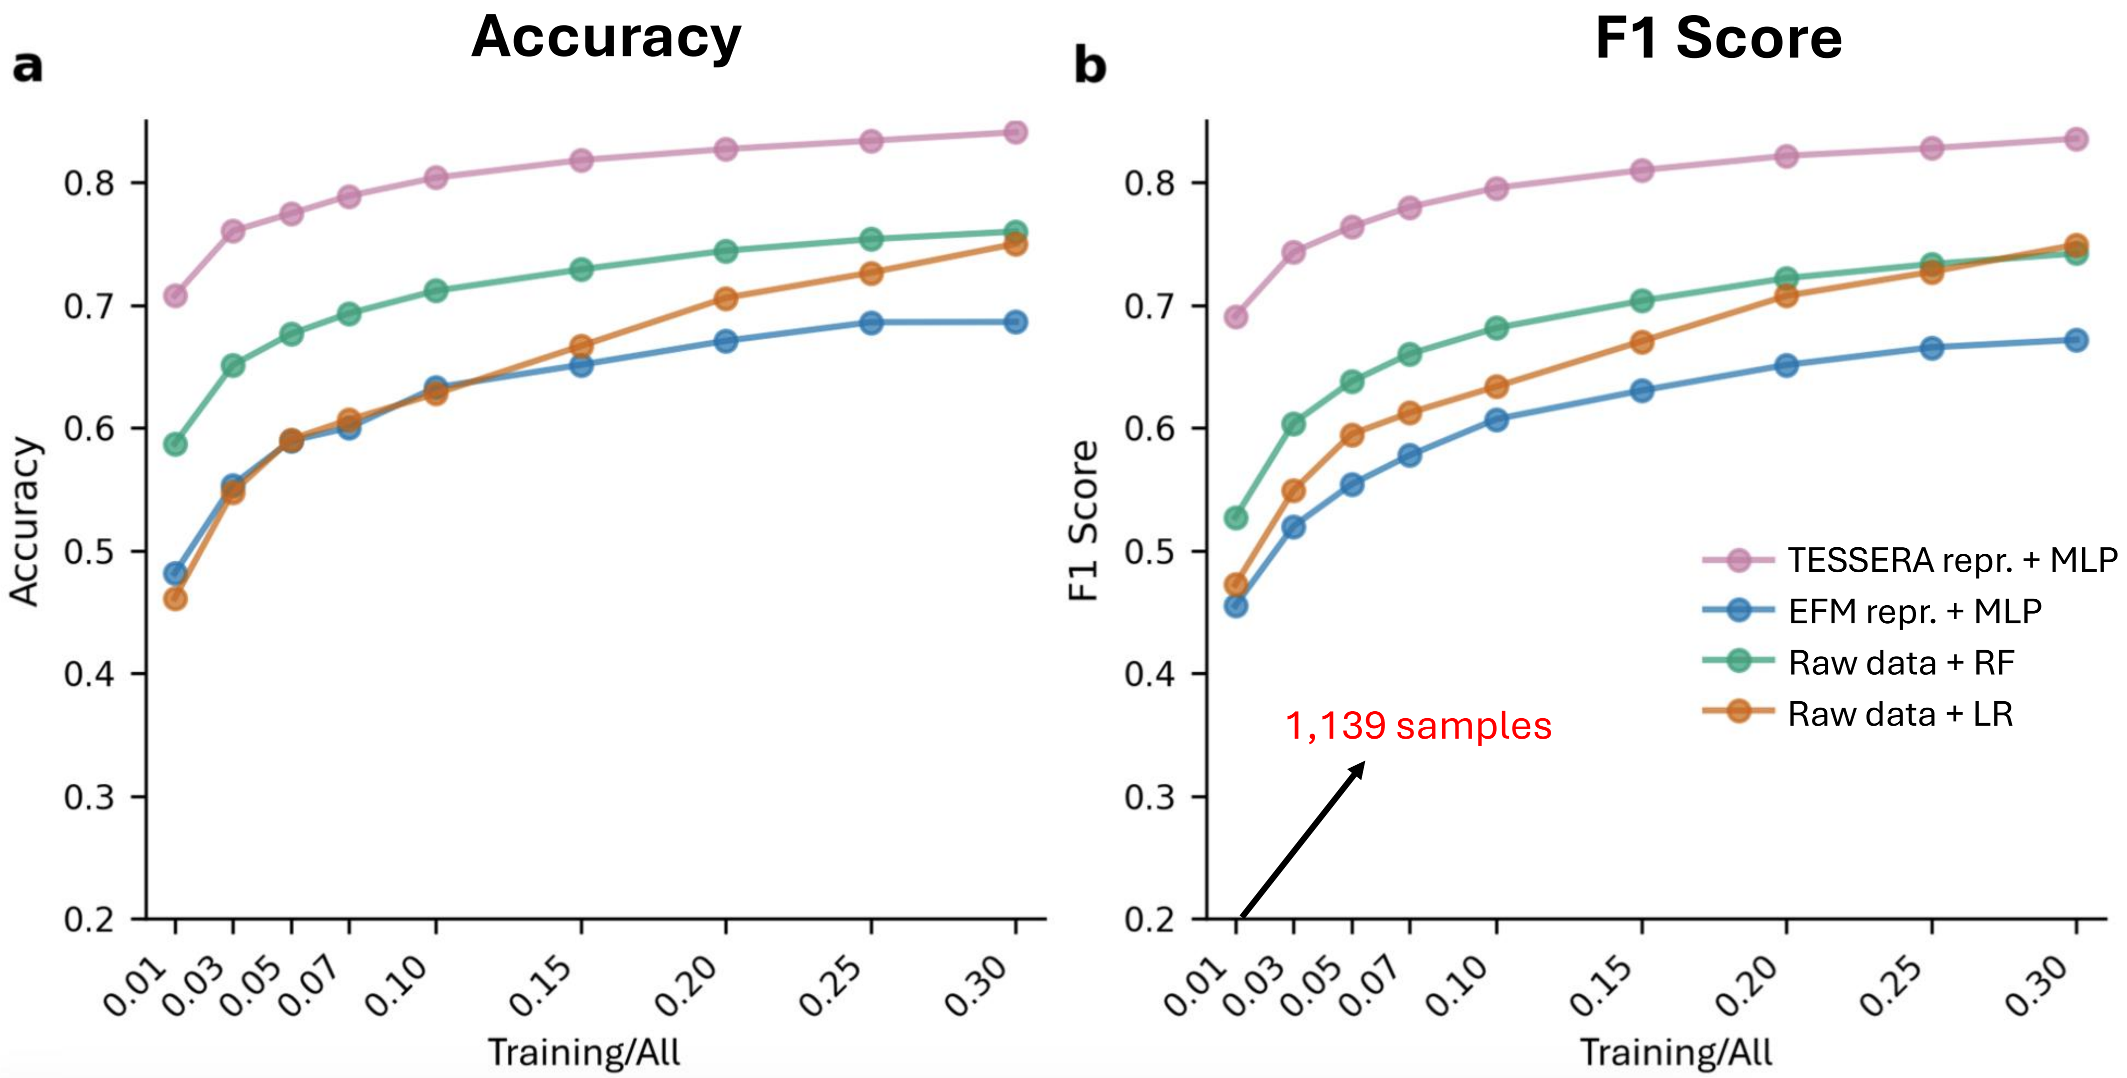

zf281@cam.ac.uk

Predicting Canopy Height

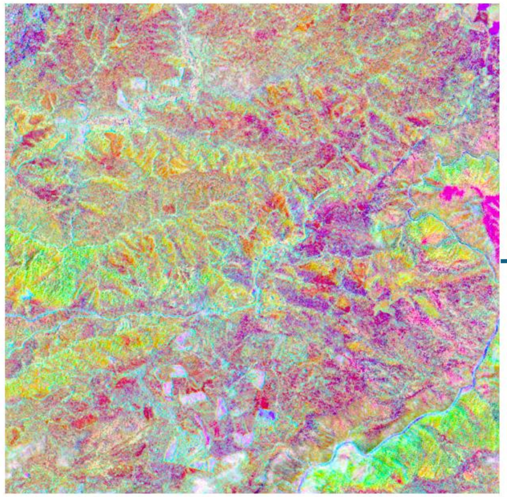

Representations

Downstream Model (UNet)

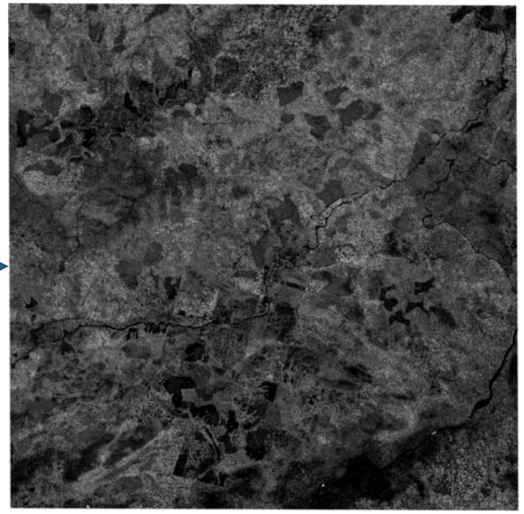

Canopy Height from Airborne LiDAR

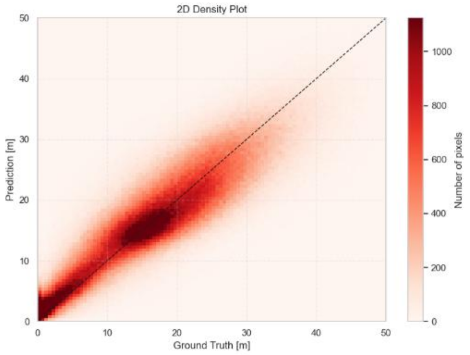

TESSERA
R^2 = 0.78
MAE = 3.7m

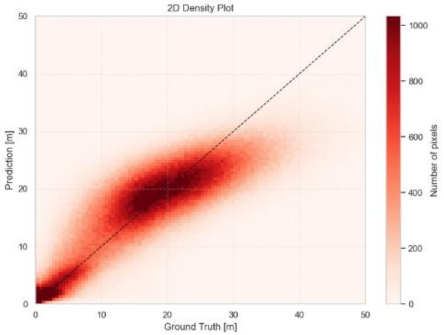

EFM
R^2 = 0.68
MAE = 4.4m

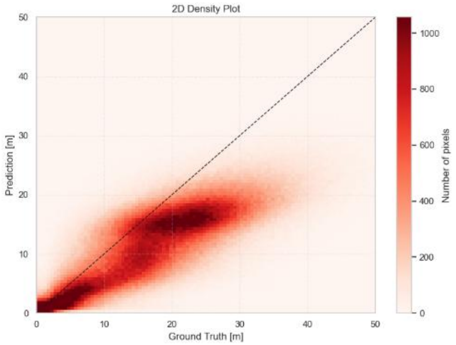

Sentinel-2 Medians
R^2 = 0.34
MAE = 6.5m

Forest in Northern California

Predicting Canopy Height

Borneo Forest

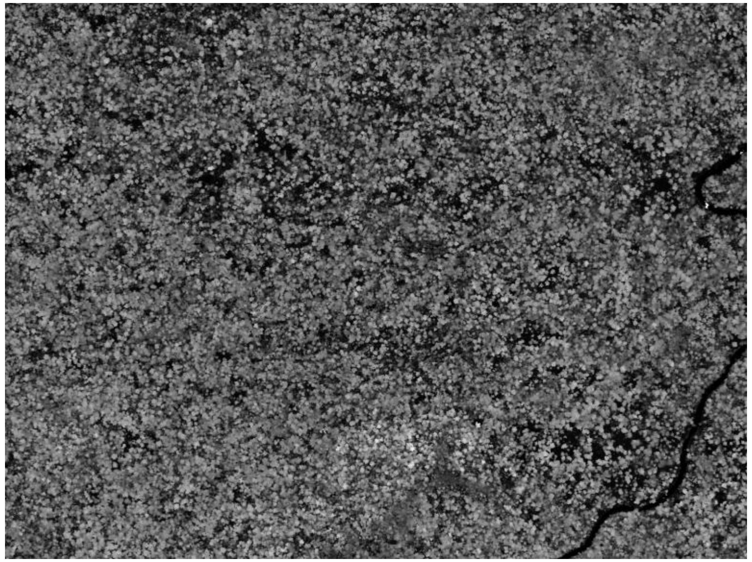

Different inputs, different predictive power

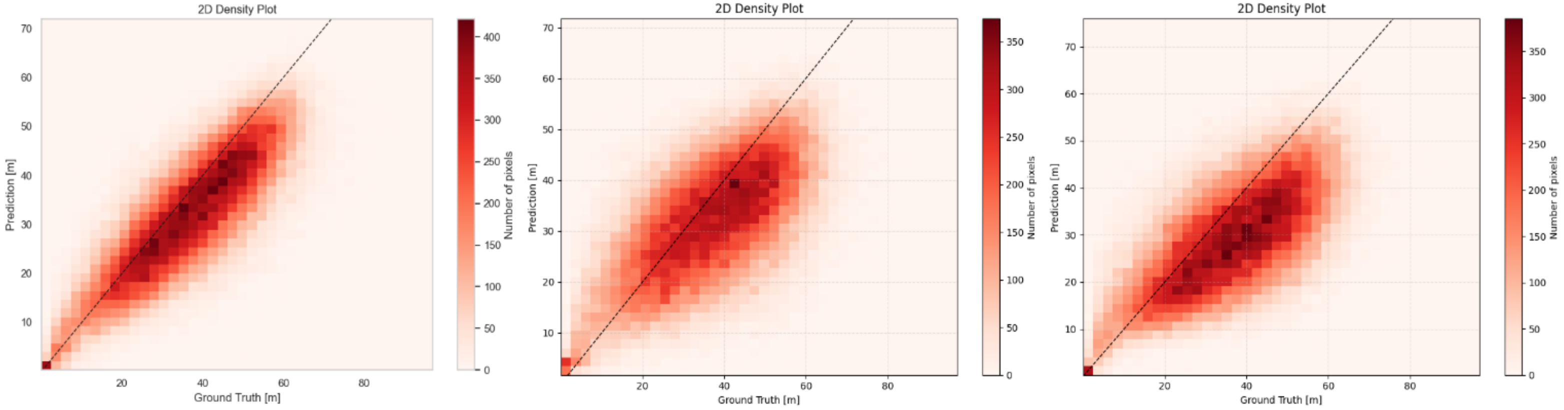

TESSERA
$R^2 = 0.55$

EFM
$R^2 = 0.25$

Sentinel-2 Medians
$R^2 = 0.17$

Ground Truth

Predicted

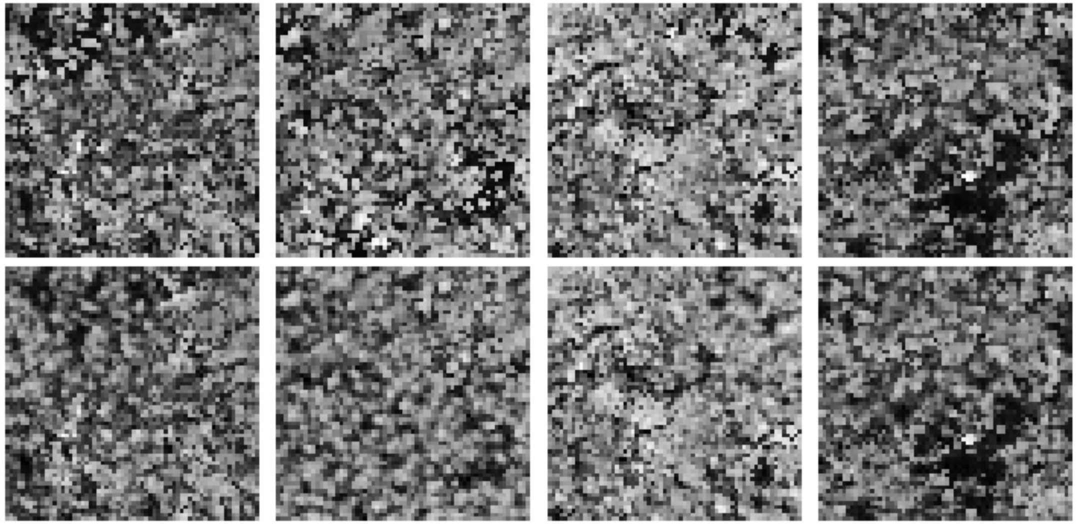

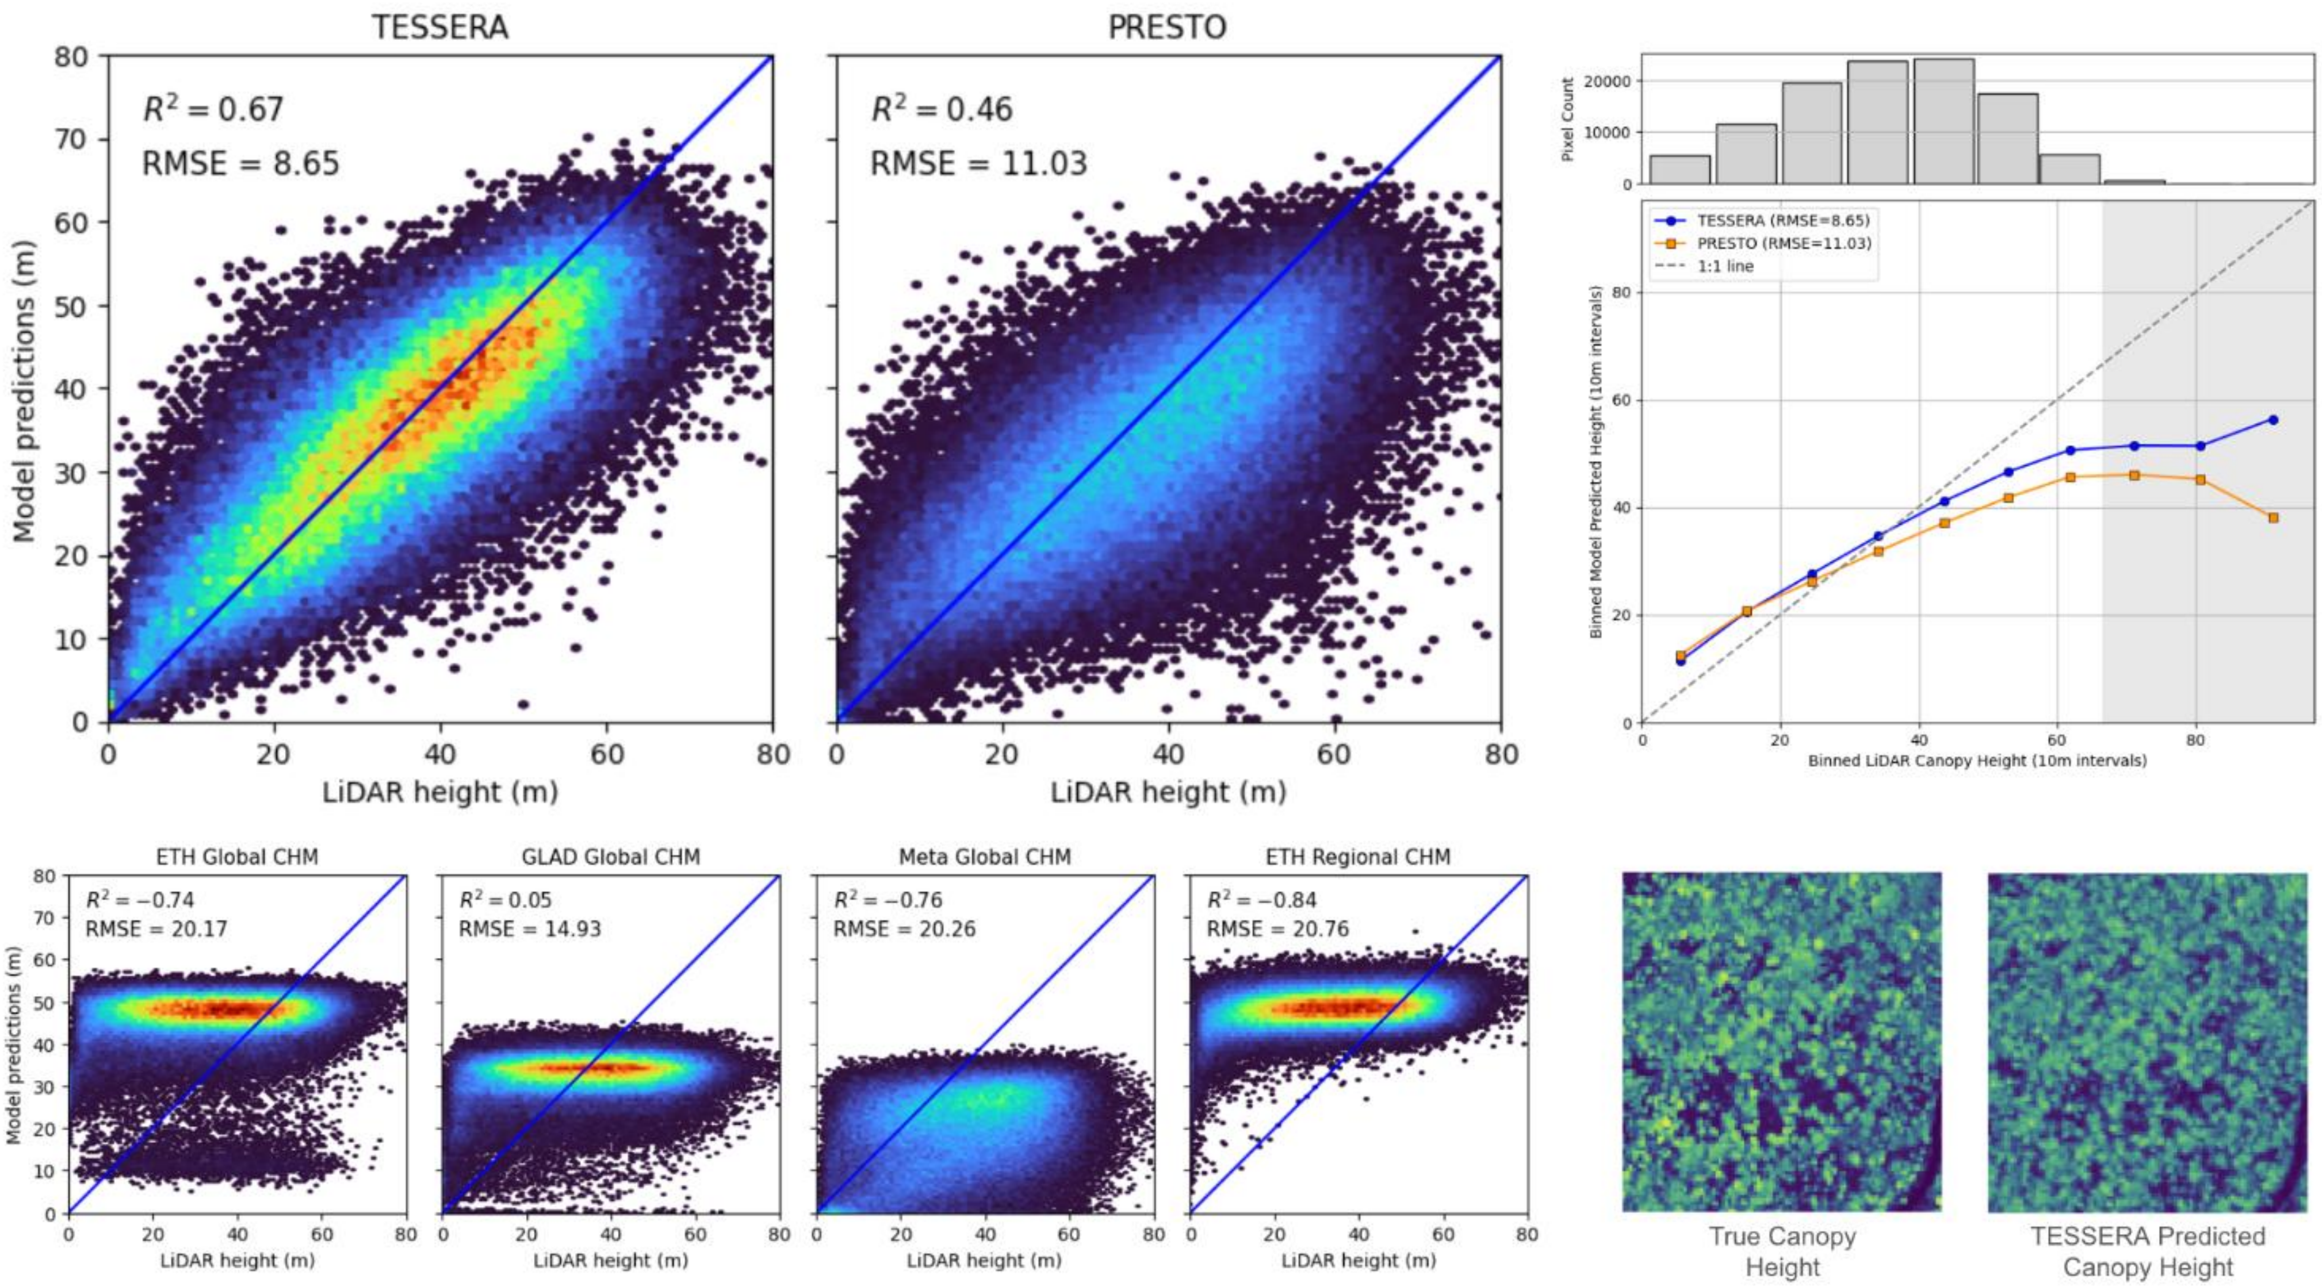
TESSERA
$R^2 = 0.67$
RMSE = 8.65
PRESTO
$R^2 = 0.46$
RMSE = 11.03
Model predictions (m)
LiDAR height (m)
Pixel Count
TESSERA (RMSE=8.65)
PRESTO (RMSE=11.03)
1:1 line
Binned Model Predicted Height (10m intervals)
Binned LiDAR Canopy Height (10m intervals)
ETH Global CHM
$R^2 = -0.74$
RMSE = 20.17
GLAD Global CHM
$R^2 = 0.05$
RMSE = 14.93
Meta Global CHM
$R^2 = -0.76$
RMSE = 20.26
ETH Regional CHM
$R^2 = -0.84$
RMSE = 20.76
True Canopy Height
TESSERA Predicted Canopy Height

Detecting Disturbances:
Fires in California

Burned Area 1

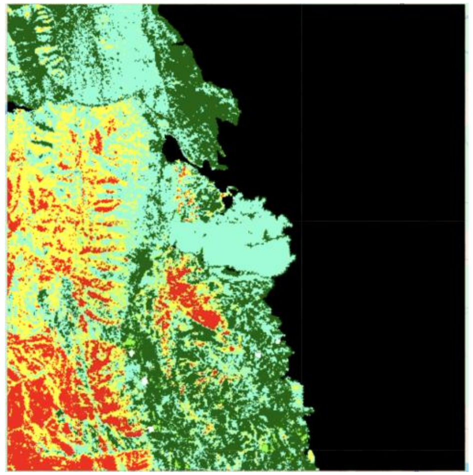

Burned Area 2

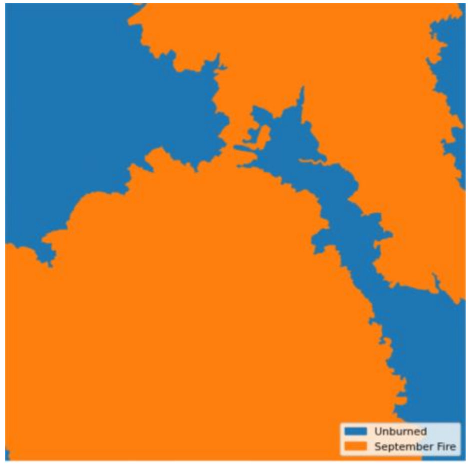

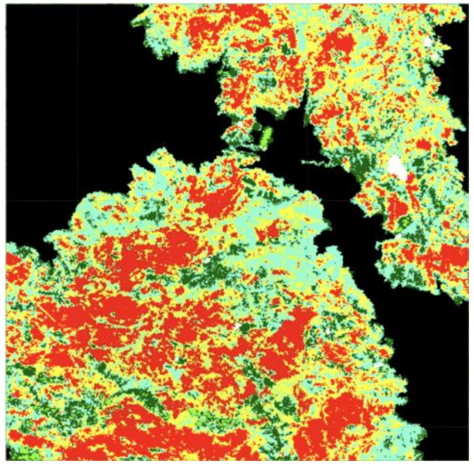

Fires in 2020

Fire Severity

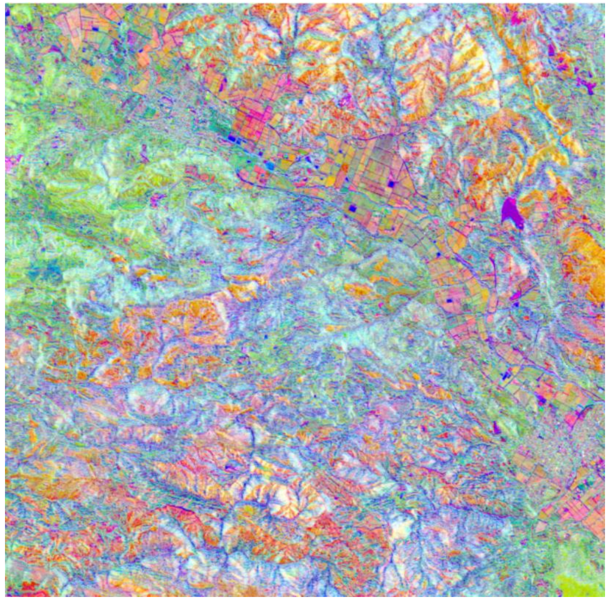

TESSERA Representations For 2020

Detecting Disturbances: *Fires in California*

UMAP (Dimensionality Reduction 128 -> 2)

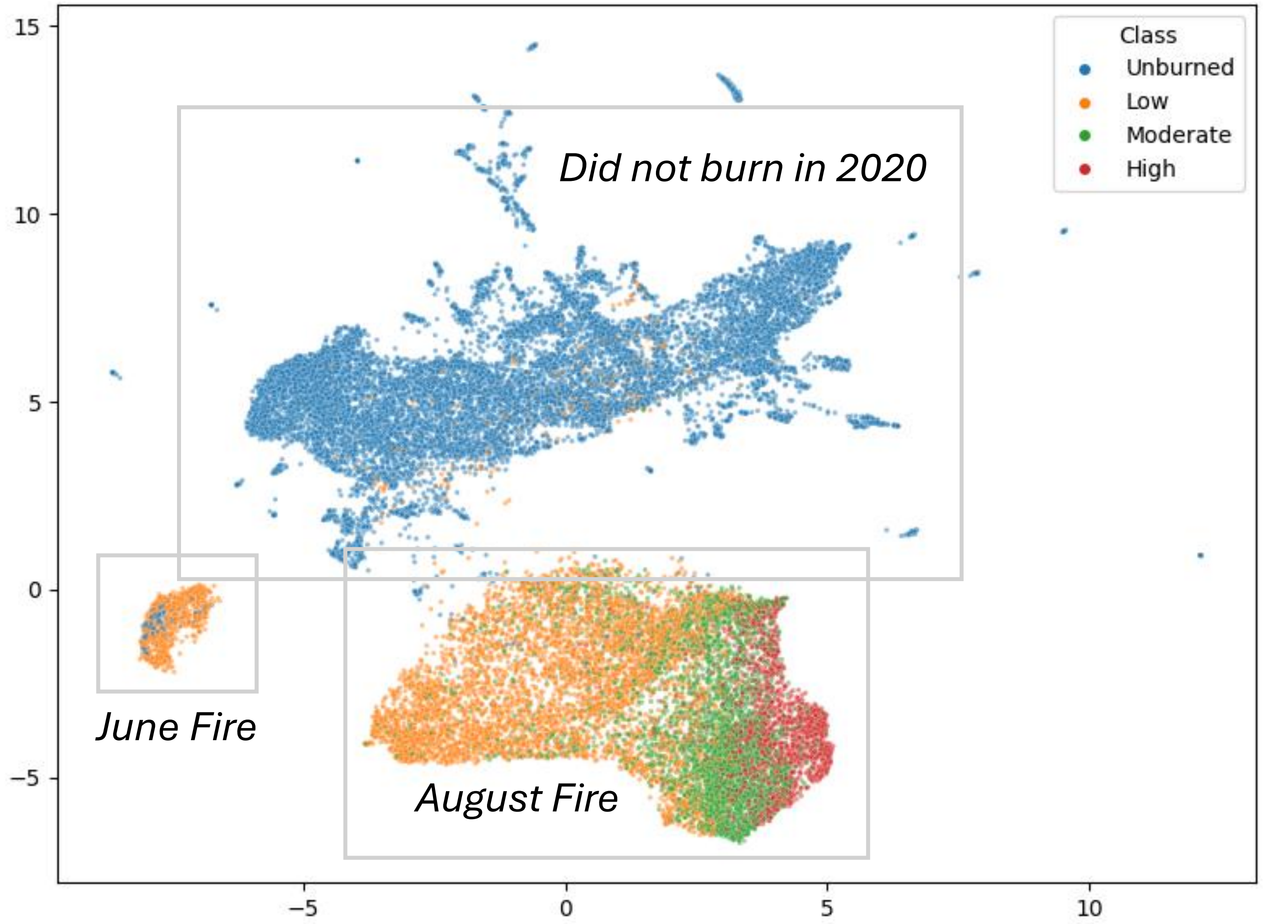

Burned Area 1

Burned Area 2

Fungal Biodiversity

Ground truth (Mycorrhizal biodiversity from SPUN)

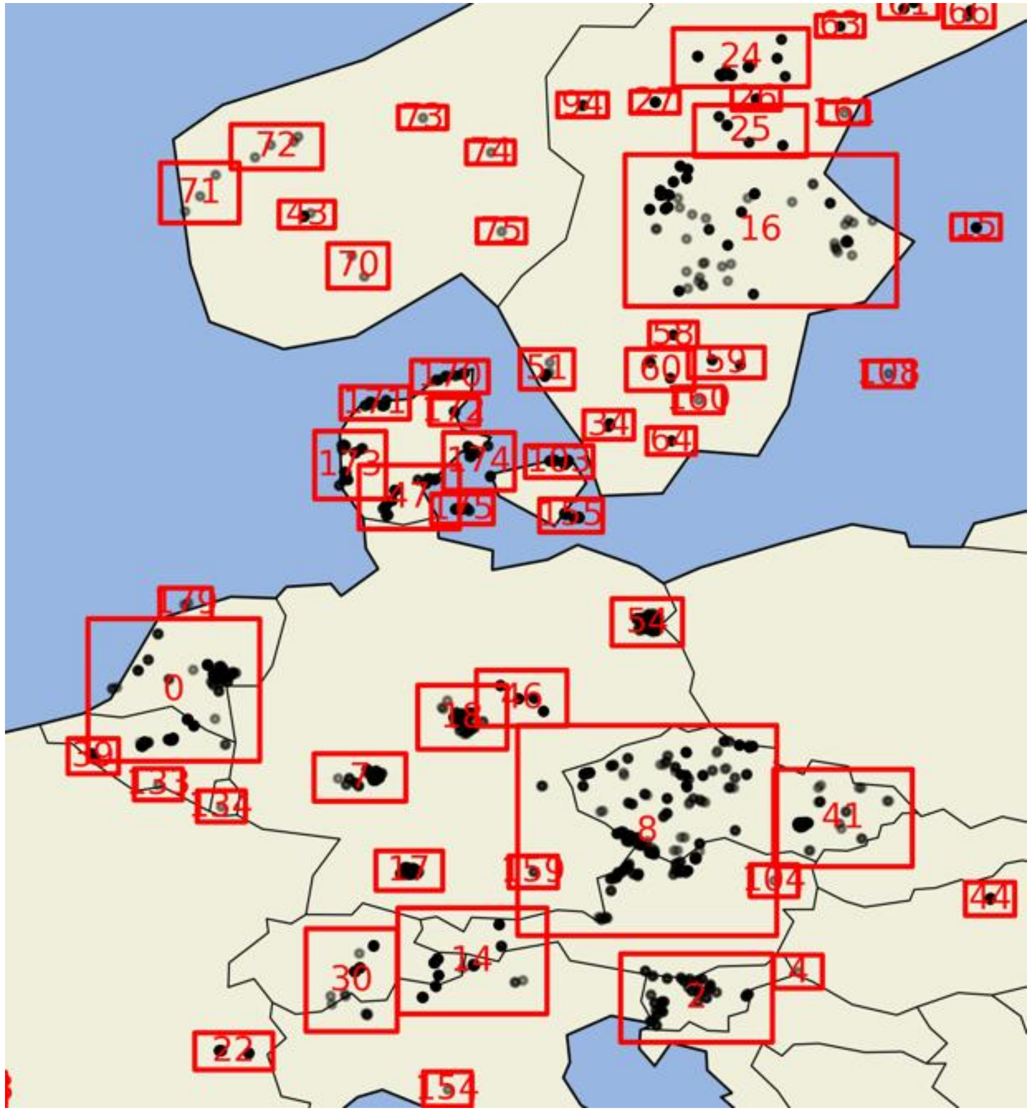

Fungal Biodiversity

Steps:

- From lat-lon data, get max,min lat-lon, create bounding box **(this is the ROI)**
- Download S1/S2 data, **create d-pixels**
- **Inference** with checkpoint, get **representation**
- Match representations to ground truth sample (x-y pair)

Apply random forest/xgboost, **train/test**

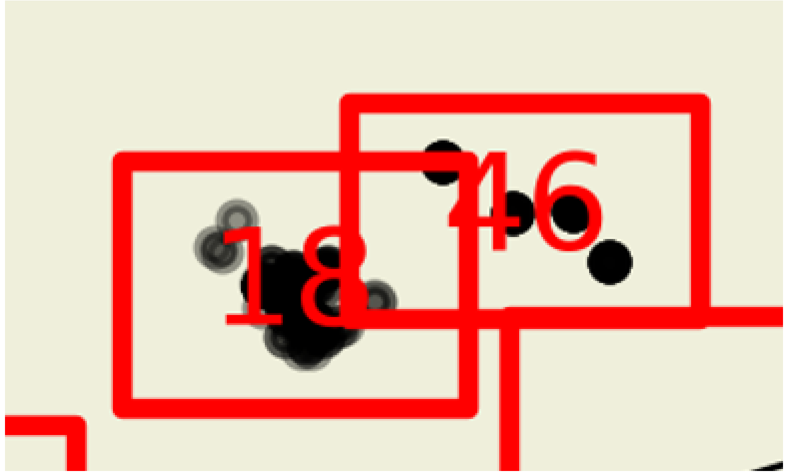

Fungal Biodiversity

Early Results

	Germany	Slovenia + Kosovo + Sicily	Estonia + Latvia
R^2	~0.65	~0.5	~0.2

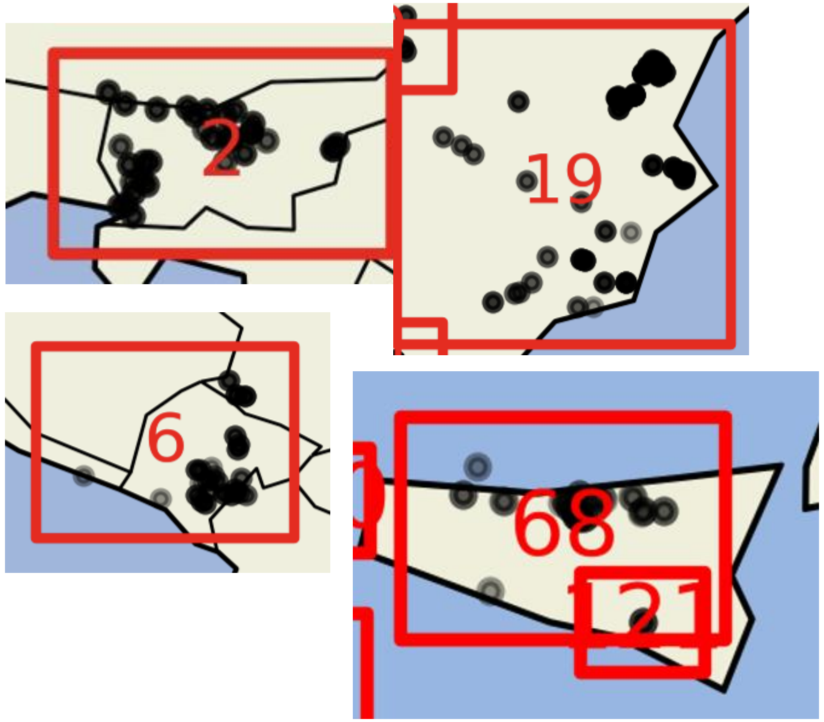

For Estonia, possibly some underlying ecological reasons
(Seasonal waterlogging messing with satellite reflectances)

Land Change Detection

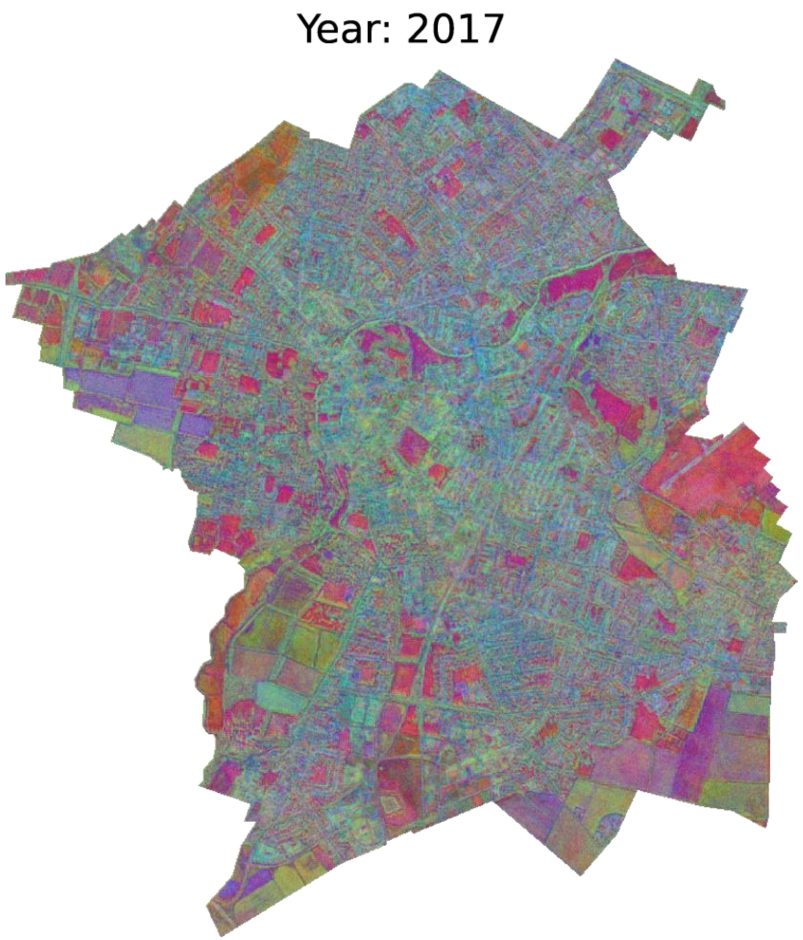

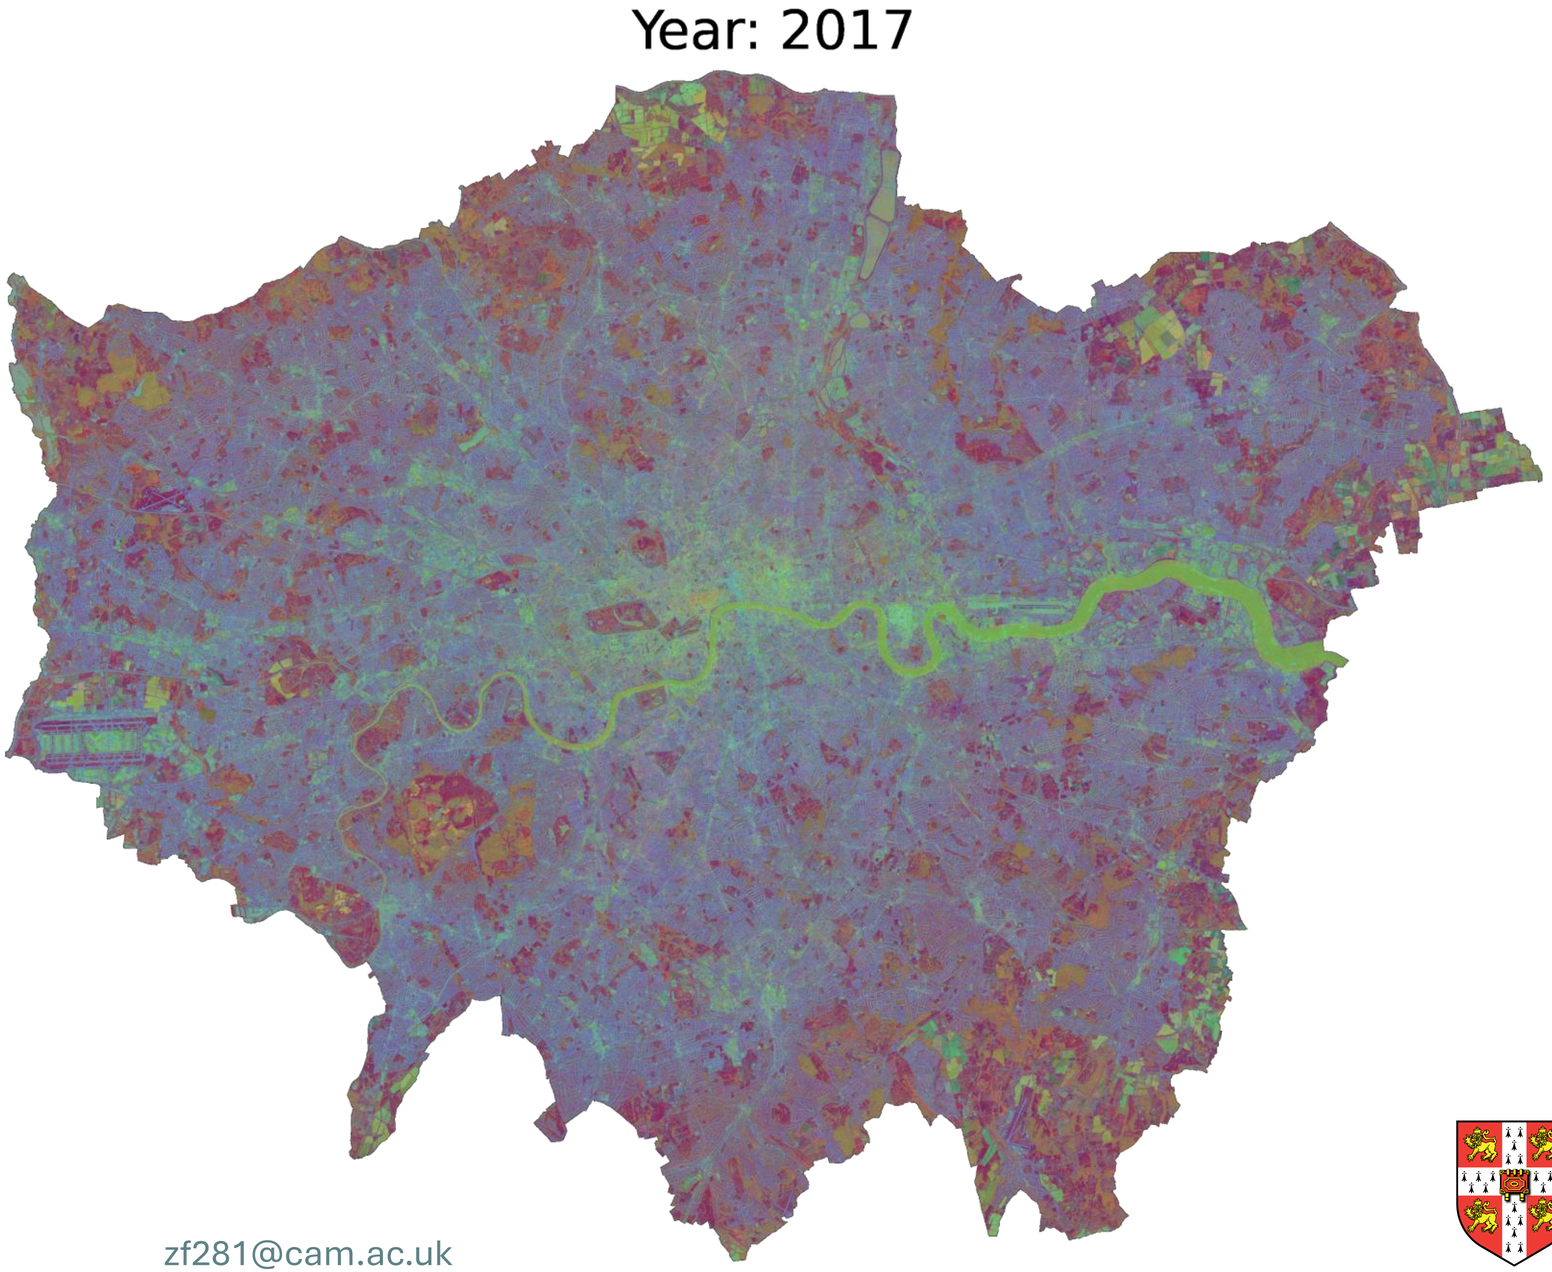

zf281@cam.ac.uk

Land Change Detection

2017-2018

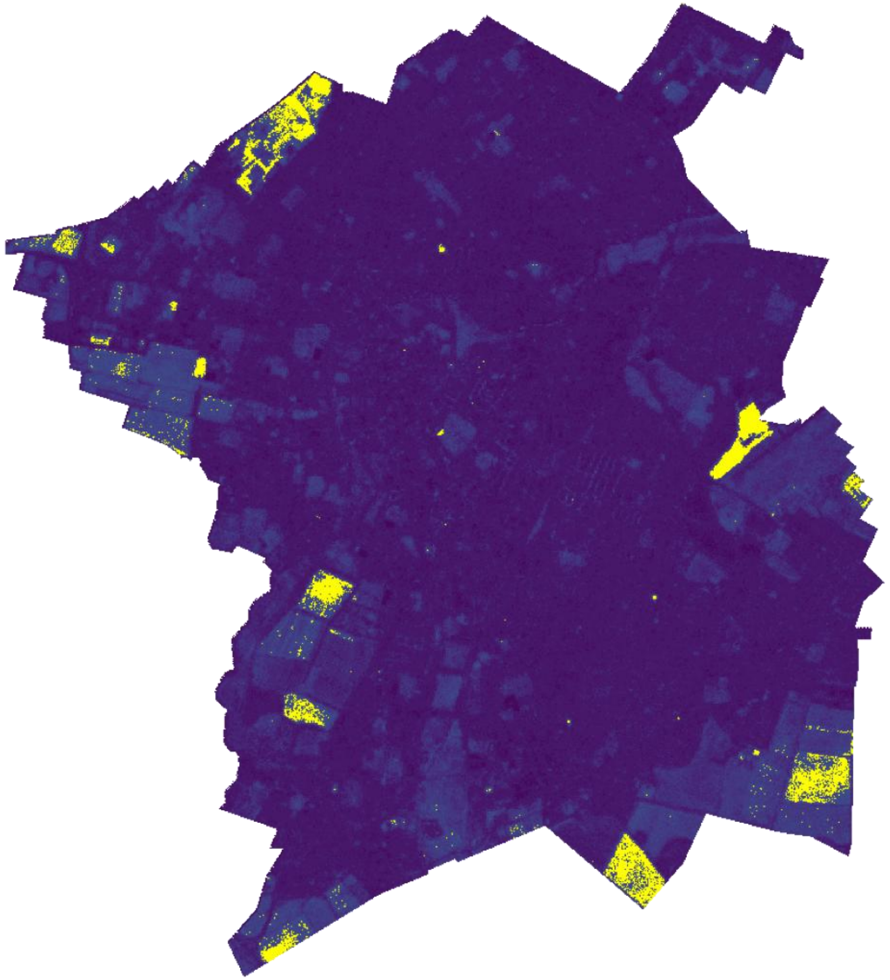

2017-2018

zf281@cam.ac.uk

Model Architecture

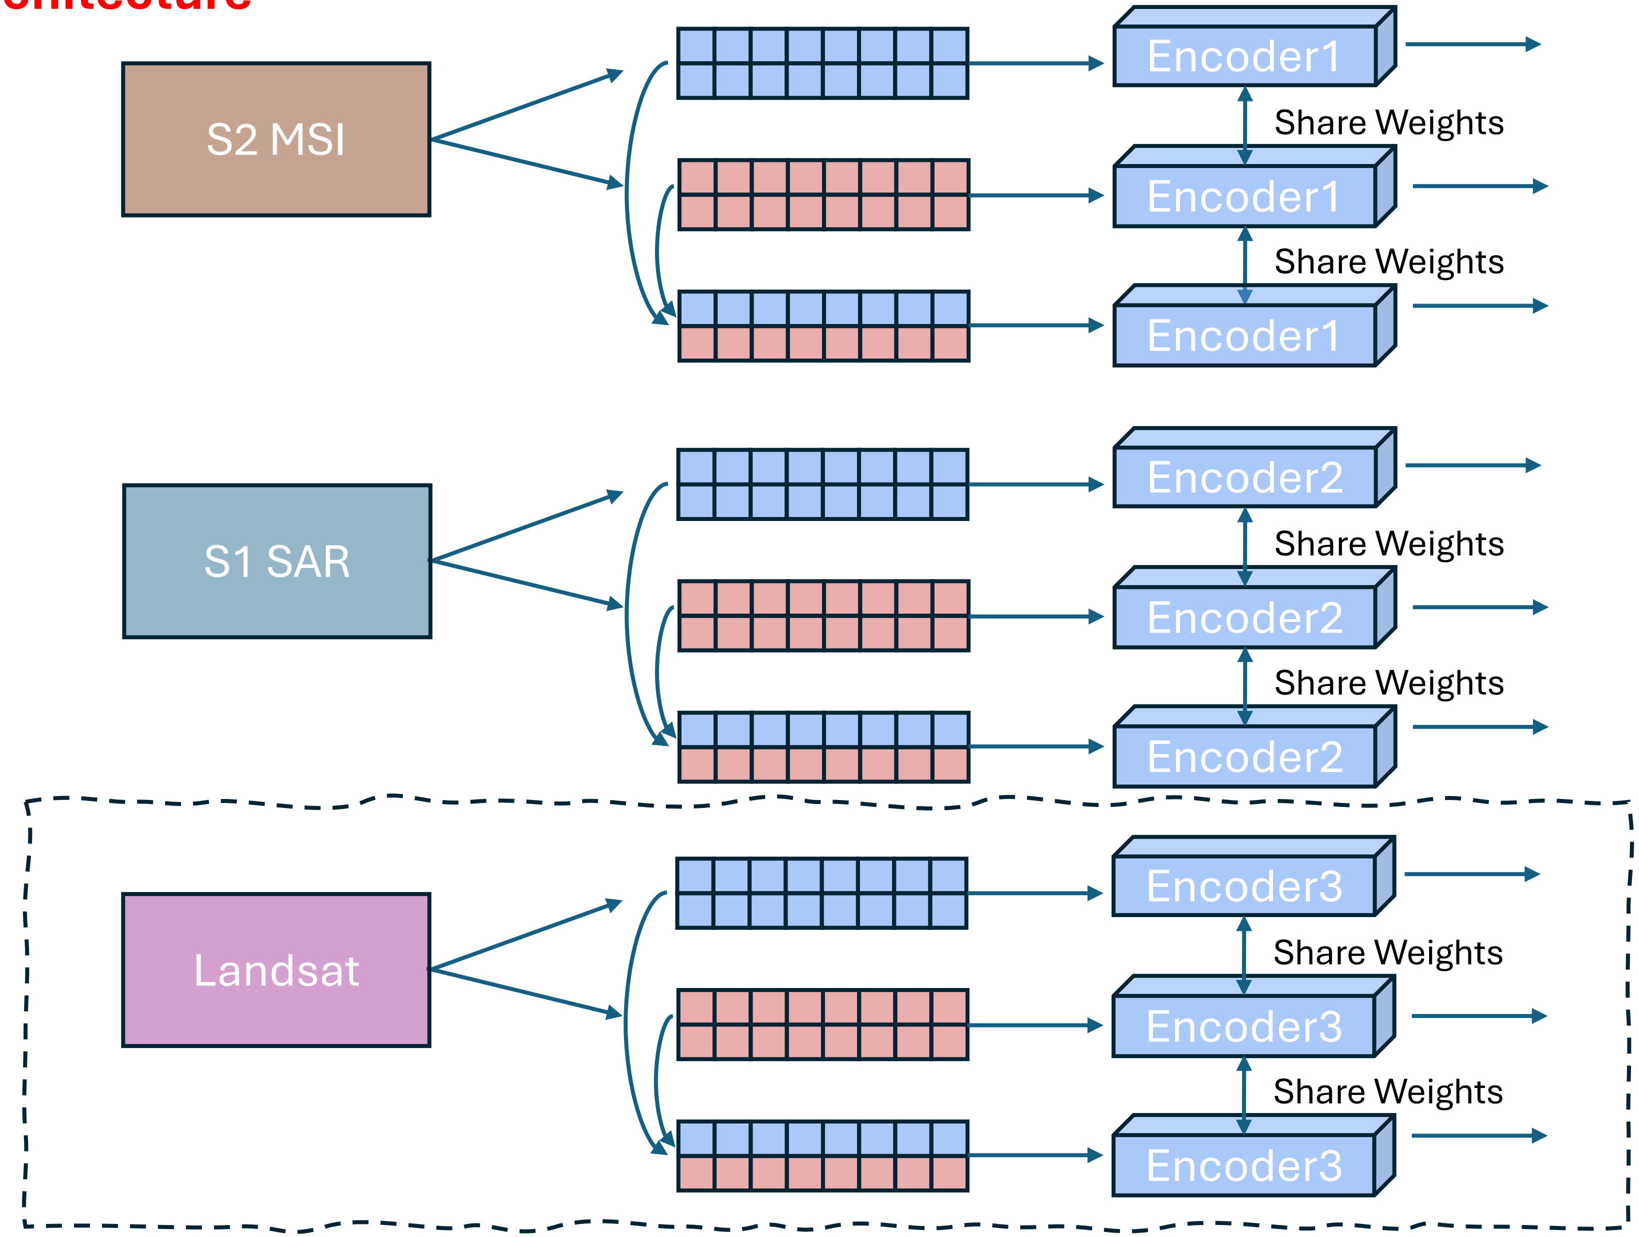

Barlow Twins

Barlow Siblings

2016-2024

1973-2024

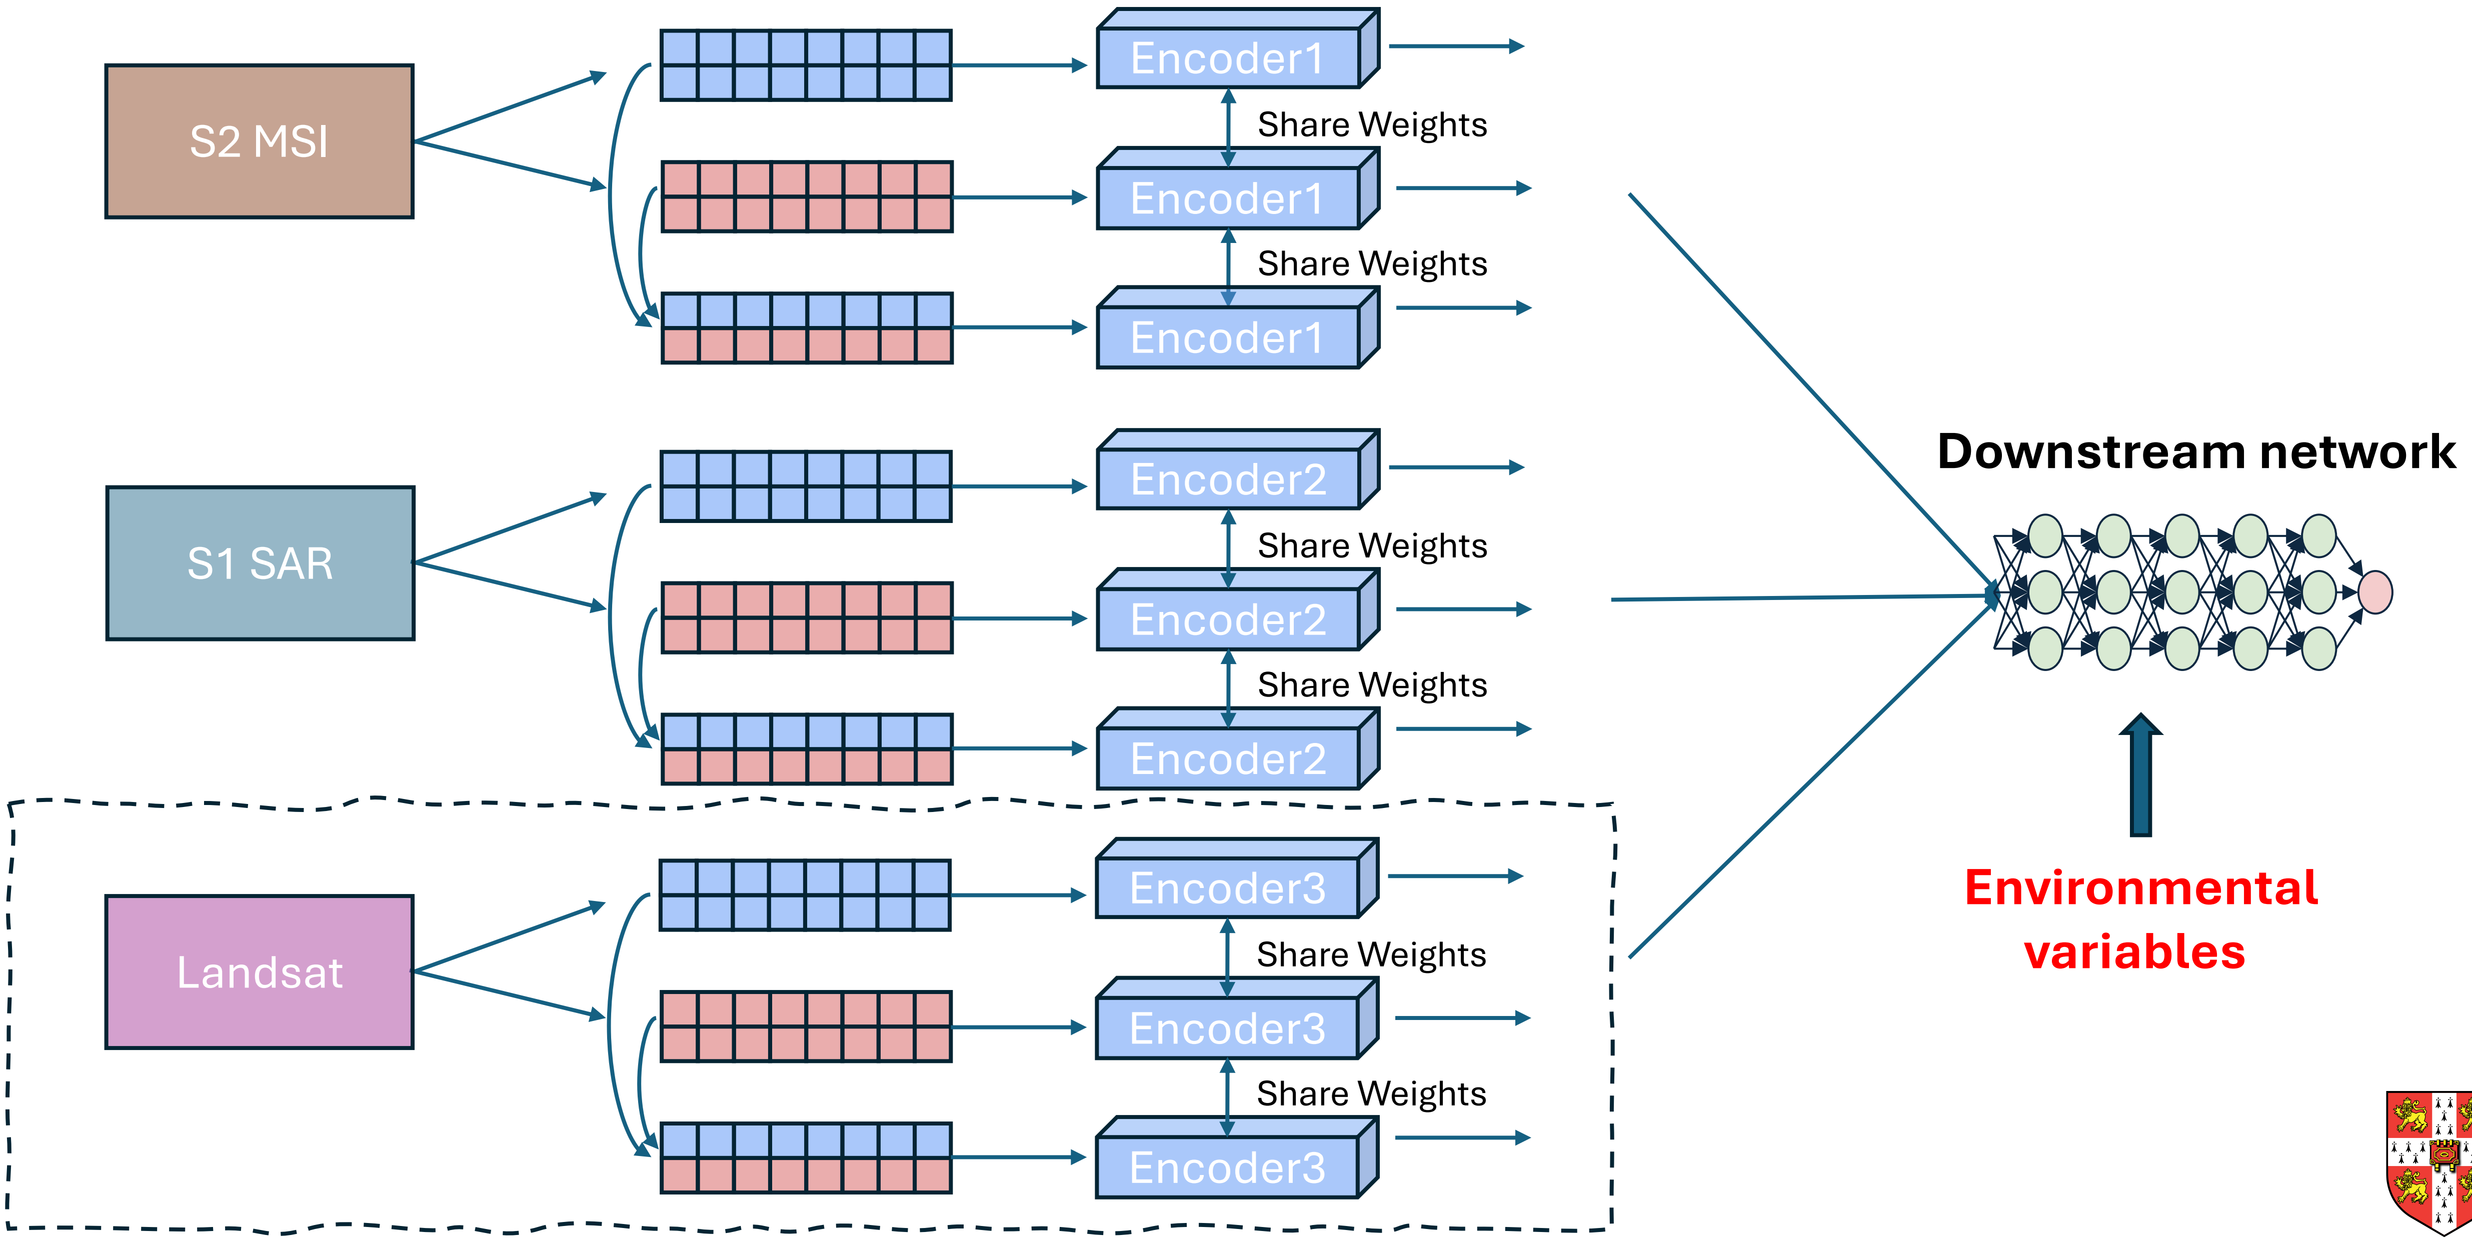
S2 MSI
S1 SAR
Landsat
Encoder1
Share Weights
Encoder1
Share Weights
Encoder1
Encoder2
Share Weights
Encoder2
Share Weights
Encoder2
Encoder3
Share Weights
Encoder3
Share Weights
Encoder3
Downstream network
Environmental variables

More downstream tasks

Land segmentation

Crop yield analysis

Counterfactual pixel matching

Using representations to understand land use change

Species distribution modelling/Robust Area of Habitat (AoH) classifiers for plant biodiversity

Comparing JRC degradation maps with the representations estimate of degradation

Global dynamic forest carbon maps

Local carbon and disturbance mapping

Global dynamic habitat maps

…

zf281@cam.ac.uk

Inter-disciplinary 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

Agriculture/Ecology

- Forest processes
- Crop behaviour

Maddy Lisaius
Clement Atberger
David Coomes

Conclusions

- SSL extracts useful insights from unlabeled EO data
- Barlow siblings approach helps improve diverse downstream tasks over SOTA
- We are seeking collaborators

zf281@cam.ac.uk

TESSERA

2017-2024 Global Representation Map

Version:0.9

Image from Google

Github Homepage

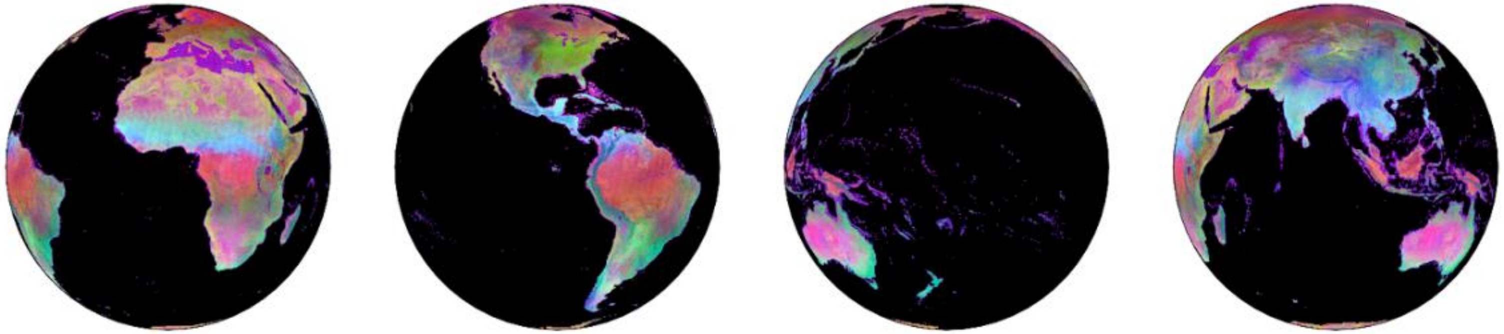

CCI Workshop

Webinar

Alpha Test

Incorporate Landsat

Early May, 2025

Jul 2025

Sep 2025

Future

```
s2_s1_global_project
├ data---------------------------------------- Folder to store raw S2 tiles
│ ├ 19FEF------------------------------------- MGRS number
│ │ └ red
│ │   └ S2A_19FEF_20220101_0_L2A.tiff
├ data_processed------------------------------ Folder to store processed S2 and S1 data
│ ├ 19FEF
│ │ ├ band_mean.npy
│ │ ├ band_std.npy
│ │ ├ bands.npy
│ │ ├ doys.npy
│ │ ├ masks.npy
│ │ ├ sar_ascending.npy
│ │ ├ sar_ascending_doy.npy
│ │ ├ sar_descending.npy
│ │ └ sar_descending_doy.npy
├ data_sar ----------------------------------- Folder to store raw S1 tiles
├ data_sar_processed ------------------------- Folder to store calibrated S1 tiff files
├ all_land_tiles.txt-------------------------- Text file including all the MGRS numbers (13672 in total)
├ Cargo.toml----------┐
├ process_tile.rs-----┤
├ Dockerfile----------┼----------------------- S2 processing program by Robin; I made minor changes
├ s2_process_tile-----┘
├ download_s2_tile.py ------------------------ Download S2 tiles
├ s1_sar_helper.py --------------------------- S1 SAR geographic alignment, downsampling, stacking, etc.
├ orchestrate.py ----------------------------- Orchestrate all the above files in parallel
```

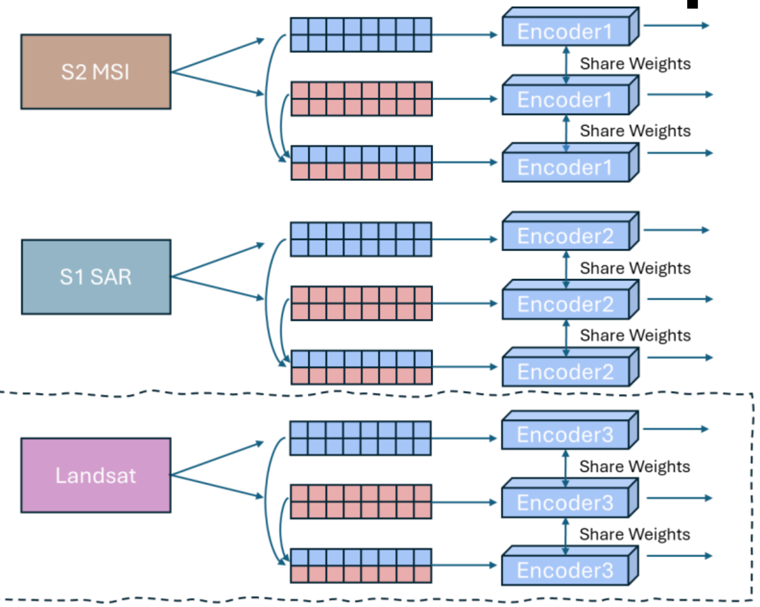

Barlow Twins

Barlow Siblings

2016-2024

1973-2024

zf281@cam.ac.uk