OpenGeoHUB
Connect • Create • Share • Repeat

Big and Useful Maps:

Land Use / land cover classification with high thematic depth

Martijn Witjes, MSc

martijn.witjes@opengeohub.org

@martijnwitjes

@MartijnWitjes

Martijn Witjes

The global mapping arms race

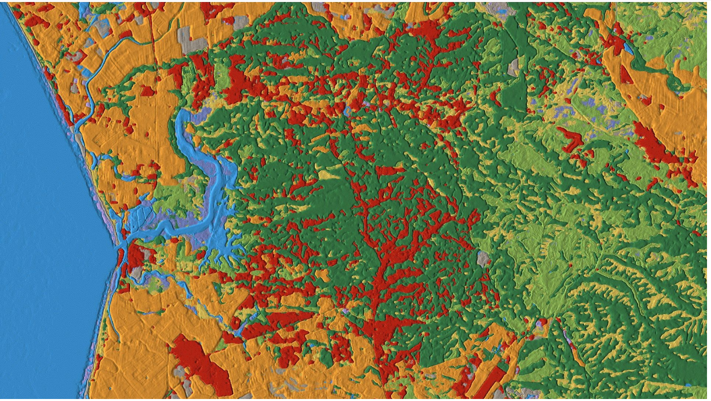

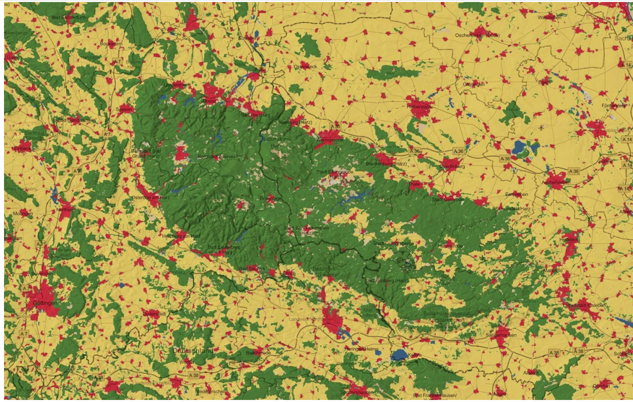

Many recent global high resolution mapping initiatives

Impressive computing

High accuracy & resolution

BUT...

Common limitations of ambitious maps

Different classification schemes / legends

1. Hard to compare **accuracy**
2. Which one to use?

Low number of classes (+- 8)

1. Limited usefulness
2. Not very innovative (other fields of ML can have many more classes)

No long-term data & uncertainty

1. Not useful for area assessment (e.g. carbon credits for forest cover)
2. Too risky to base any decisions on

Free to download, but not open

1. Hard to independently **validate**
2. Hard to **reproduce or improve**

Needs

Main gaps:

1. Comparison
2. More classes
3. Areas and trends
4. Open data & open source

Limitation:

Training data = expensive

Large legends = difficult ML problem

If only we had ImageNet for land cover....

ImageNet: 15m training images, 1000 classes, hierarchical ‘legend’

Open for anyone to use since 2006 -> **Revolutionized image recognition**

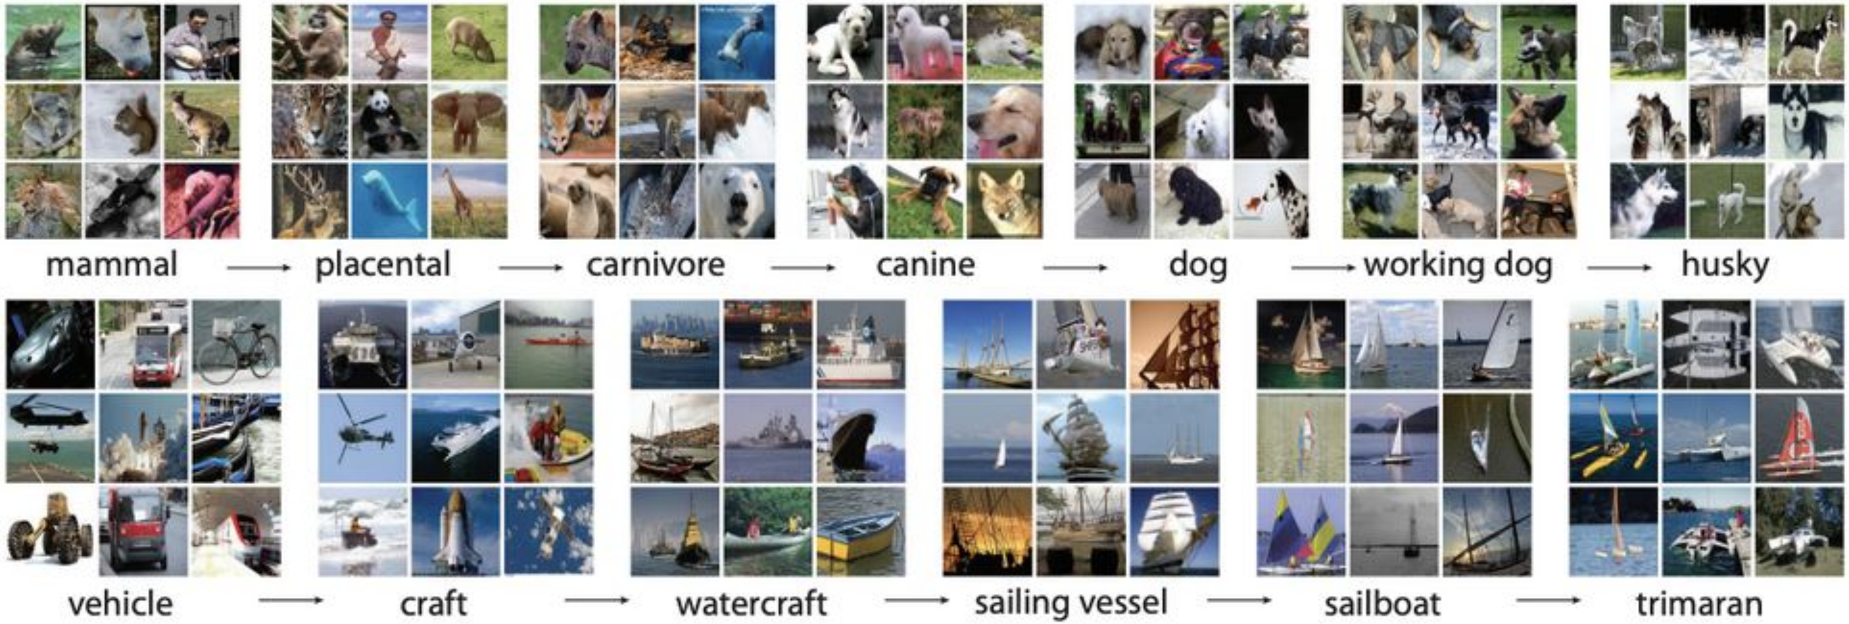

Training data

1.5 million LUCAS land cover points

5 million CORINE land cover points

Hierarchical legend:

Lvl 1: 5 classes

Lvl 2: 14 classes

Lvl 3: 43 classes

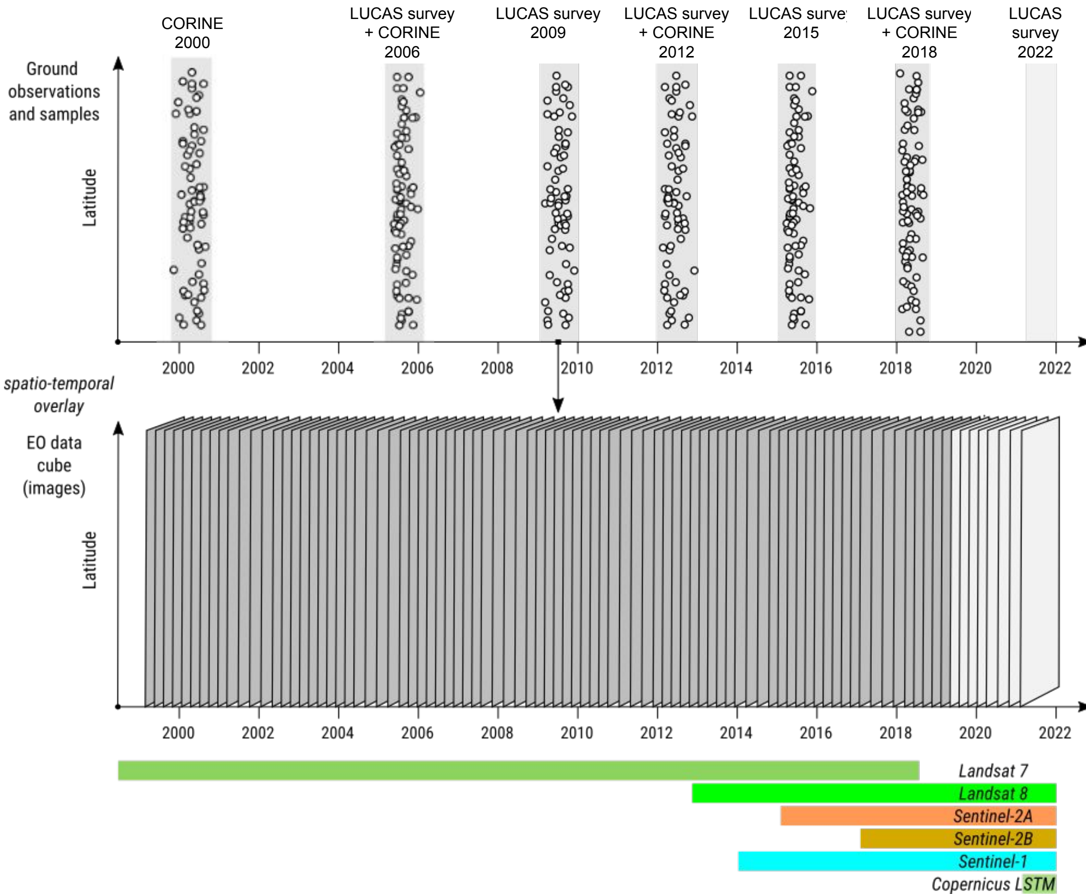

Features

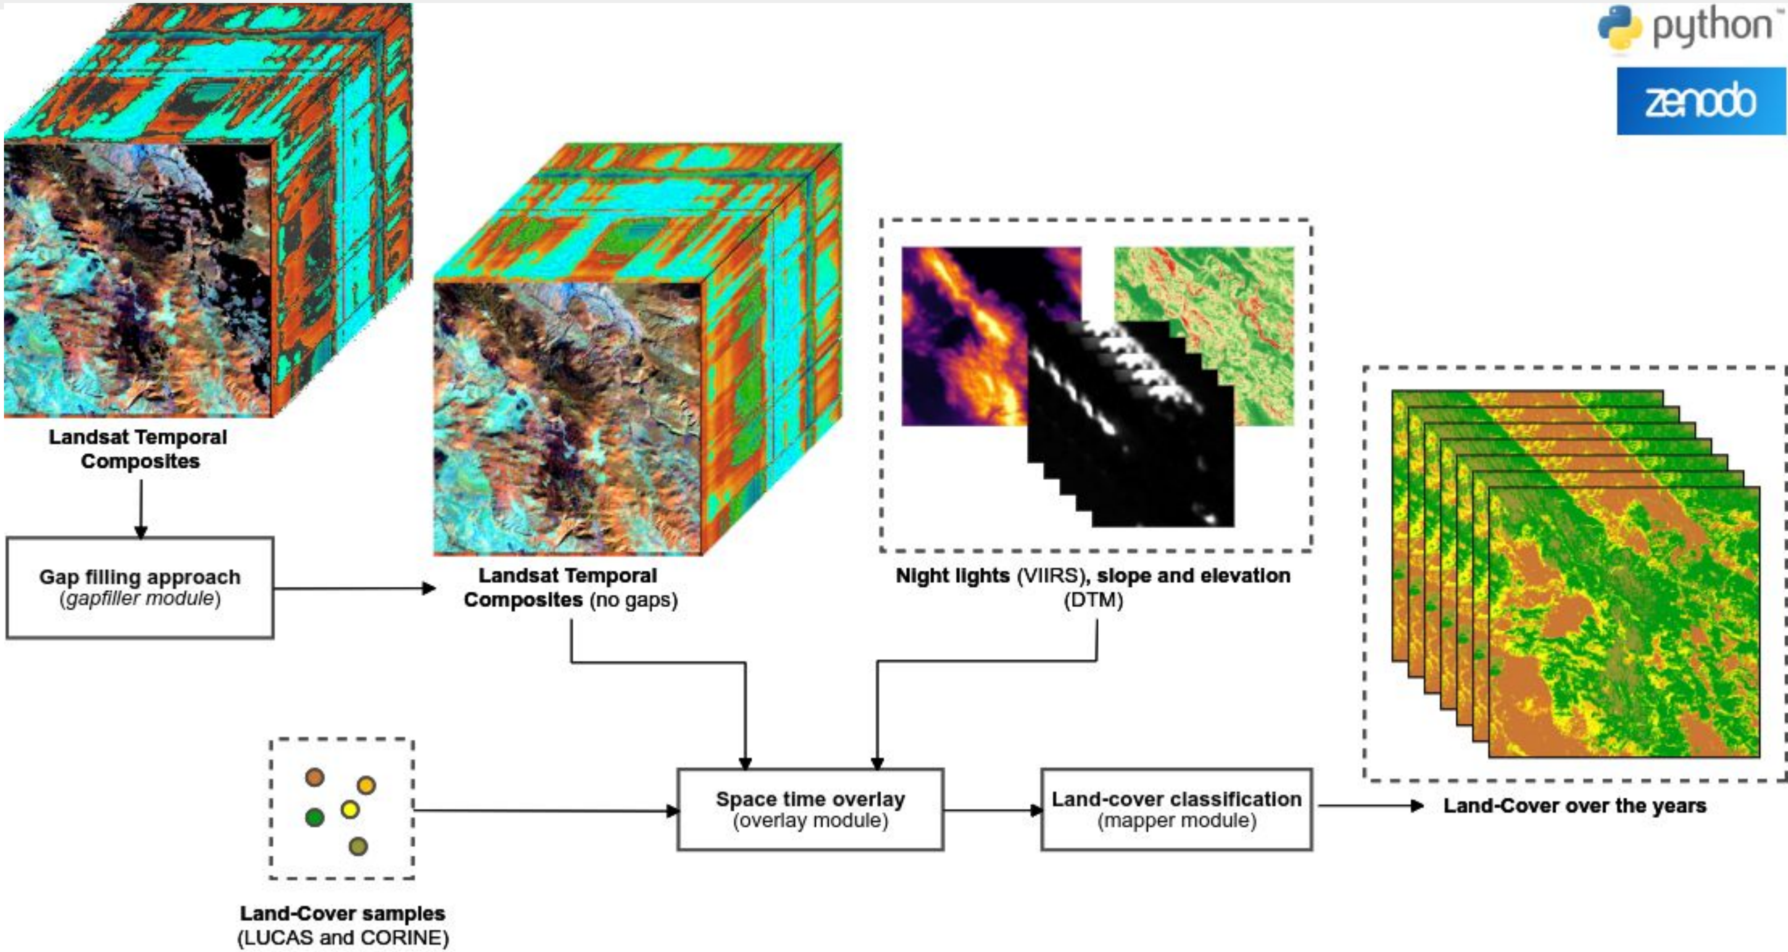
Landsat Temporal Composites
Gap filling approach (gapfiller module)
Landsat Temporal Composites (no gaps)
Night lights (VIIRS), slope and elevation (DTM)
Land-Cover samples (LUCAS and CORINE)
Space time overlay (overlay module)
Land-cover classification (mapper module)
Land-Cover over the years

Workflow

OpenGeoHUB
Connect • Create • Share • Repeat

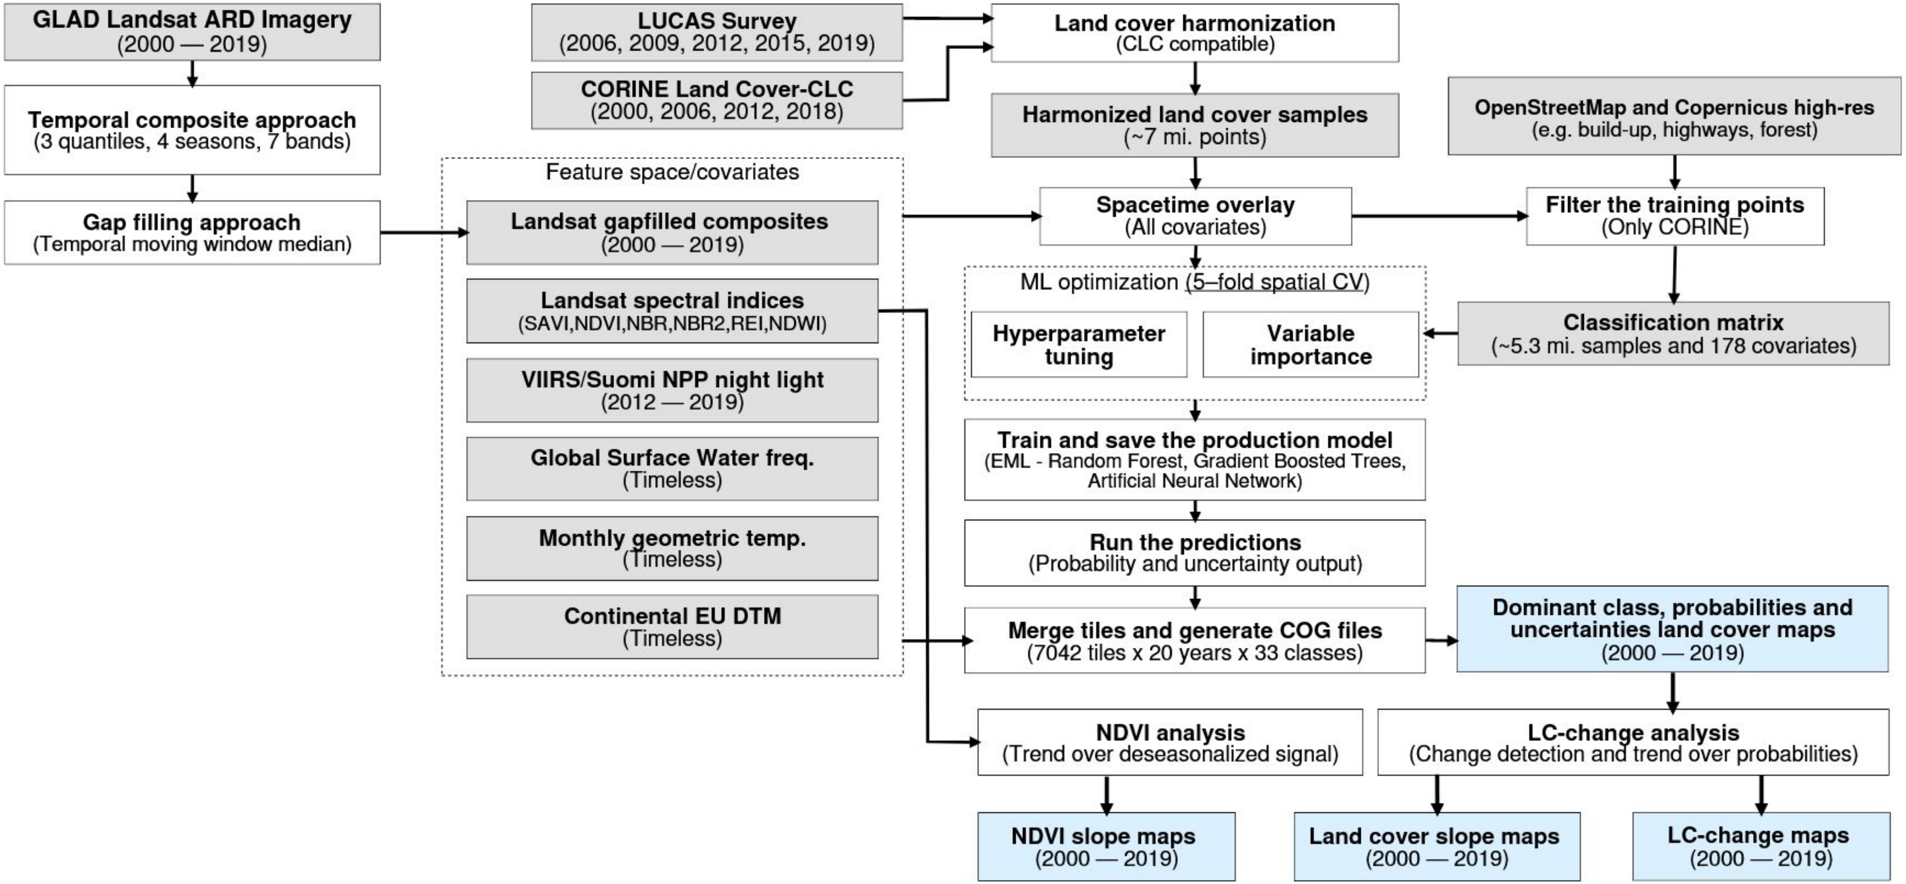

Legend-agnostic performance metric

OpenGeoHUB
Connect • Create • Share • Repeat

More classes = more difficult to classify.

But is that really lower performance?

Most performance metrics: *‘Yes’*

Baseline Log Loss = How much better than a random guess, given X number of classes?

1 - (model / random)

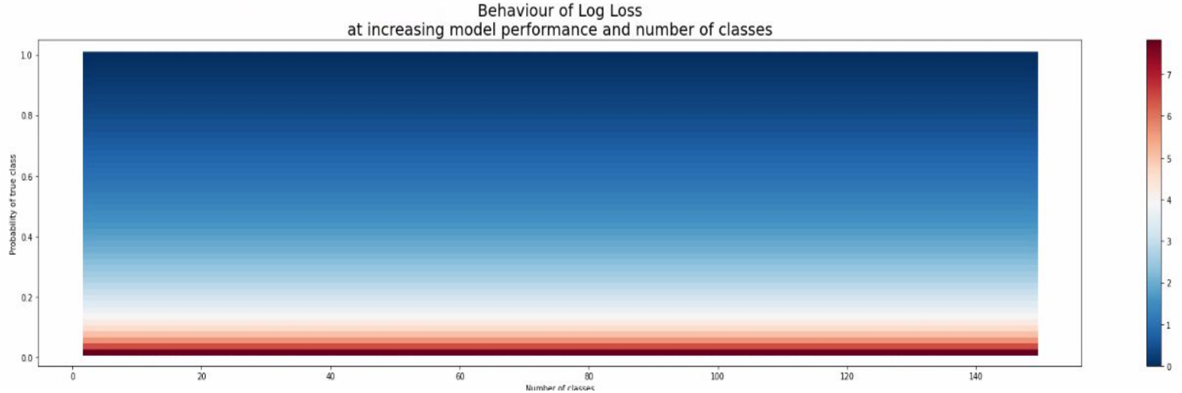

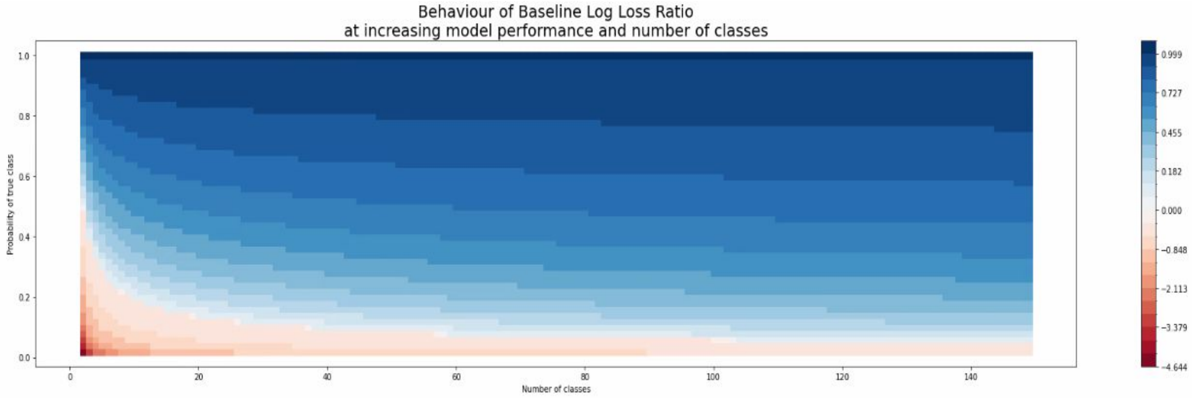

Results

Lvl	Classes	F1	Baseline Log Loss
1	5	0.83	0.77
2	14	0.63	0.71
3	43	0.49	0.70

σ F1: time: 0.135; space: 0.15

Distribution of training data

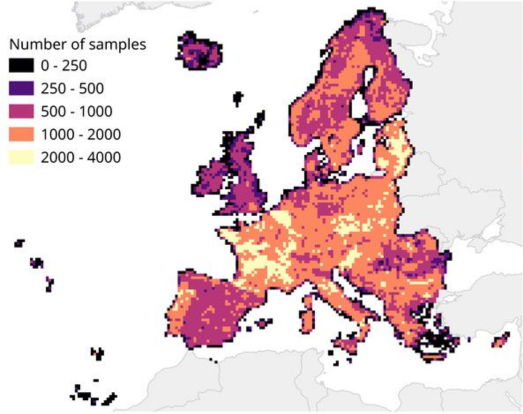

Localized accuracy

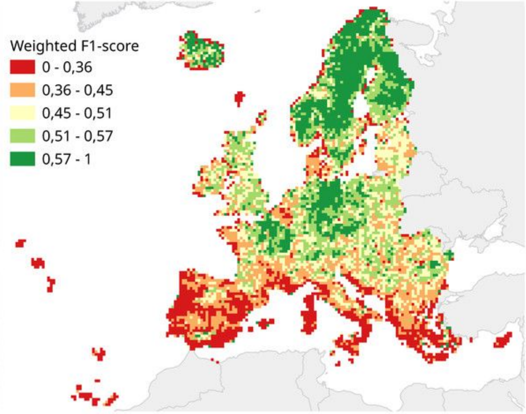

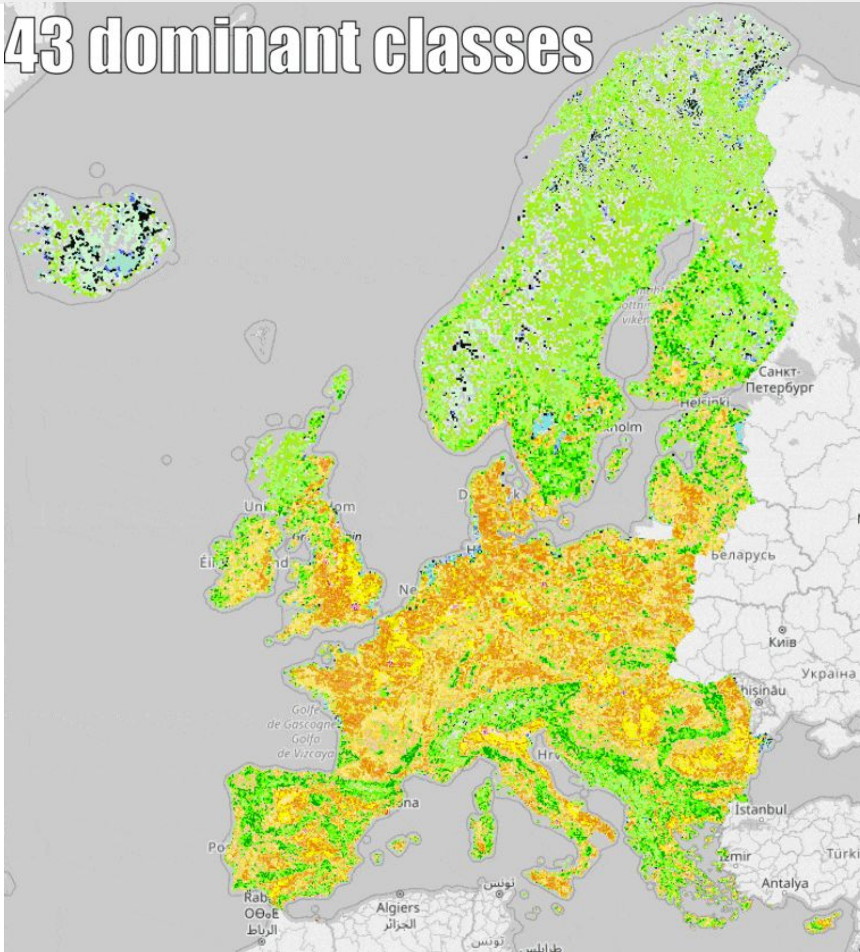

30-m Land use / land cover mapping

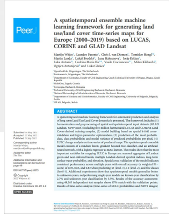

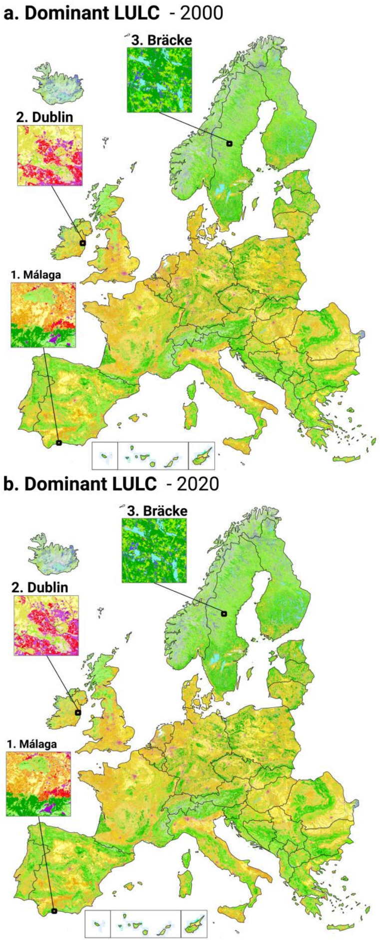

43 CLC classes (2000—2020)

https://doi.org/10.7717/peerj.13573

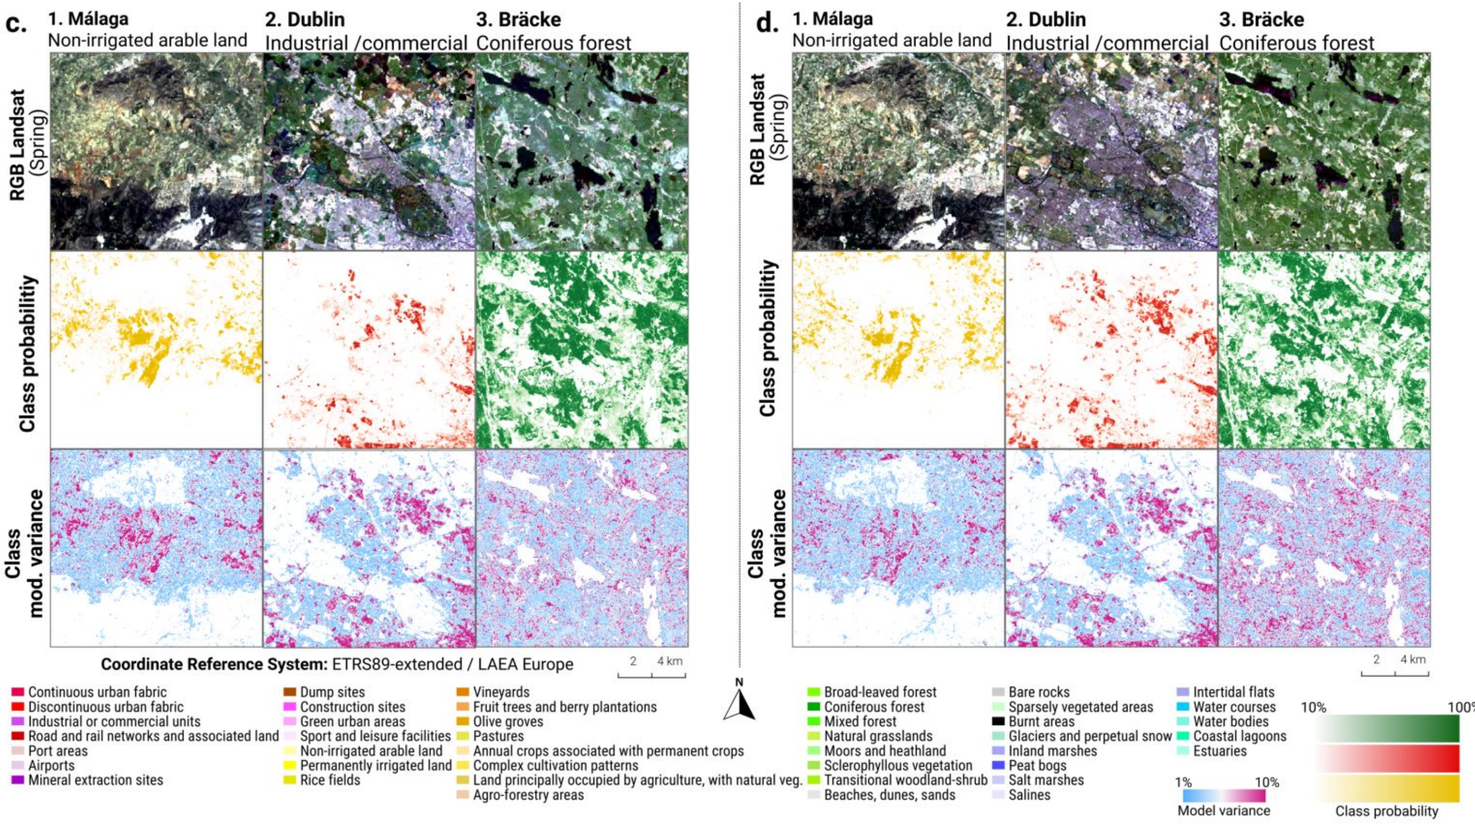

Next steps

OpenGeoHUB
Connect • Create • Share • Repeat

- Investigate **Baseline log loss**
  - Theoretical underpinnings
  - **Actual performance loss** when training on more classes
  - **Objectively compare** high-profile maps

- Investigate **prediction uncertainty:**
  - How likely is the model correct **where and when?**
  - Combine with time-series to determine **land cover area change**

- Investigate **hierarchical legends:**
  - How to **leverage for performance gain**
  - How to design them for **humans and algorithms**

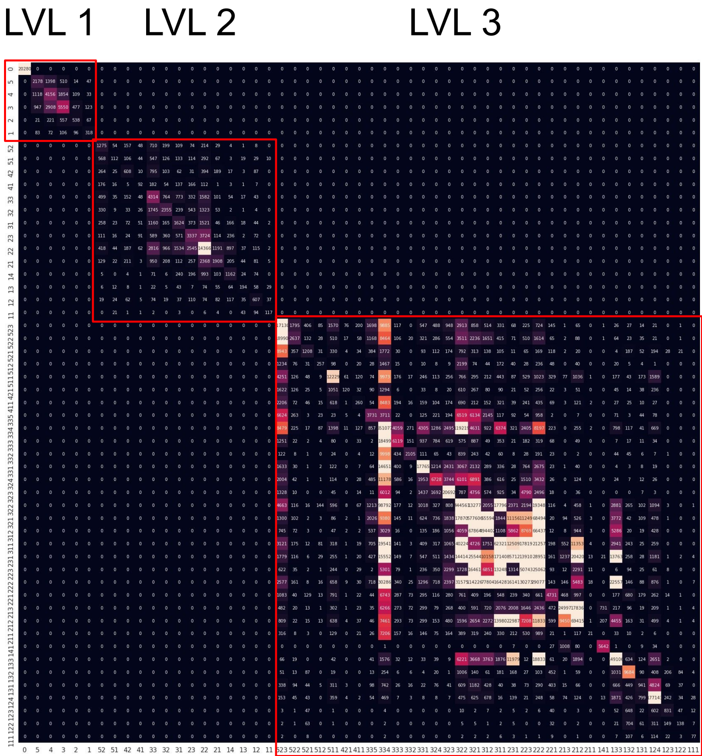

Hit me up! martijn.witjes@opengeohub.org