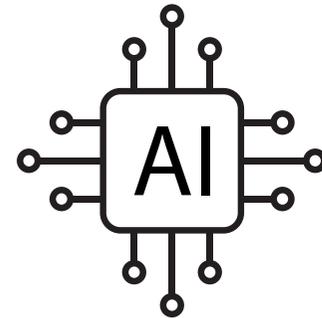
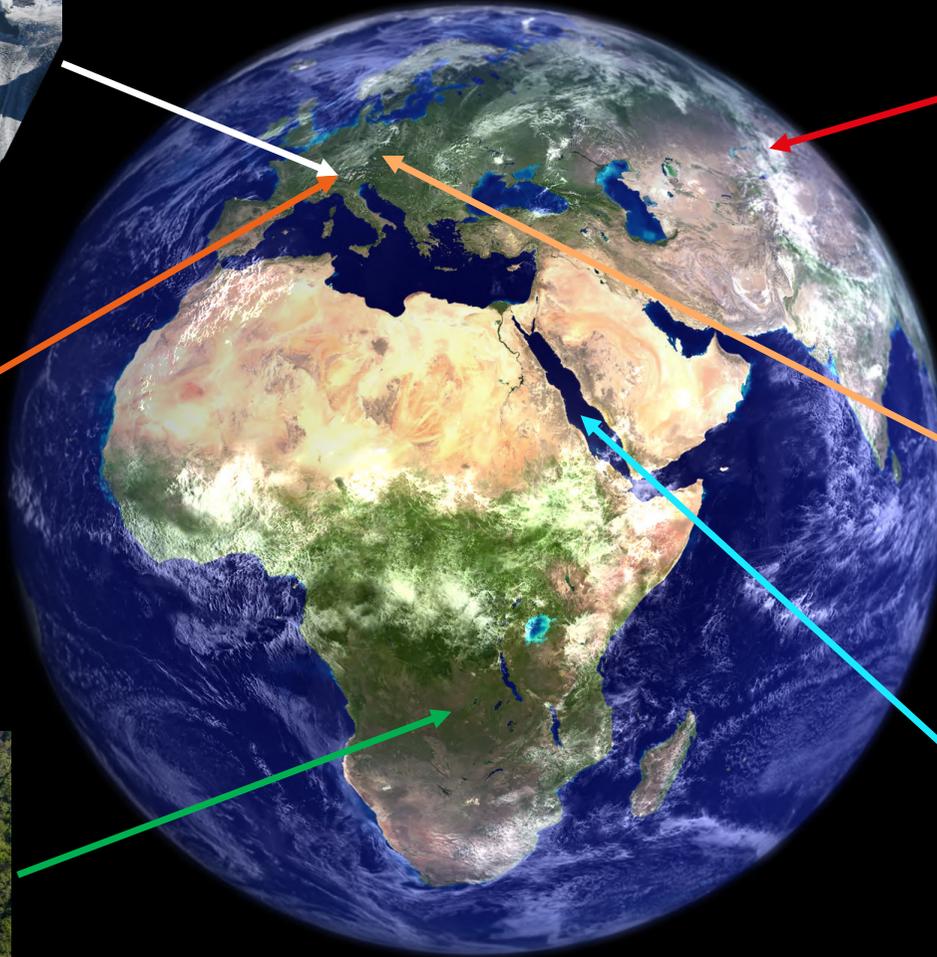
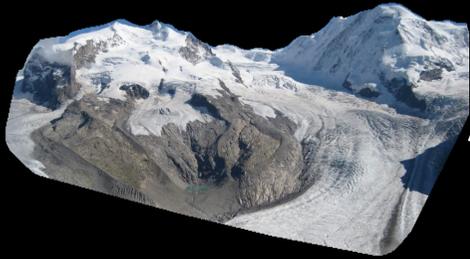


Machine learning for Earth: monitoring the pulse of our Planet with sensor data, from your phone all the way to space

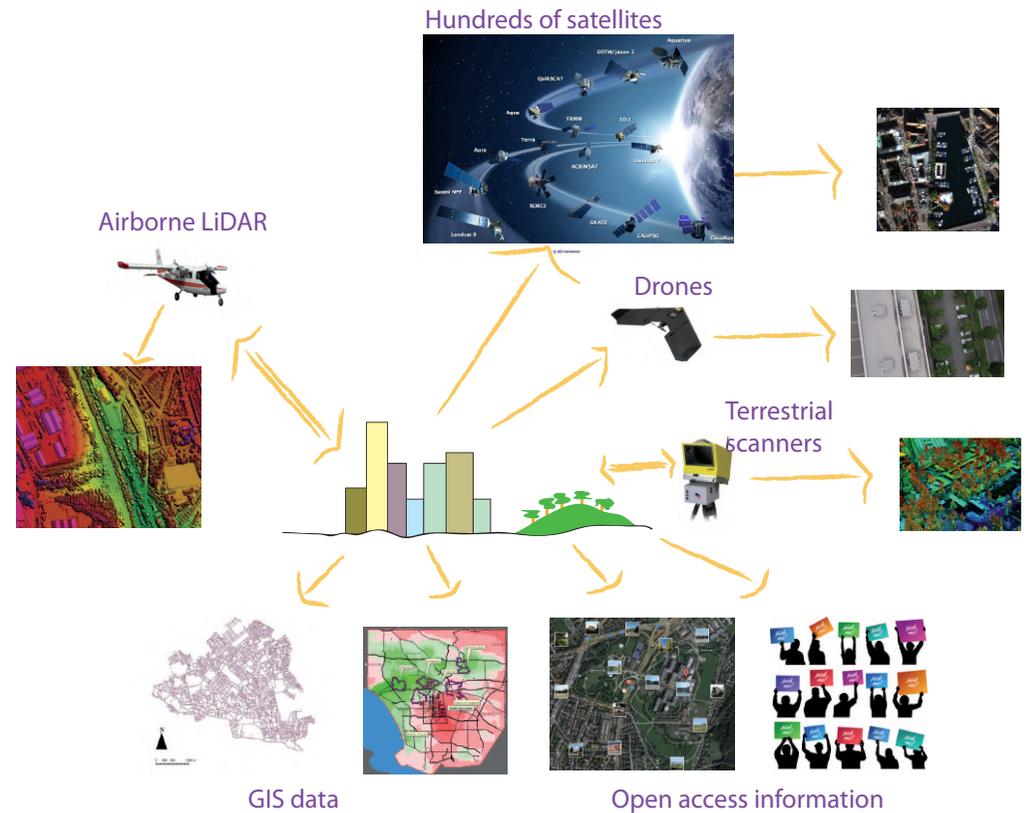
Prof. Devis Tuia, EPFL





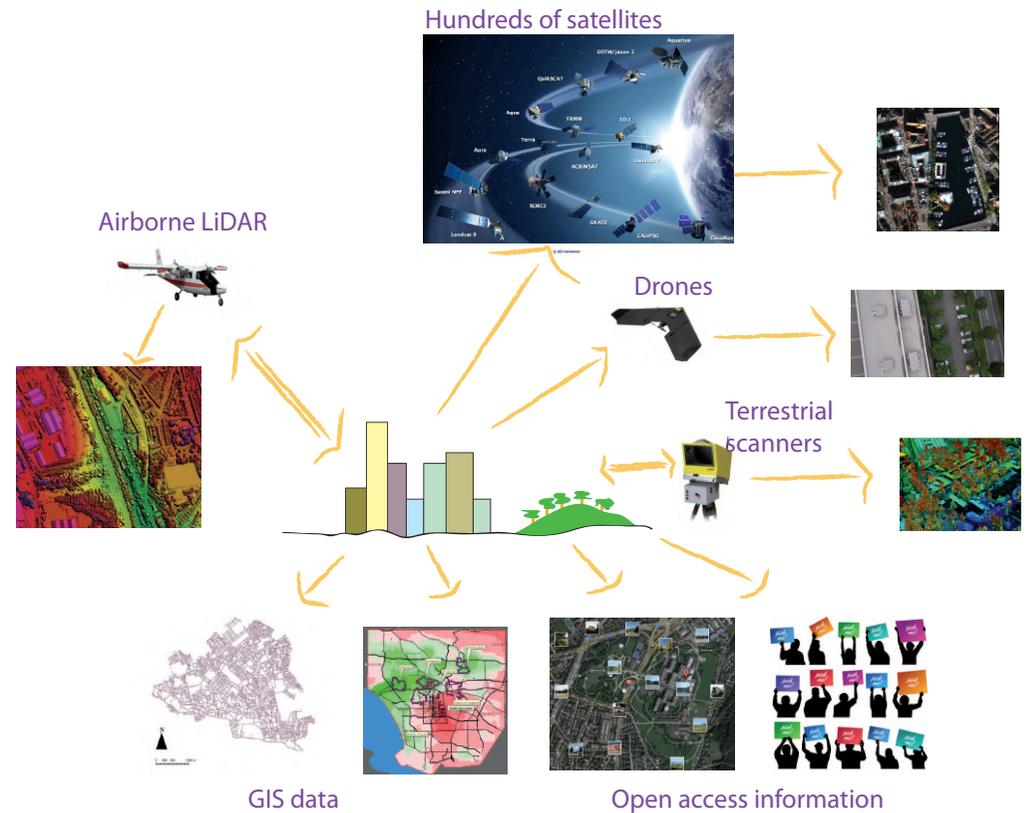
There were many sensor data to monitor Earth in 2015

- 333 Earth Observation satellites in orbit in 2015 [ucsusa.org].
- 10'000 recreational drones registered in the U.S. by 2020 [FAA].
- 20 Pb of oblique photos in Google Street View in 2015 [Google Maps].

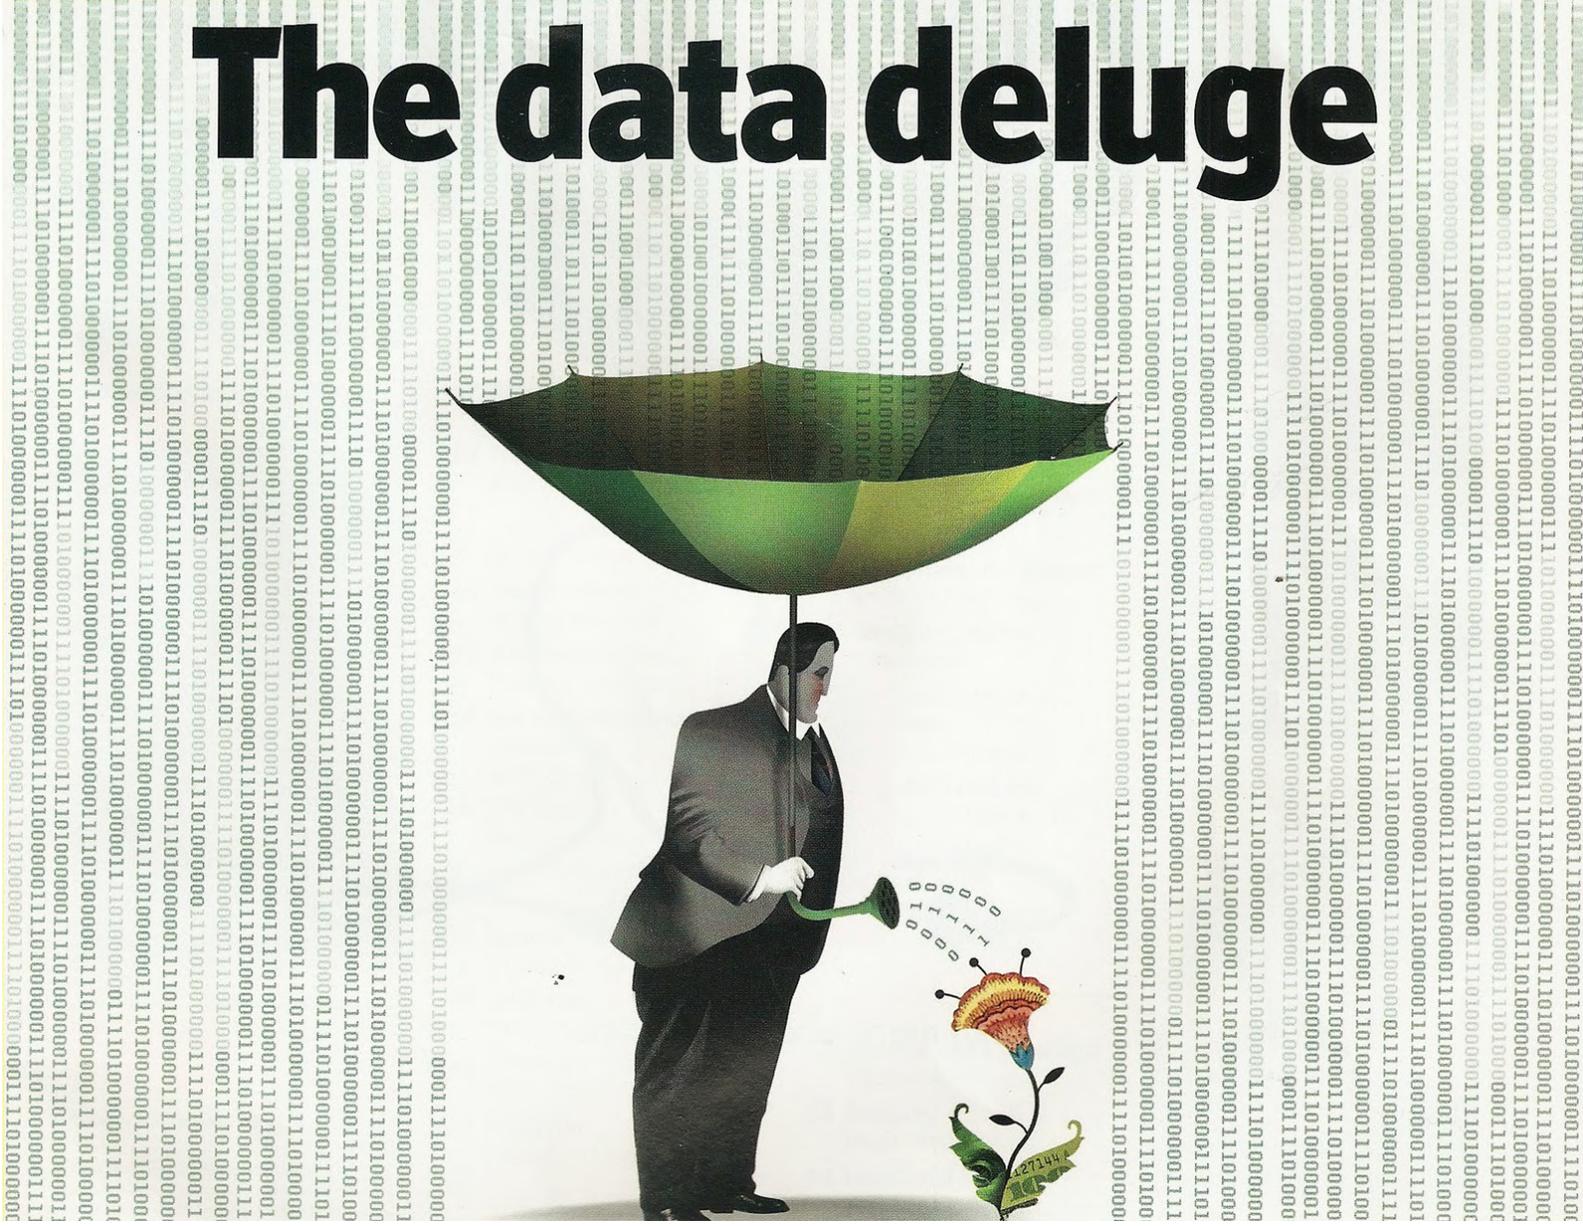


There are many sensor data to monitor Earth in ~~2015~~ 2024

- 333 **1'005** Earth Observation satellites in orbit in 2023 [ucsusa.org].
- 10'000–**1'100'000** recreational drones registered in the U.S in 2023. [FAA].
- **170 billions of oblique photos** in Google Street View in 2020 [Google Maps].



The data deluge



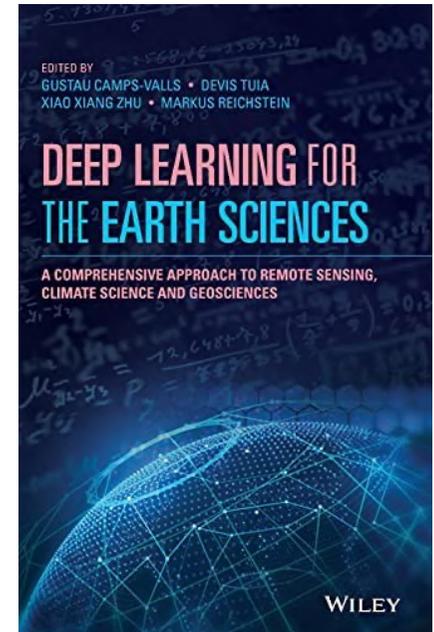
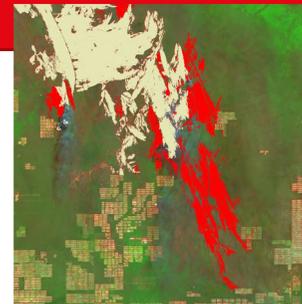
Why now : statistical and computational models are good enough...

- Machine learning has reached a certain maturity... and percolated in many fields of science.

2022

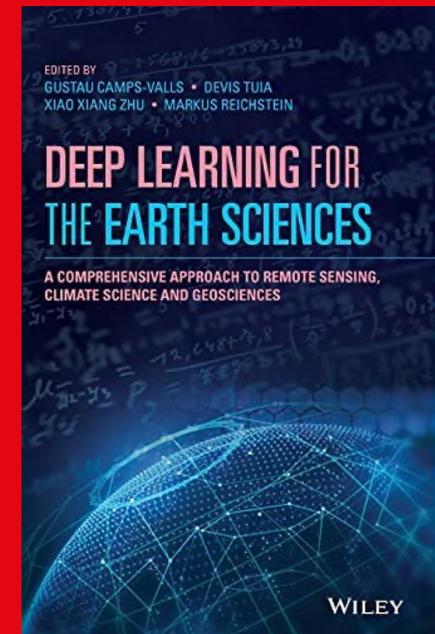


2015



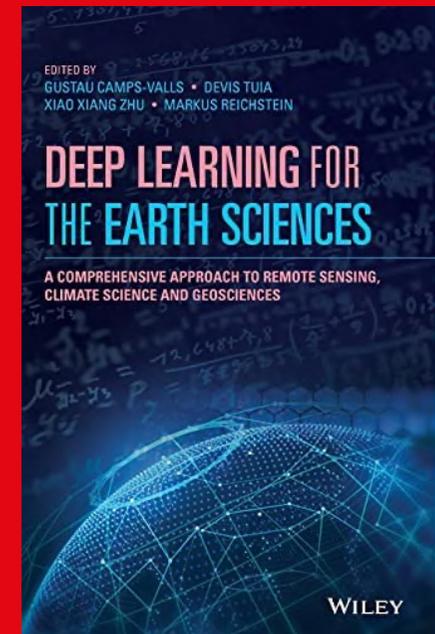
With Earth observation and AI, we can

develop computational approaches
to the environmental sciences
that are **accurate**



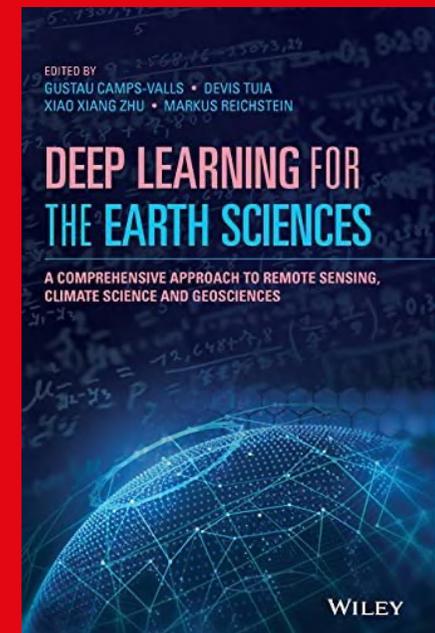
With Earth observation and AI, we can

develop computational approaches
to the environmental sciences
that are **accurate**



With Earth observation and AI, we can

develop computational approaches
to the environmental sciences
that are **accurate**, but also
scalable,
knowledge-based and
accessible to everyone.





Towards environmental machine learning that is

Scalable

Knowledge-driven

Accessible to anyone

How many species of trees are there on Earth?

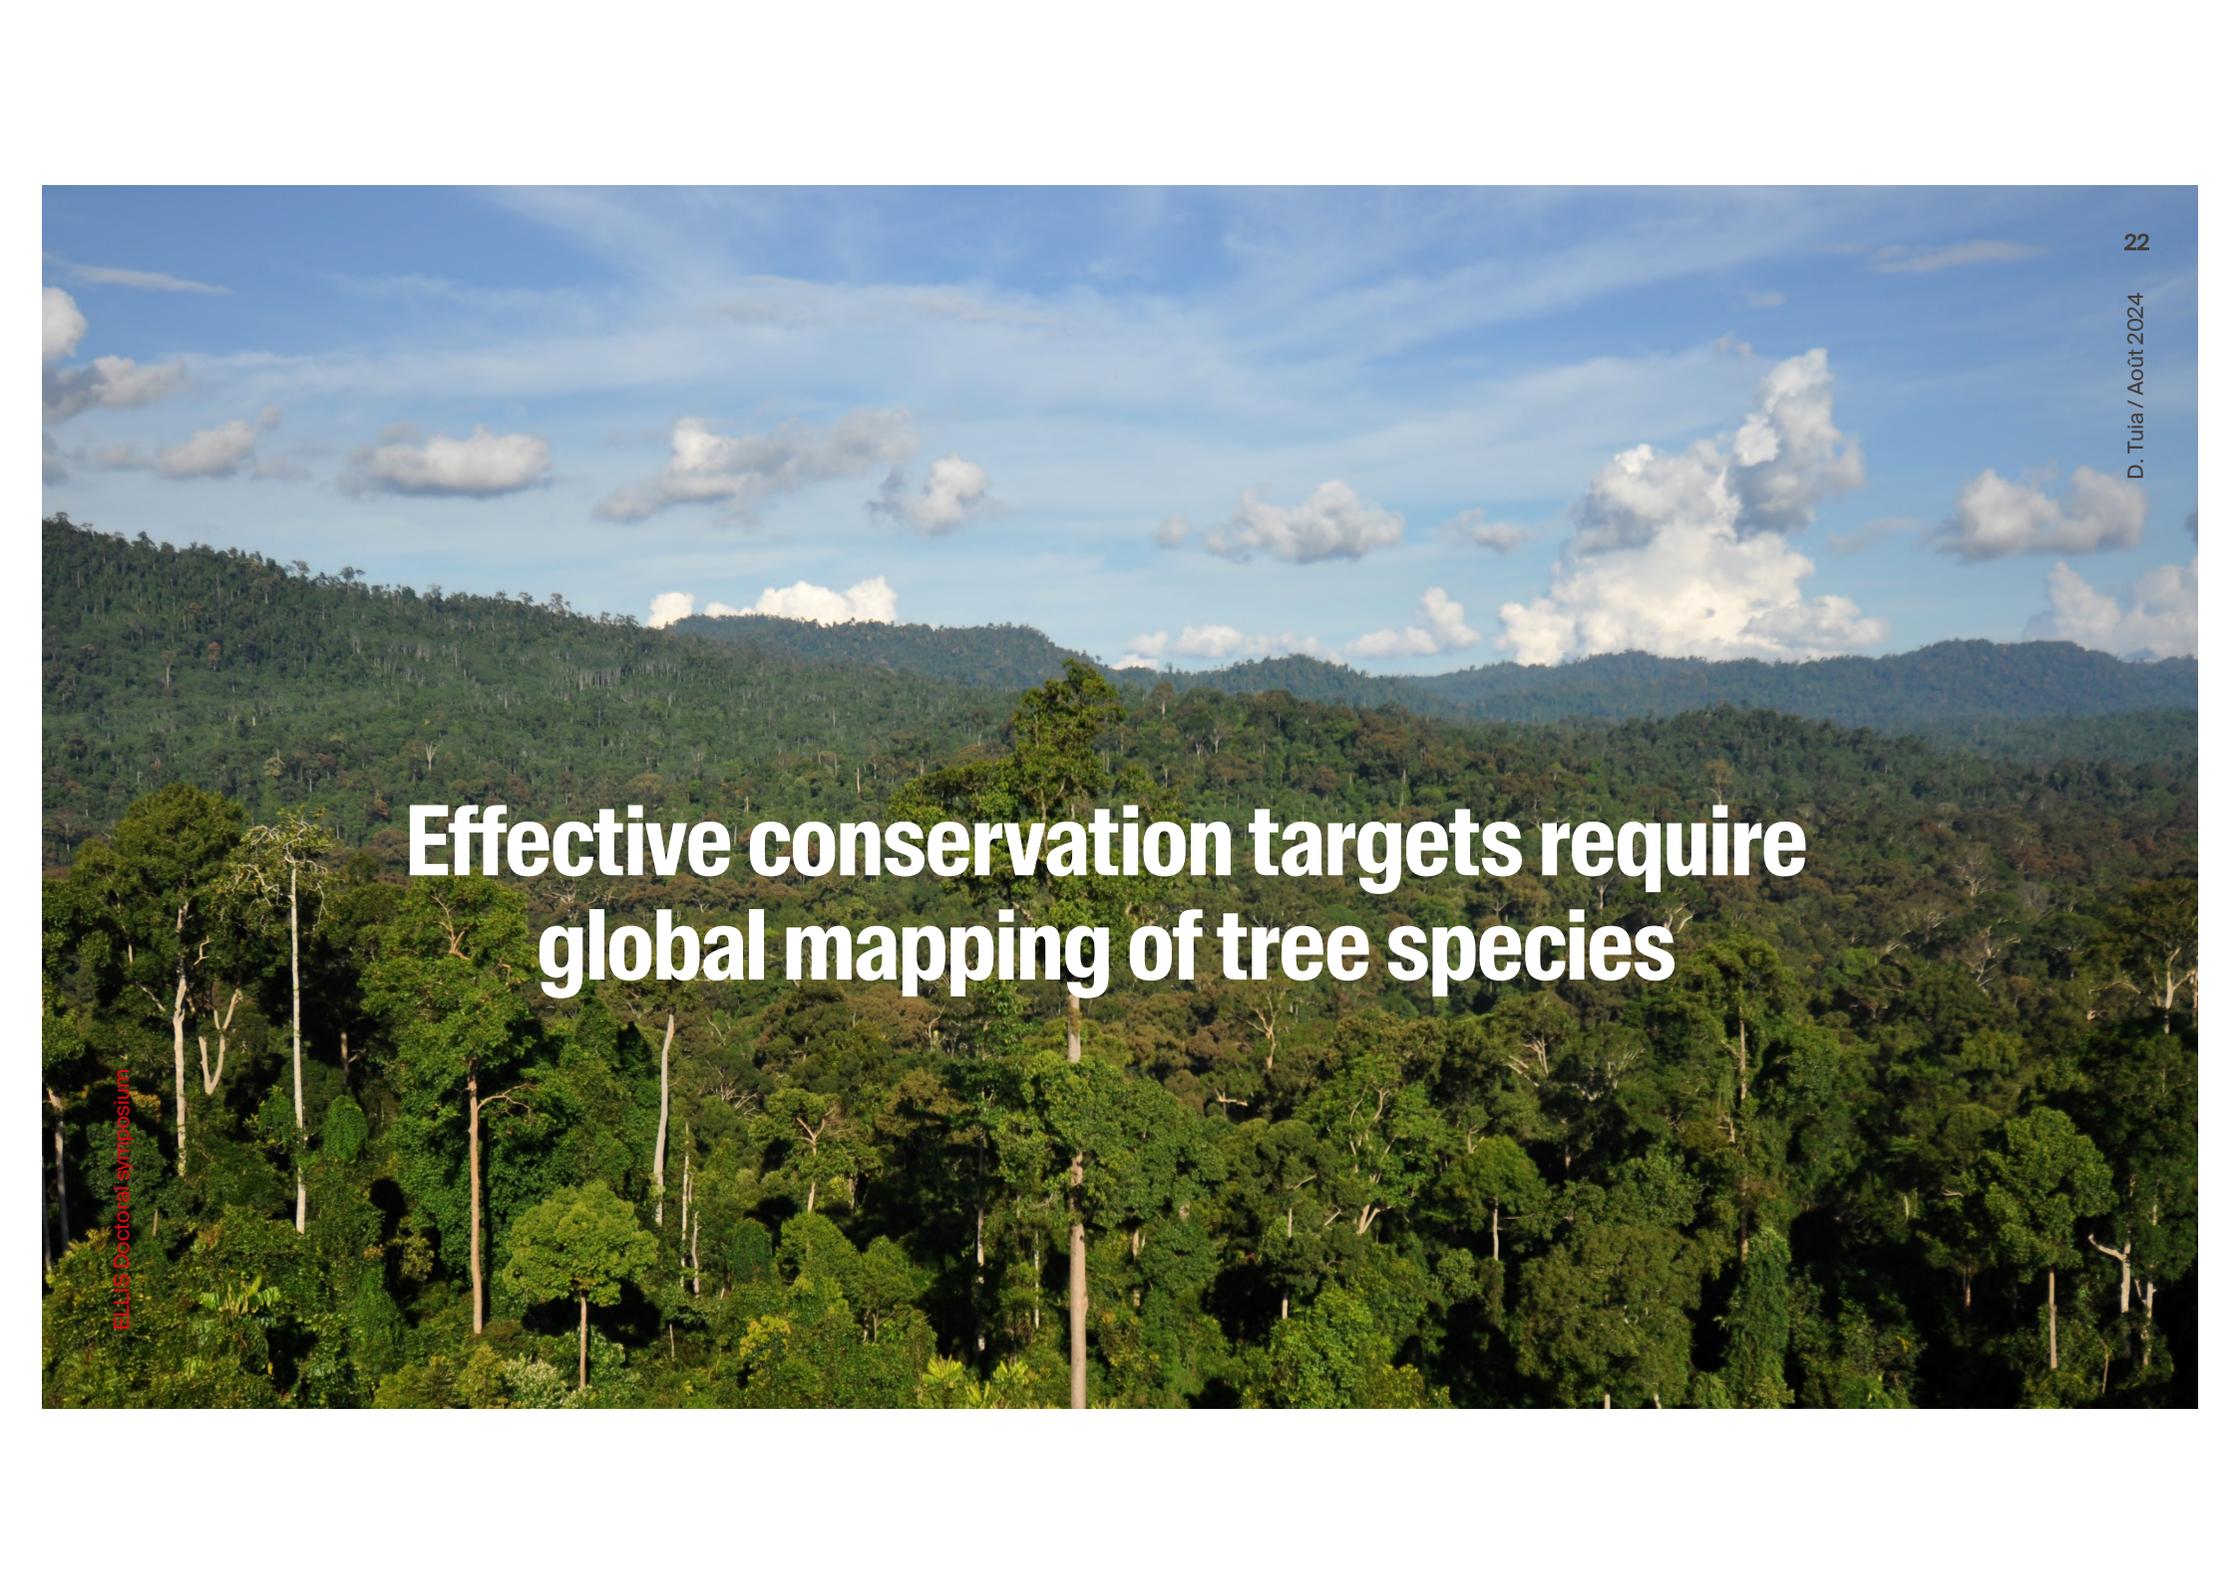
- More or less 73'000

Where are they?

- We don't know exactly. But we have field measurements for today's situation

Where will they be?

- It will depend on how the climate evolves. But we know that 30% of them are at risk of extinction.



**Effective conservation targets require
global mapping of tree species**

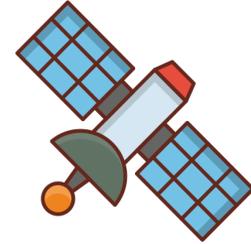
Species distribution models (1): species data



- We can't sample the whole world
- We need to use data from field campaigns, stored in 13 databases
 - 30 millions observations
- These tell us about where the species has been observed (one-class problem).

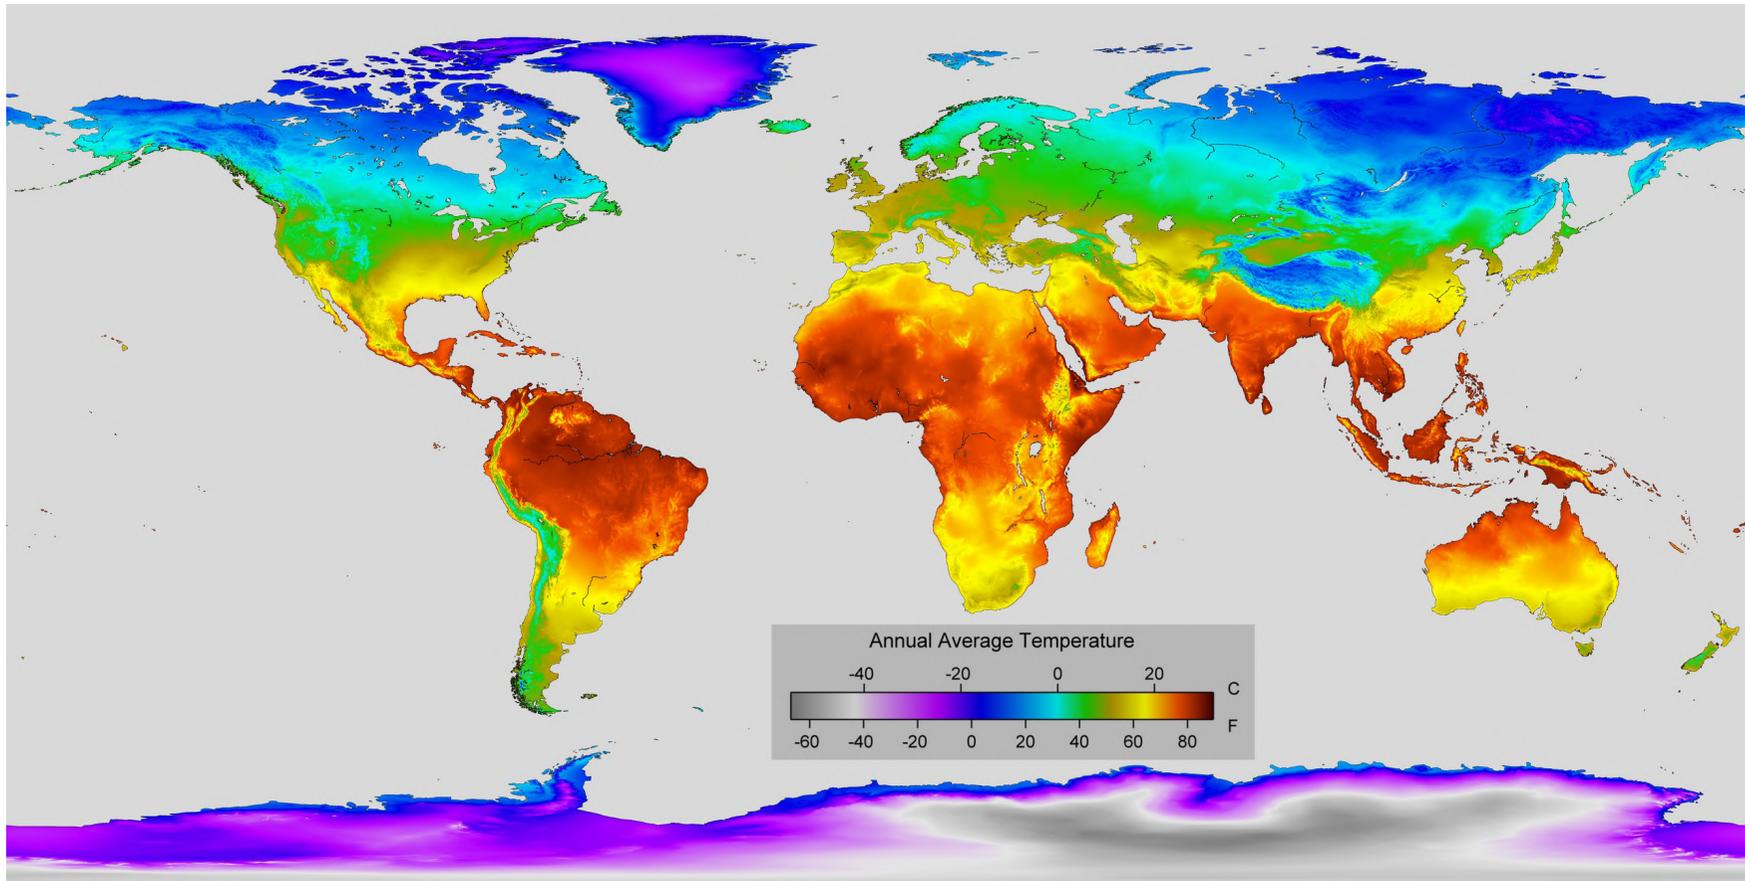
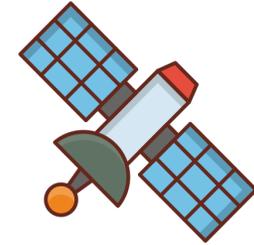
- We have 10'590 species with 90+ observations.
- For each species
 - We used training data in a buffer 1'000km around the native range
 - We sampled 5'000 pseudo-absences samples

Species distribution models (2): environmental predictors



- We use a $\sim 1\text{km}^2$ (30 arc sec) grid, global
- As predictors we used
 - Climate variables (temperature, precipitation, growing season, ..) from the CHELSEA V2.1 data
 - Soil variables (Soil Ph, content, ...) from SoilGrids.

Species distribution models (2): environmental predictors



Species distribution models (3): the model

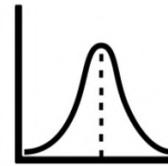
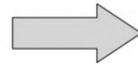
- Ensemble of ML models (random forests and gradient tree boost)

	x^1	x^2	...
location	temp. [°C]	elev. [m]	prec. [mm/a]
A1	16	128	4320
A2	18	134	4560
B1	11	1950	1381
...			

point-wise predictors
(climate, elevation, *etc.*)

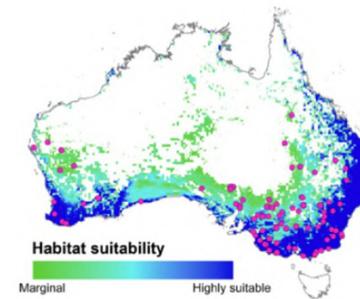
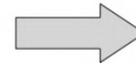


species occurrence records
(field sampling,
museum collections, *etc.*)



$$y = f(x^1, x^2, \dots)$$

Species Distribution Model
(generalised linear model,
random forest,
Bayesian model, *etc.*)



habitat suitability
(prediction, extrapolation)

Where are they?

- 10'590 species maps

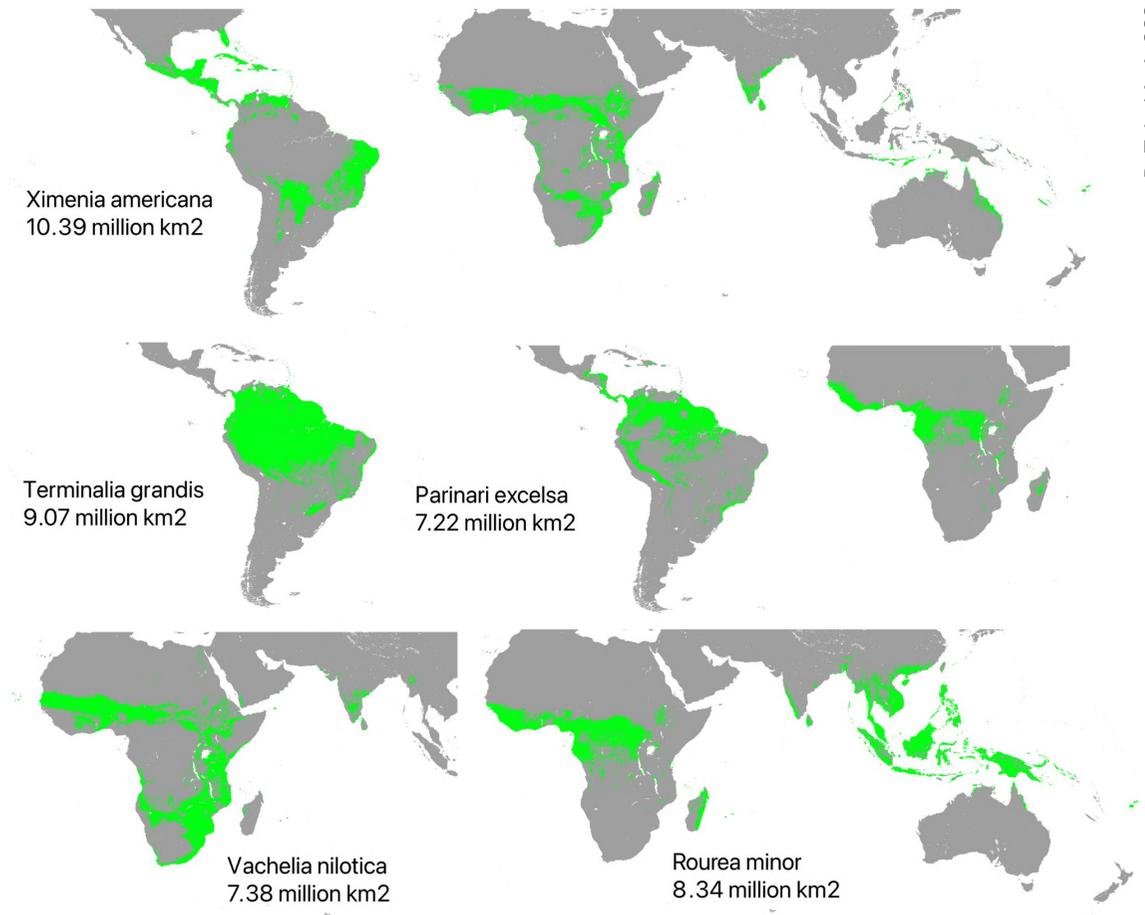
<https://doi.org/10.5281/zenodo.10911892>

- Xval scores:

- True Skill Statistics: 0.77
- AUC : 0.93

- Independent validation

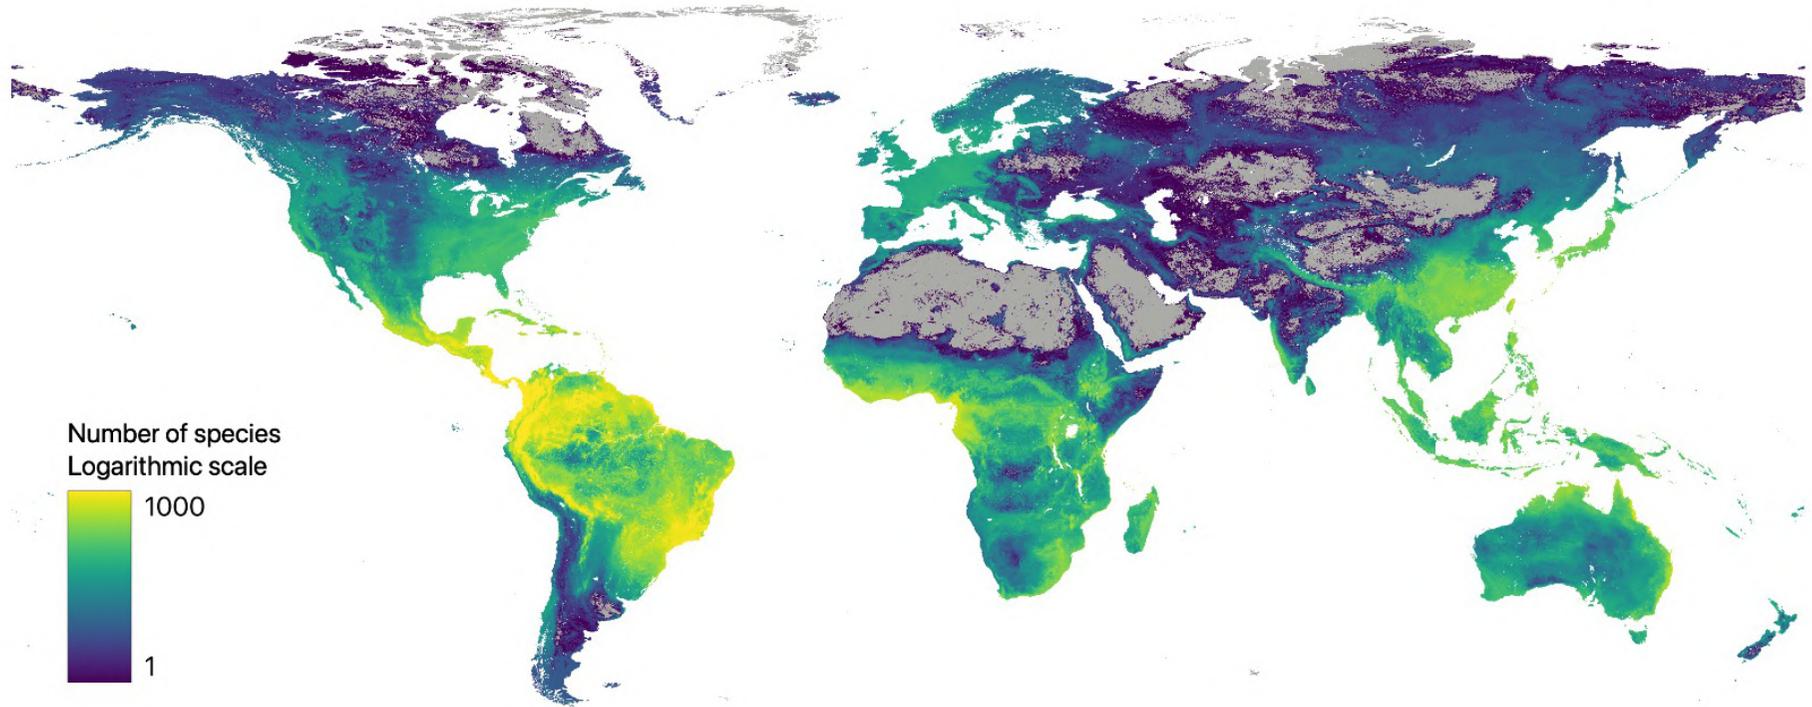
- TSS: 0.53



N. van Tiel, L. Lyu, F. Fopp, P. Brun, J. van der Hoogen, D. N. Karger, C. M. Casadei, **D Tuia**, N. E. Zimmermann, T. Crowther, and L. Pellissier. Regional uniqueness of tree species composition and response to forest loss and climate change. *Nature Comm.*, 15(4375), 2024.

<https://www.nature.com/articles/s41467-024-48276-3>

Where are they?

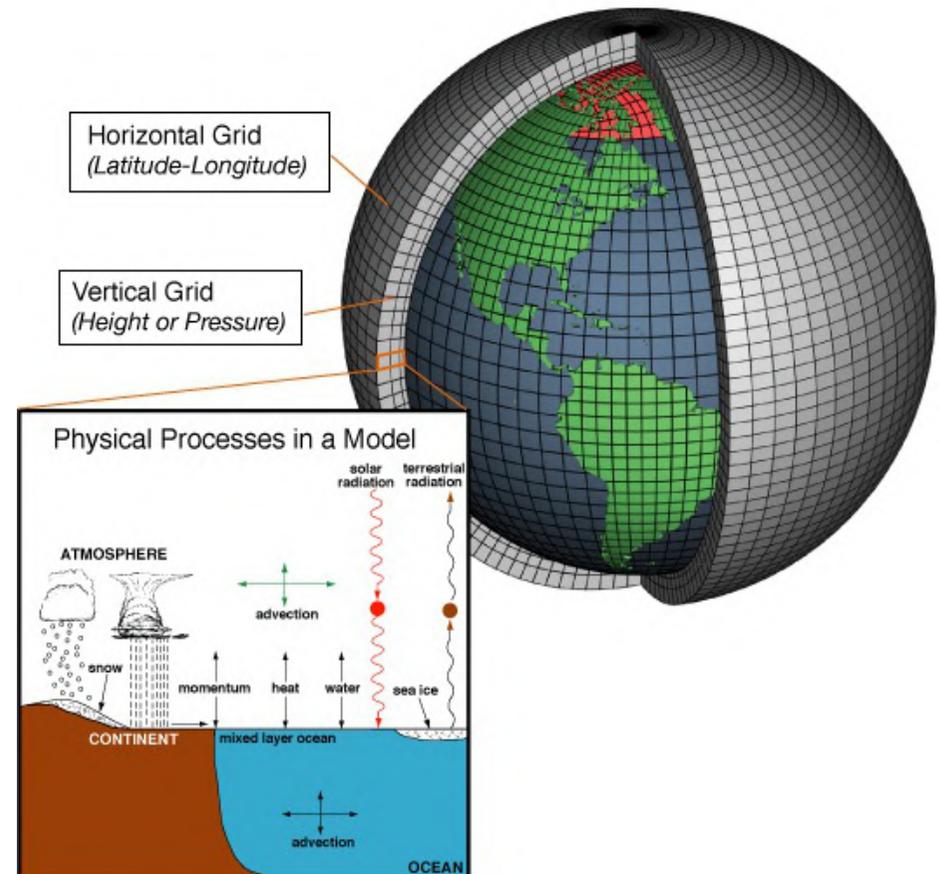


N. van Tiel, L. Lyu, F. Fopp, P. Brun, J. van der Hoogen, D. N. Karger, C. M. Casadei, **D Tuia.**, N. E. Zimmermann, T. Crowther, and L. Pellissier. Regional uniqueness of tree species composition and response to forest loss and climate change. *Nature Comm.*, 15(4375), 2024.

<https://www.nature.com/articles/s41467-024-48276-3>

Where will they be?

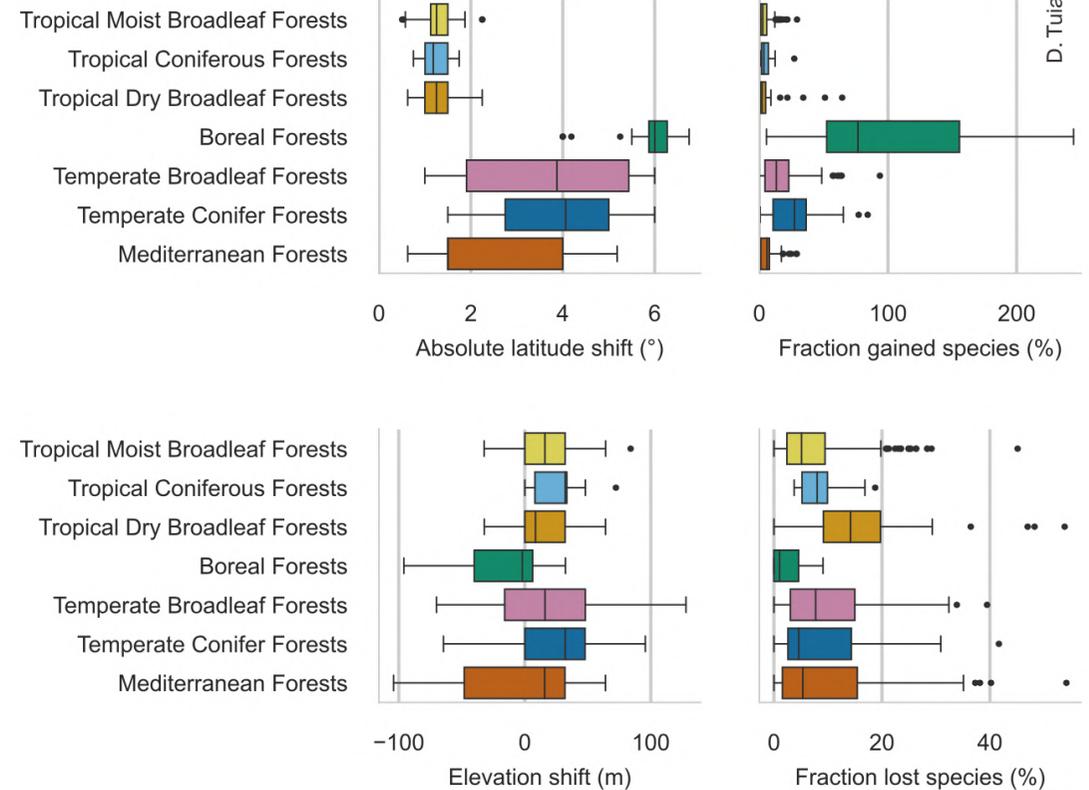
- We now simulate species distribution under future climate.
- We use General Circulation Models (GCMs)
- GCMs represent physical processes of the atmosphere and ocean to simulate response of climate to greenhouse gas emission
- Scenarios (SSP) represent an increase in Solar radiation entering Earth between 1750 and 2100





Where will they be?

- We run climate models to obtain climatic variables in 2071-2100.
- These are for SSP 5.81 (+5.81 W/m²)
- (SSP 4.5 is the scenario where we stay roughly on the same path)



N. van Tiel, L. Lyu, F. Fopp, P. Brun, J. van der Hoogen, D. N. Karger, C. M. Casadei, **D Tuia.**, N. E. Zimmermann, T. Crowther, and L. Pellissier. Regional uniqueness of tree species composition and response to forest loss and climate change. *Nature Comm.*, 15(4375), 2024.

■ <https://www.nature.com/articles/s41467-024-48276-3>

Enabling new scientific discovery

- Such mapping enables scientific discovery.
- We deliver the means to ecologists to scale up their valuable fieldwork
- We can now study:
 - biodiversity richness in space and time
 - interactions between species
- For practitioners, we can help focus conservation and restoration efforts

There is still a lot to do!

- The ML models can be improved.
- There is a lively community in “deep SDMs”:

New Results

 [Follow this preprint](#)

The Performance and Potential of Deep Learning for Predicting Species Distributions

 Benjamin Kellenberger,  Kevin Winner,  Walter Jetz

doi: <https://doi.org/10.1101/2024.08.09.607358>

- The climate models remain static, the species don't adapt to climate change.



Towards environmental machine learning that is

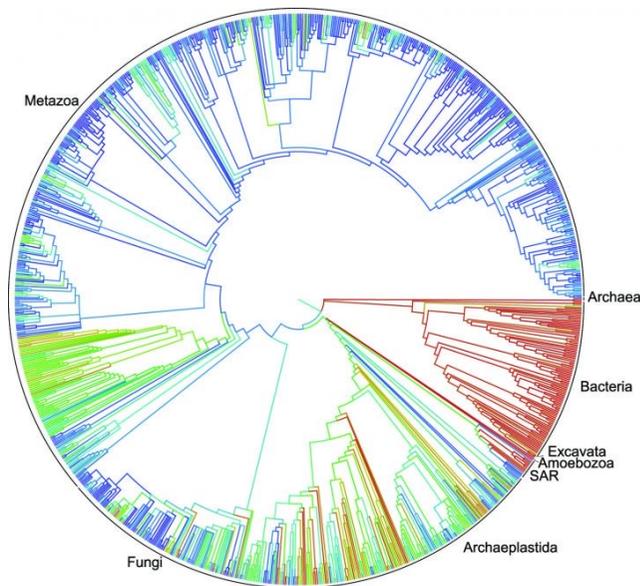
Scalable

Knowledge-driven

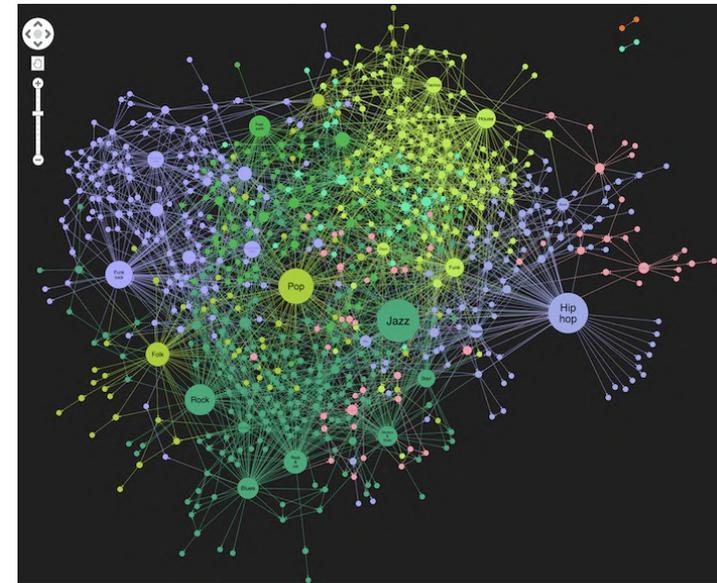
Accessible to anyone

But do we need to extract all information from data?

- Many things about the world, we know them from knowledge



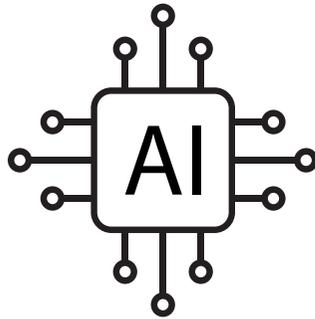
Scientific knowledge



Common sense knowledge

Do we need to extract all information from data?

- Many things about the world, we know them from knowledge.
- Machine learning models tend to learn everything from data, as if we knew nothing about the world.



=

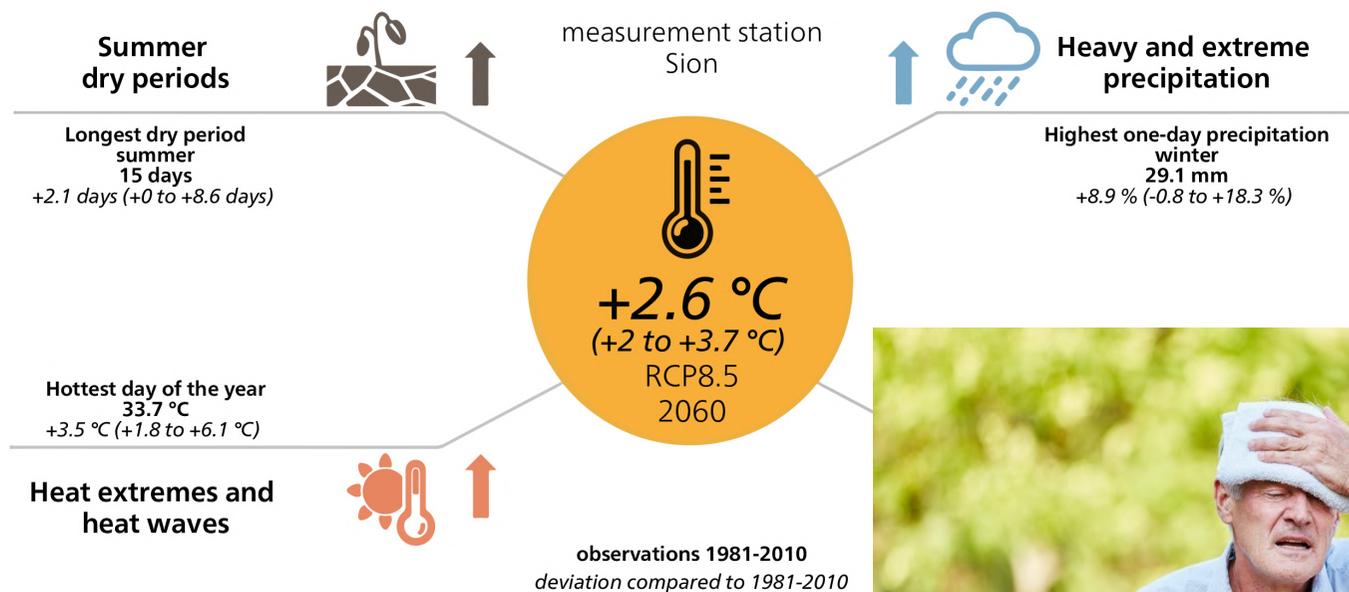


Do we need to extract all information from data?

- Many things about the world, we know them from knowledge.
- Machine learning models tend to learn everything from data, as if we knew nothing about the world.
- **Integrating domain knowledge** is crucial for models that are meaningful



It's getting hotter everywhere! In Valais, Switzerland:



Nccs.admin.ch



Familydoctor.org

When it's warm what do you do?



Flickr



Flickr

What can the forest do?



<https://www.youtube.com/watch?v=8vjDjCv-ekM>

What can the forest do?

- For trees, it is slower. But still it moves!
- Those movements most often happen at the upper **treeline**

- **What's a treeline?**

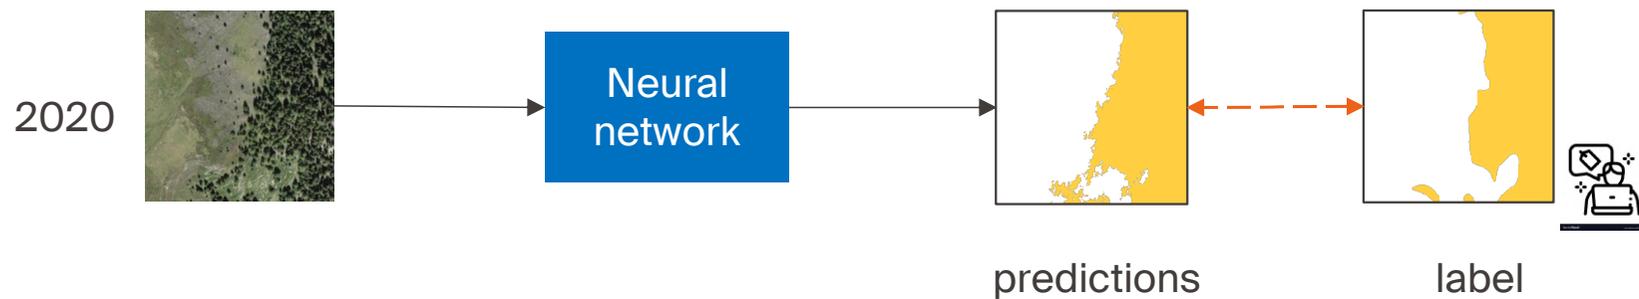
- But first, we need to map the forest.



<https://farsouthecology.com/do-treelines-in-the-southern-hemisphere-follow-the-rules/>

We could rush into semantic segmentation

- The technology is there
- We have labels
- It's just a matter of compute, right?



Results reproduce well the ground truth, but...



Aerial image



SwissTLM3D labels



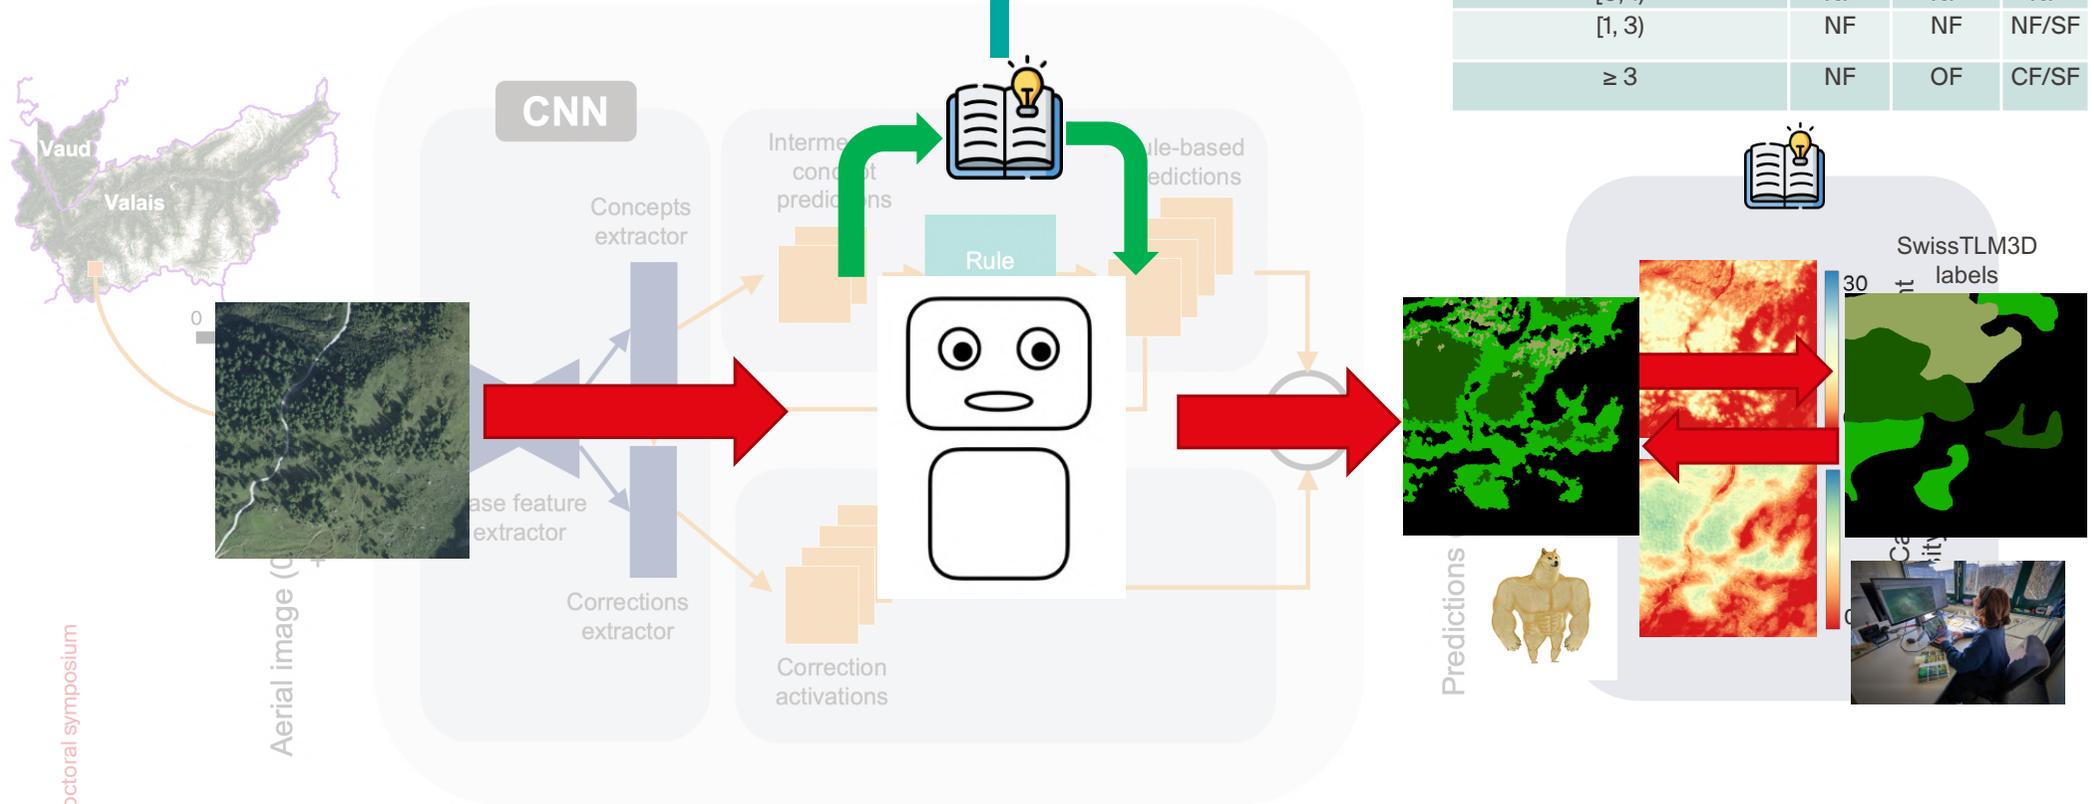
"Black-box" model predictions



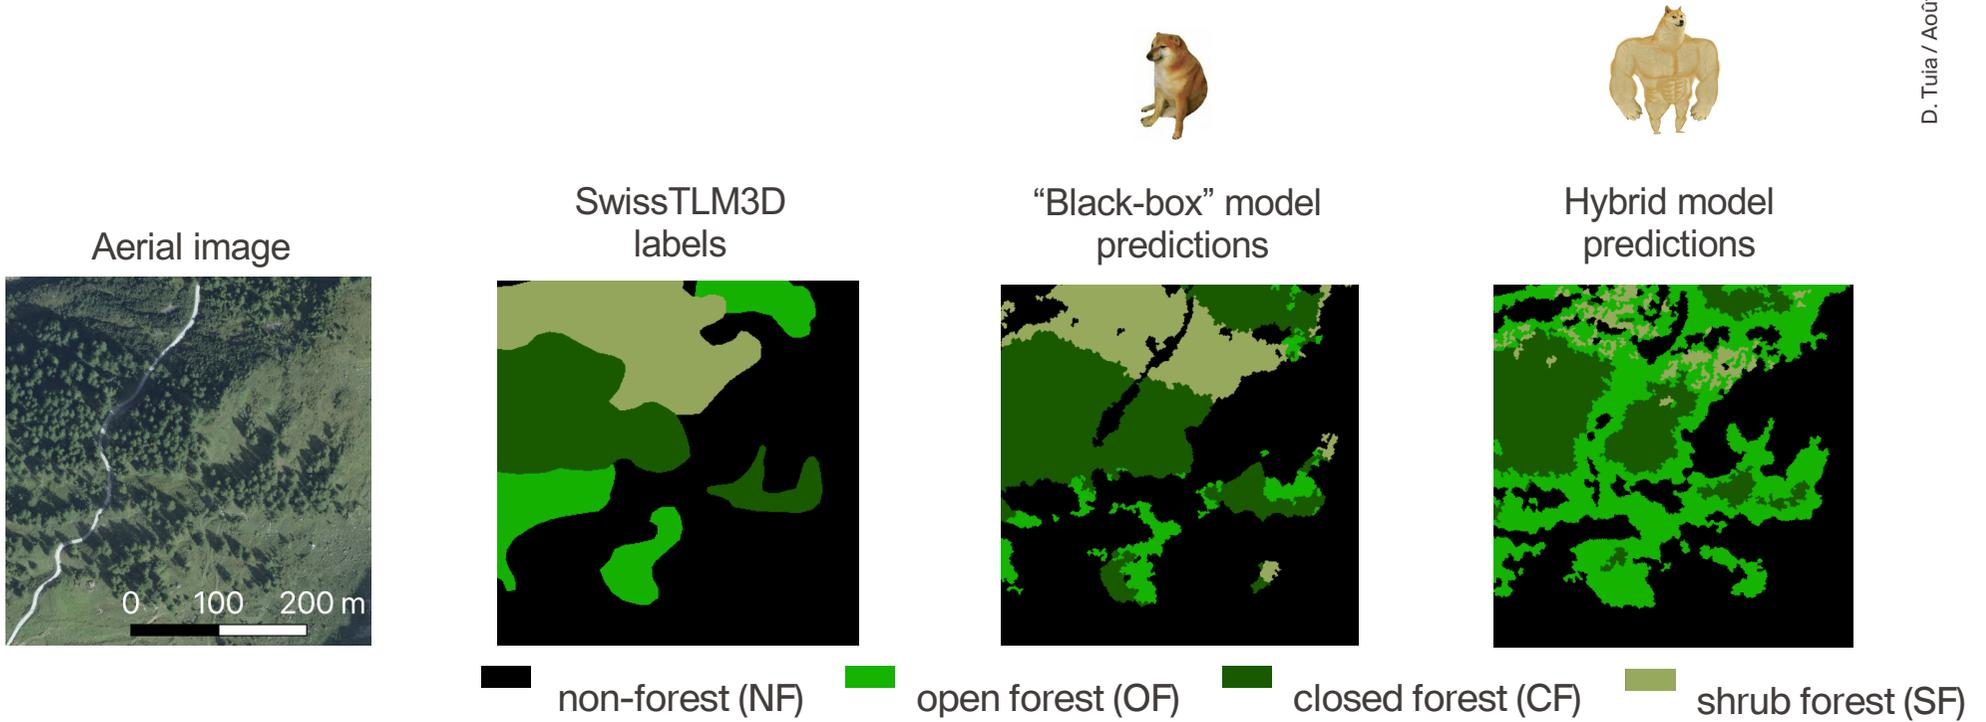
T.-A. Nguyen, B. Kellenberger, and D. Tuia. Mapping forest in the Swiss Alps treeline ecotone with explainable deep learning. *Remote Sens. Environ.*, 281(113217), 2022.



Integrating forest definitions in segmentation models



Results better align to forest definitions



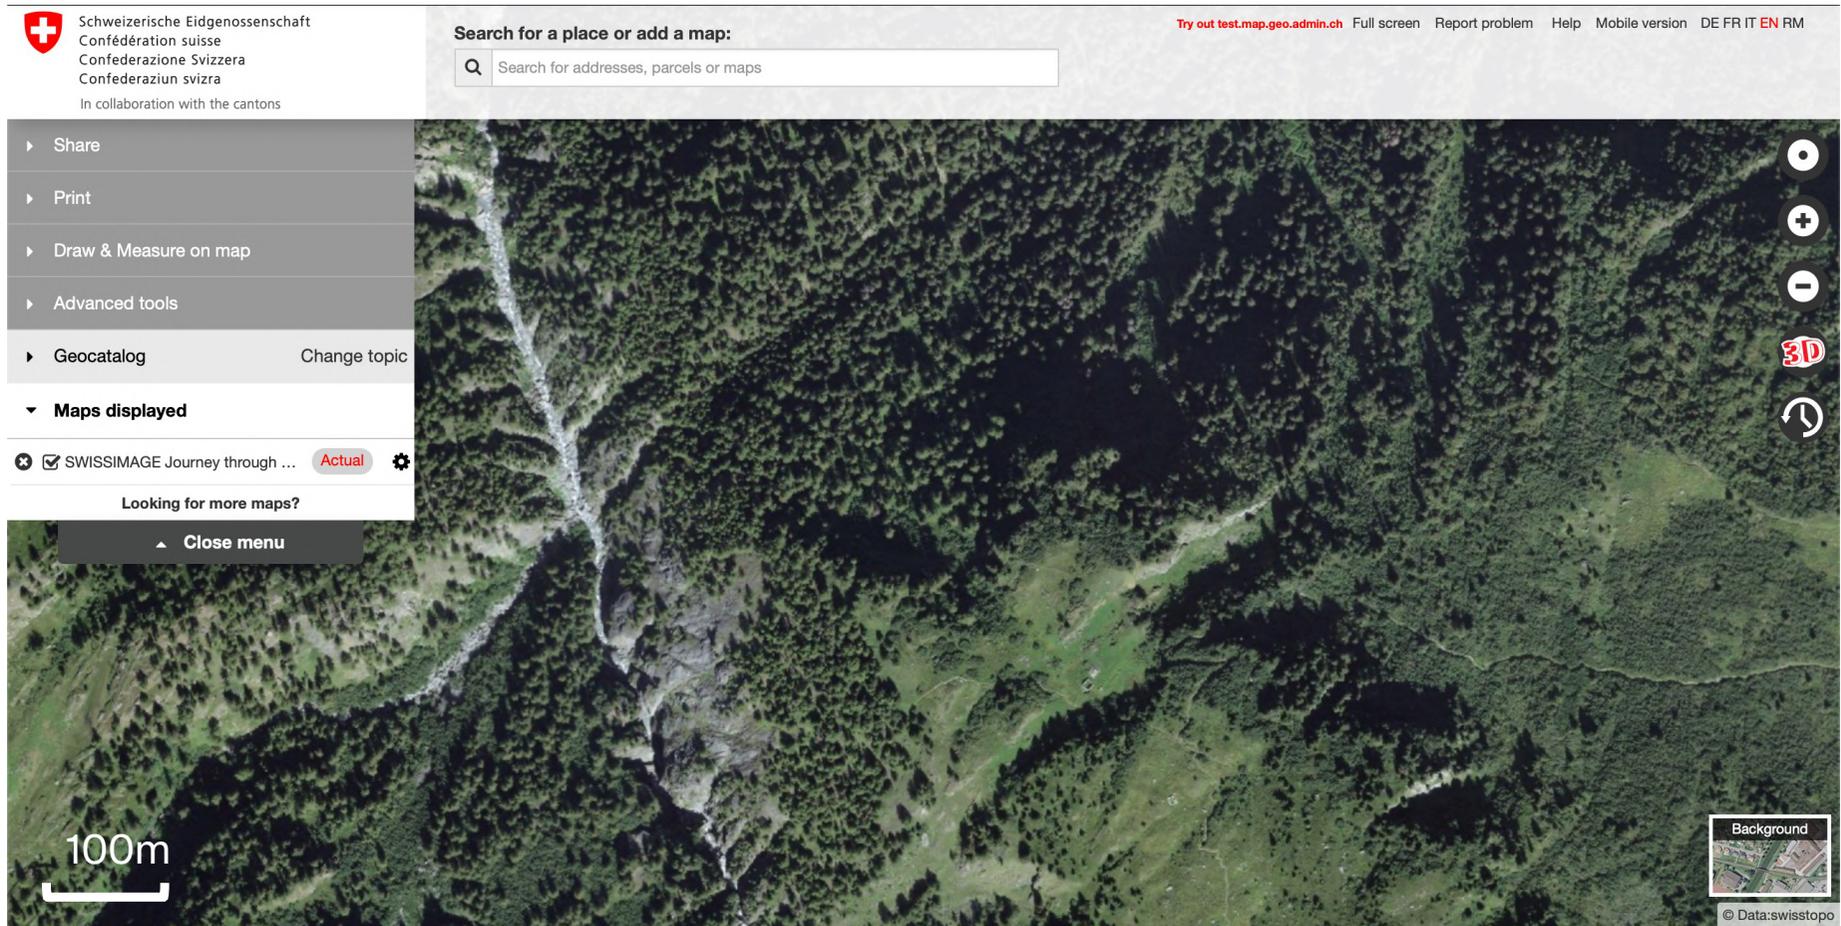
T.-A. Nguyen, B. Kellenberger, and D. Tuia. Mapping forest in the Swiss Alps treeline ecotone with explainable deep learning. *Remote Sens. Environ.*, 281(113217), 2022.



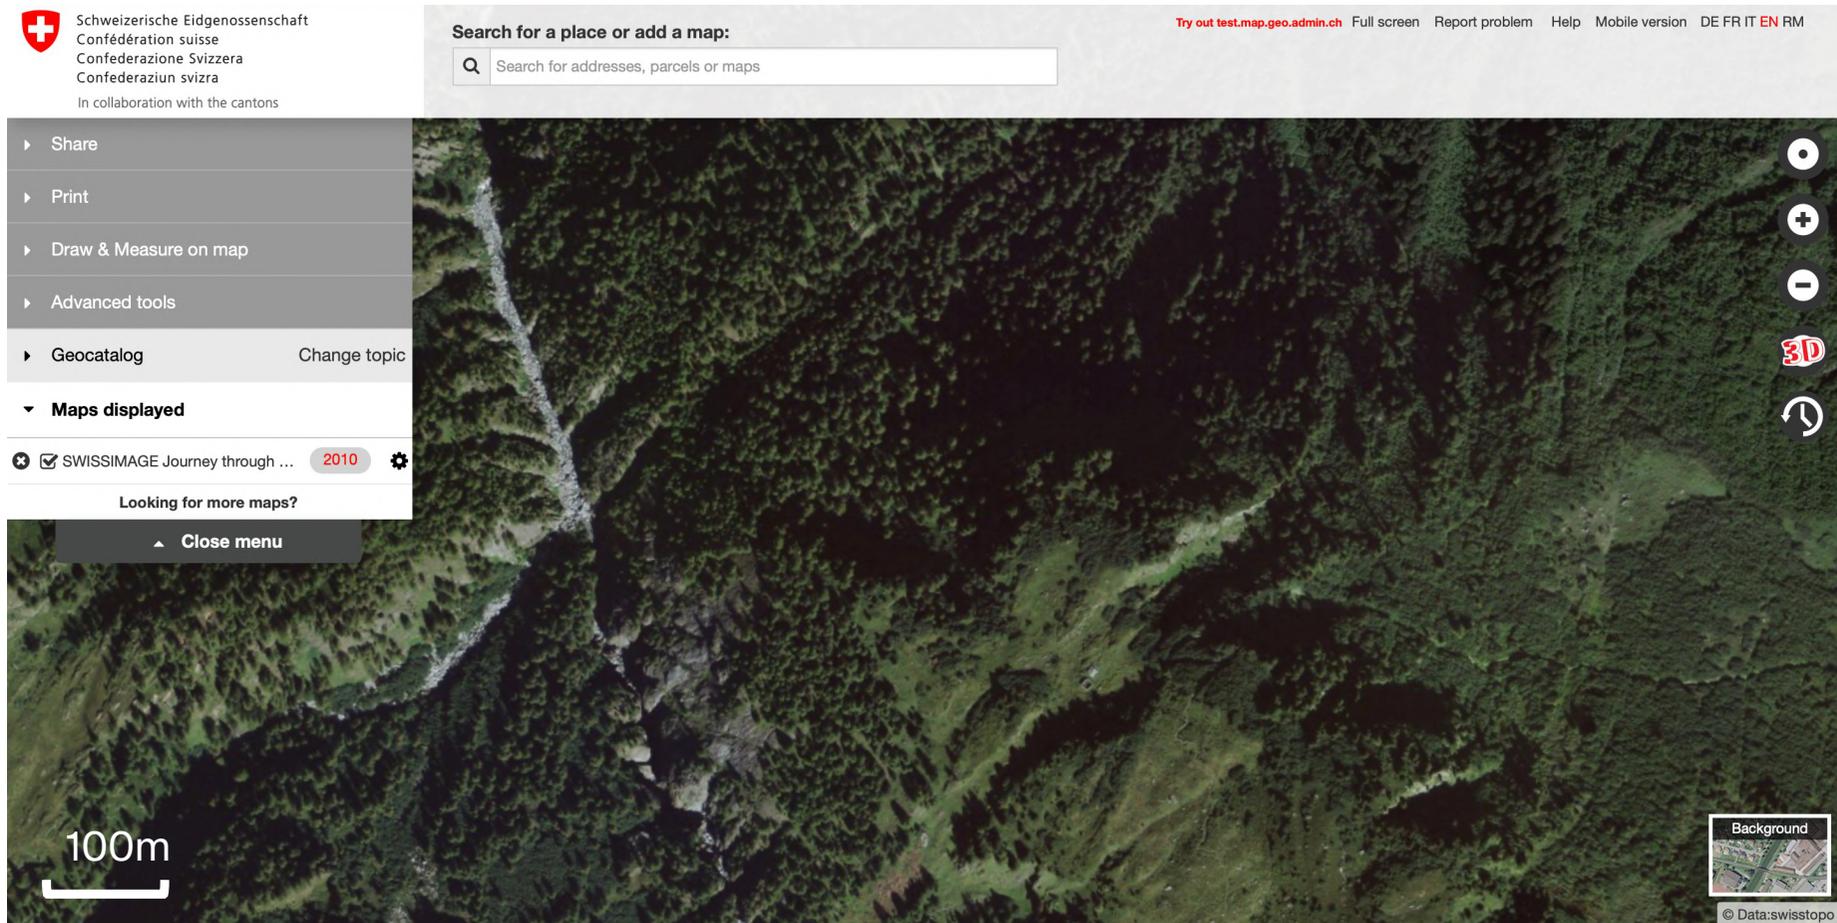
What about the time dimension now?



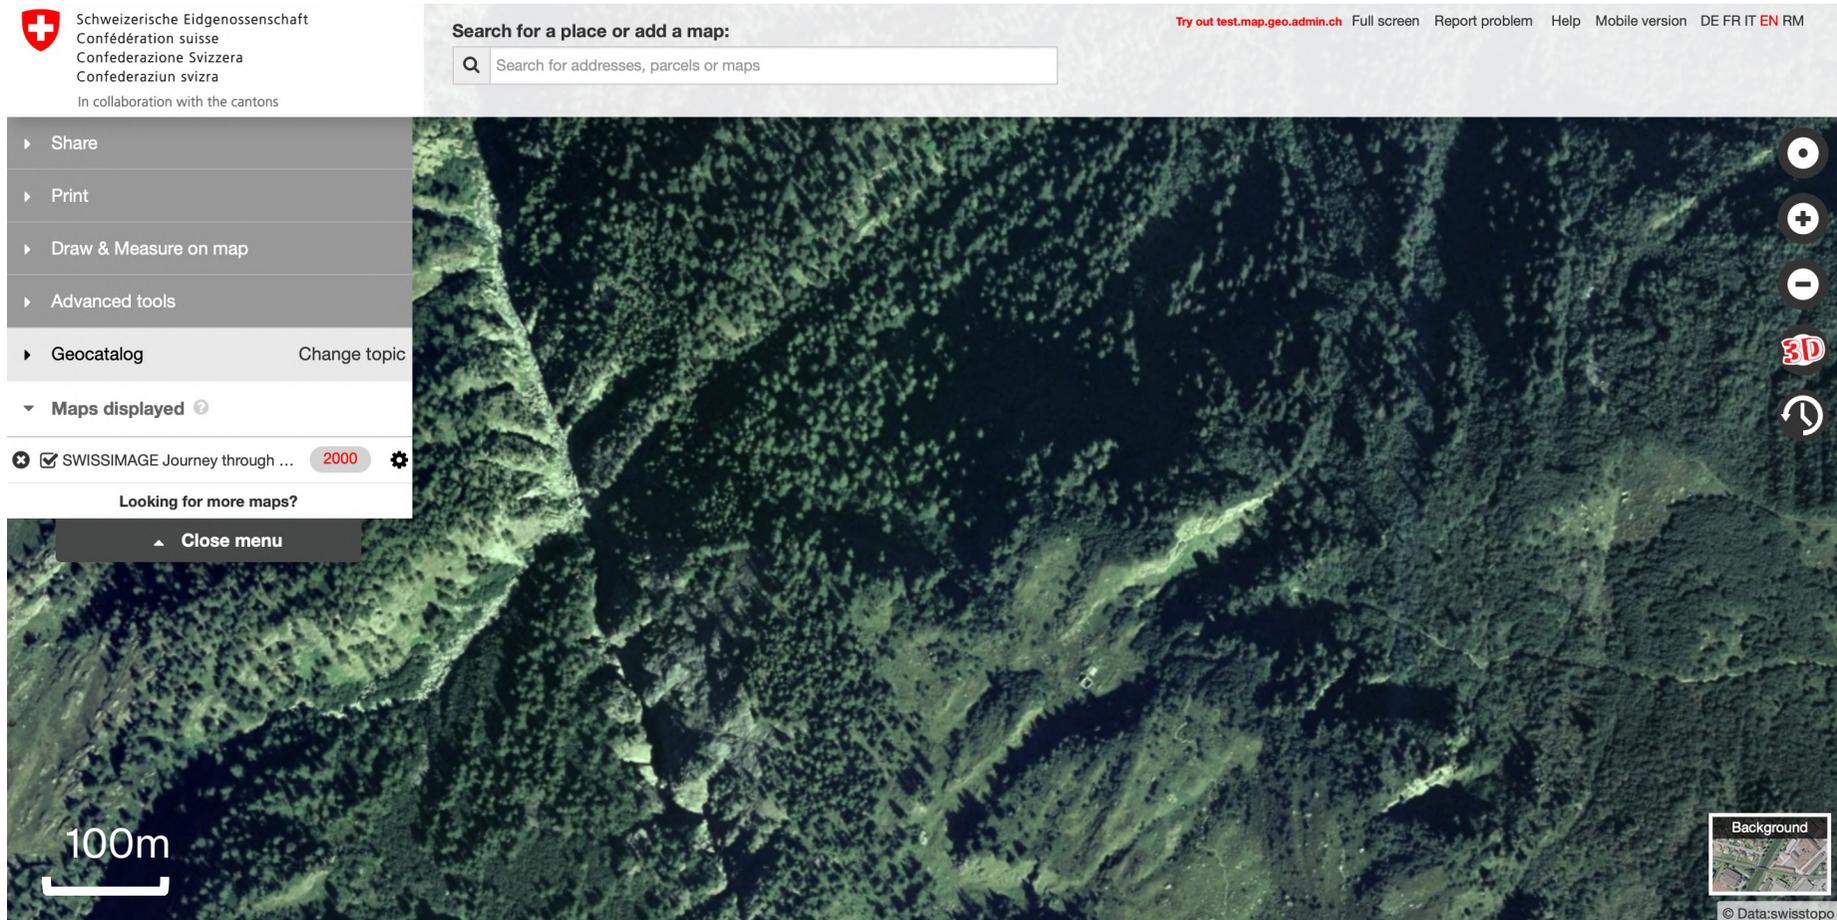
In Switzerland we have a super resource : aerial photography since the 1940s!



In Switzerland we have a super resource : aerial photography since the 1940s!



In Switzerland we have a super resource : aerial photography since the 1940s!



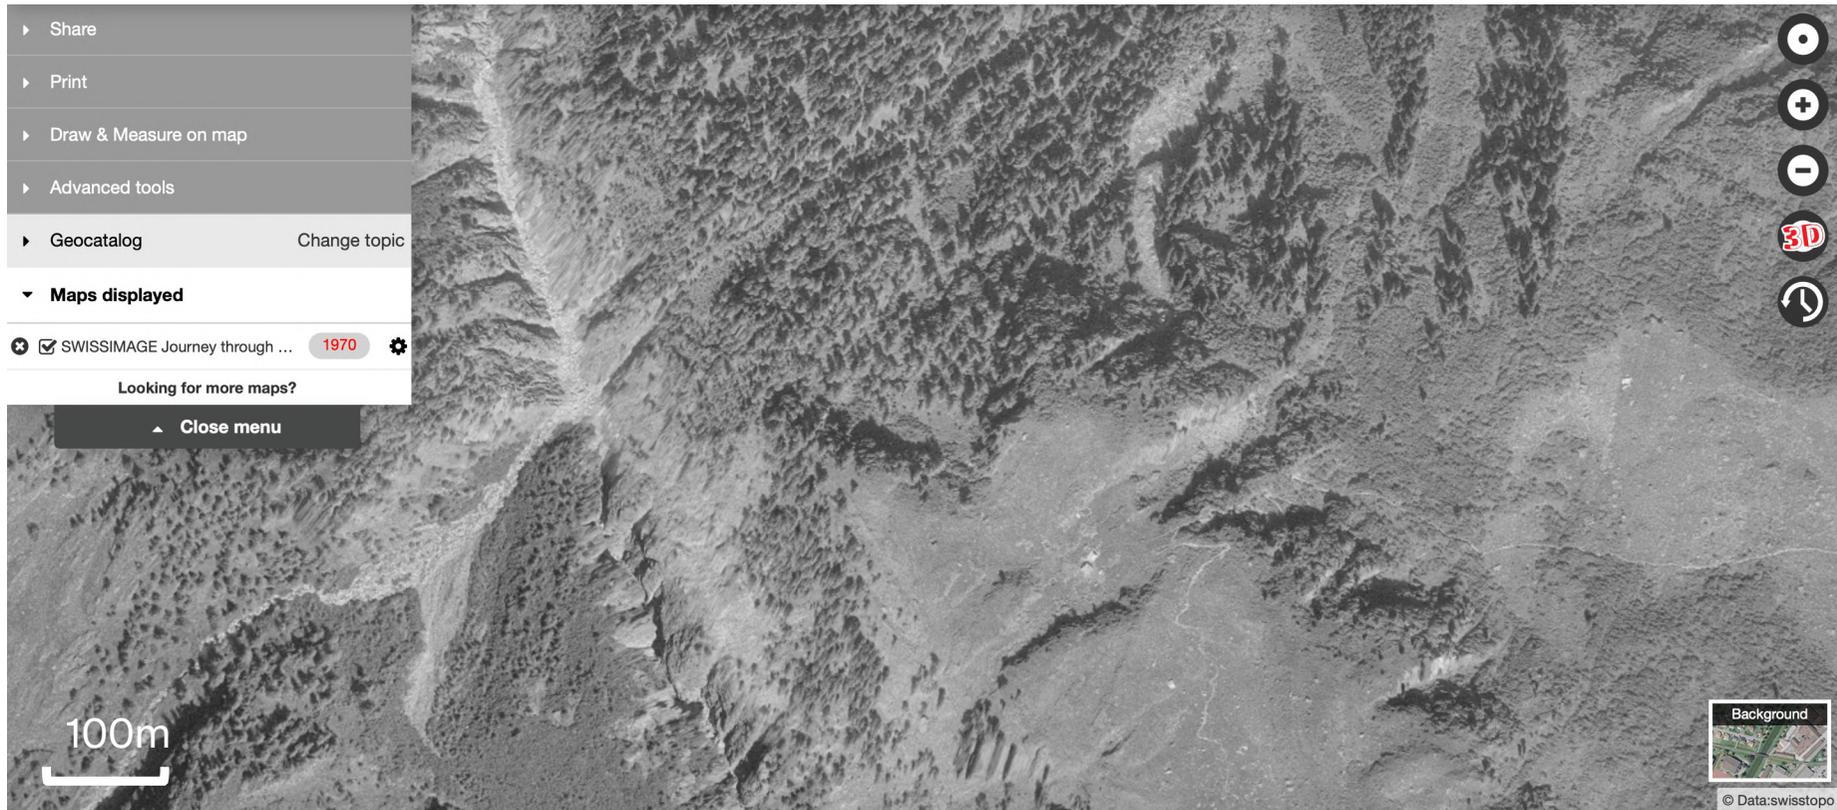
In Switzerland we have a super resource : aerial photography since the 1940s!



In Switzerland we have a super resource : aerial photography since the 1940s!



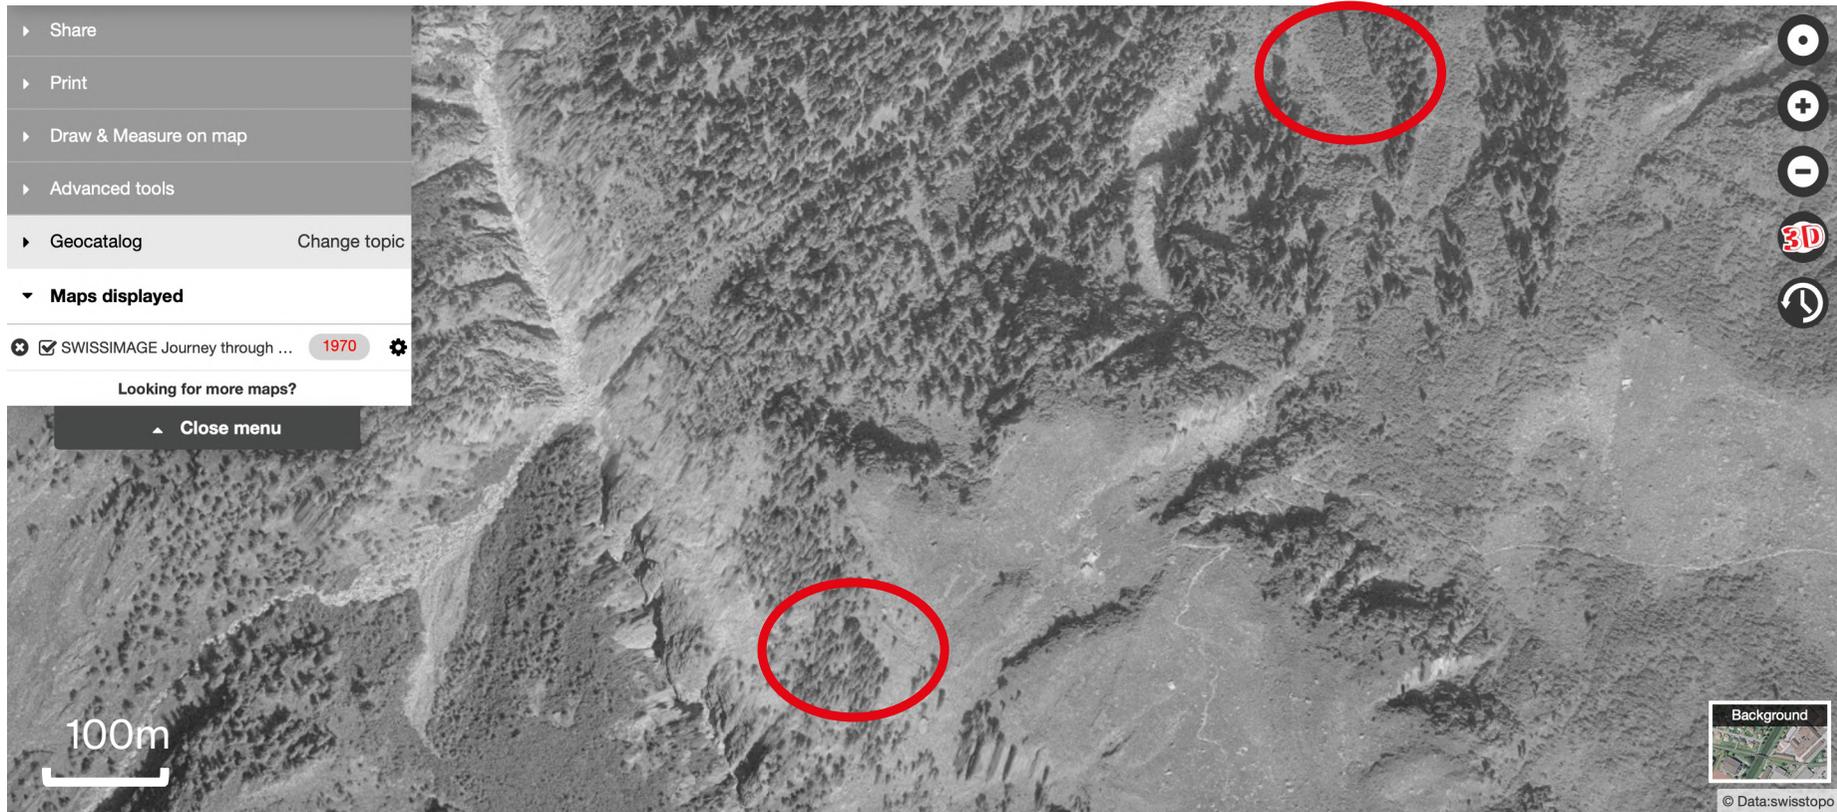
In Switzerland we have a super resource : aerial photography since the 1940s!



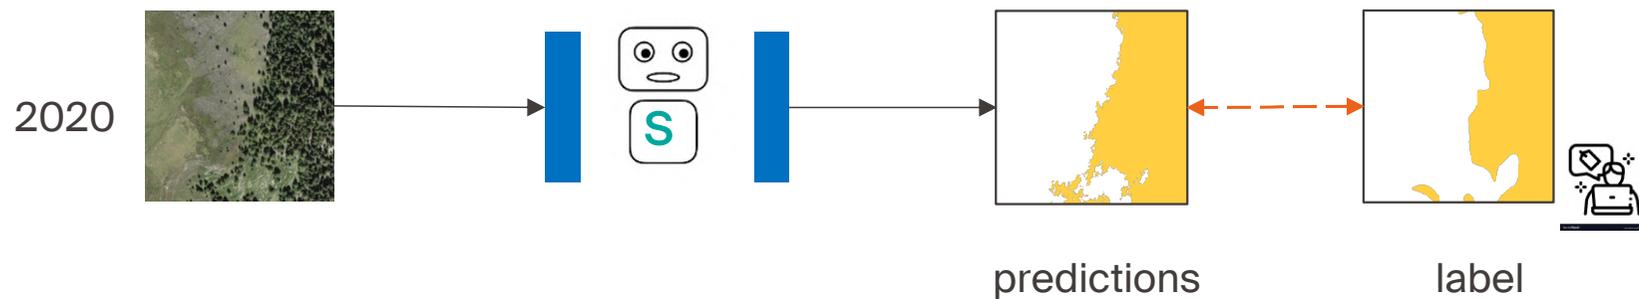
In Switzerland we have a super resource : aerial photography since the 1940s!



In Switzerland we have a super resource : aerial photography since the 1940s!



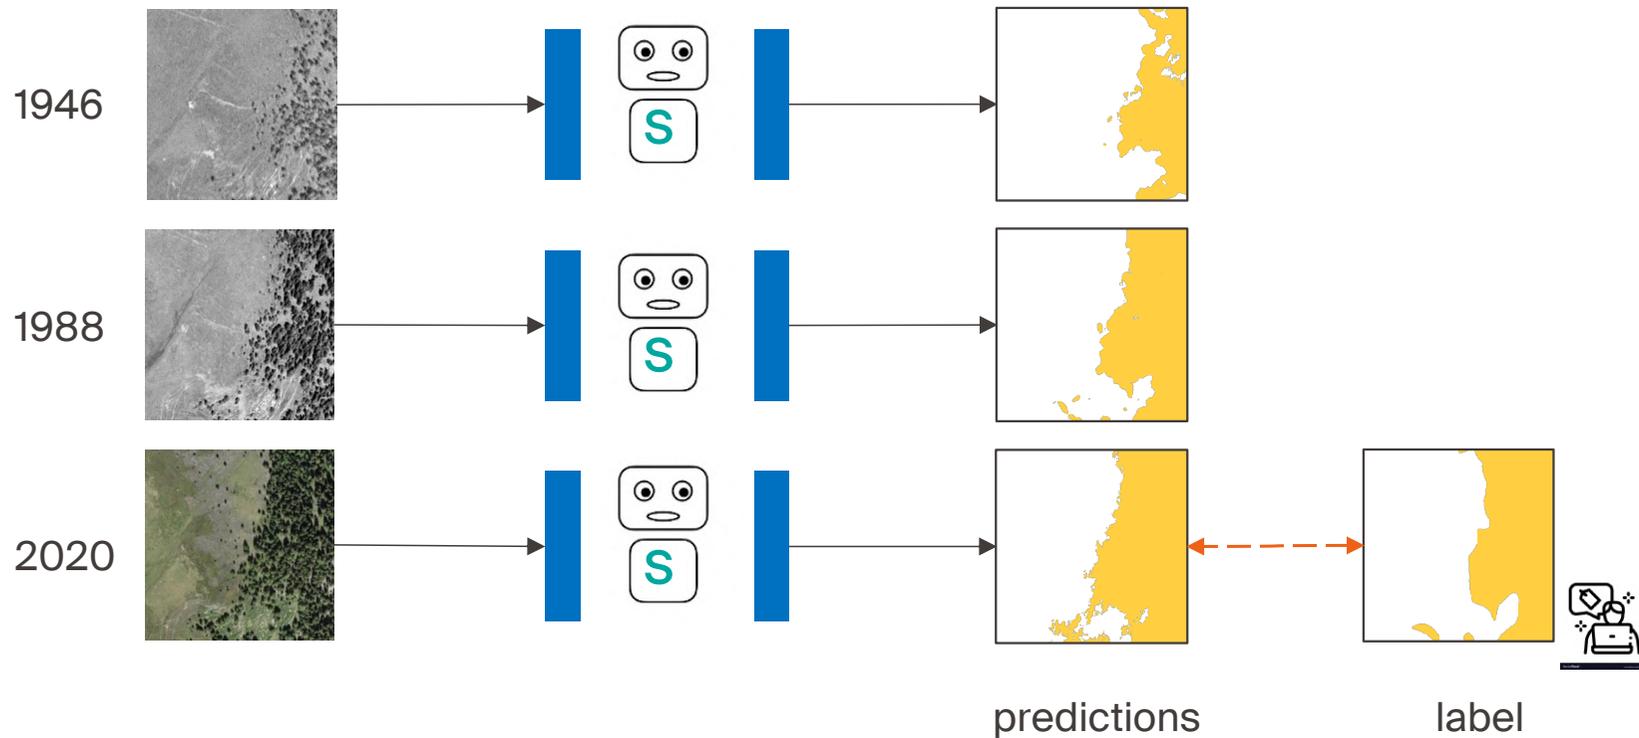
A single time classifier



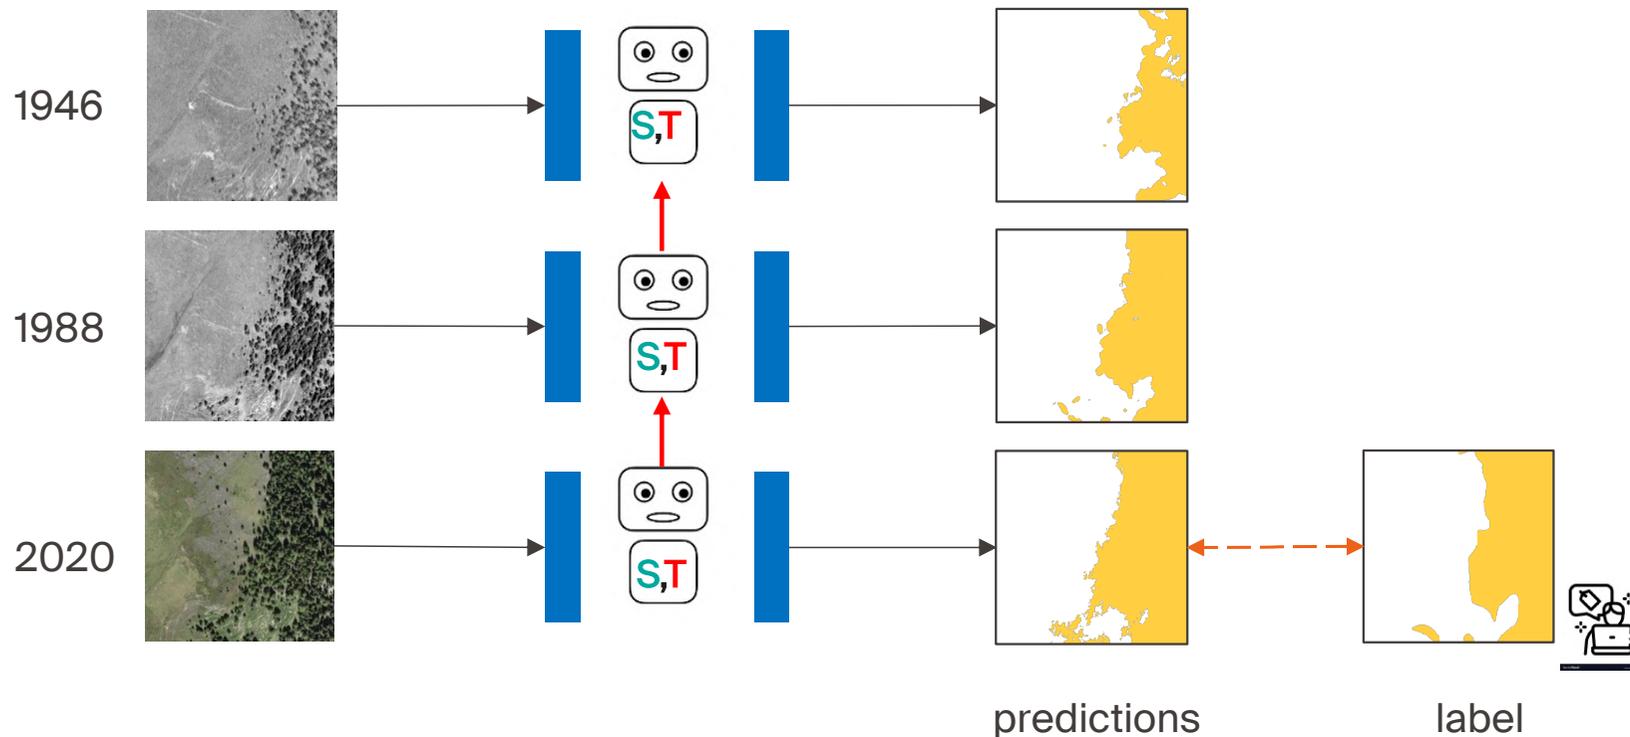
T.-A. Nguyen, B. Kellenberger, and D. Tuia. Mapping forest in the Swiss Alps treeline ecotone with explainable deep learning. *Remote Sens. Environ.*, 281(113217), 2022.

<https://doi.org/10.1016/j.rse.2022.113217>

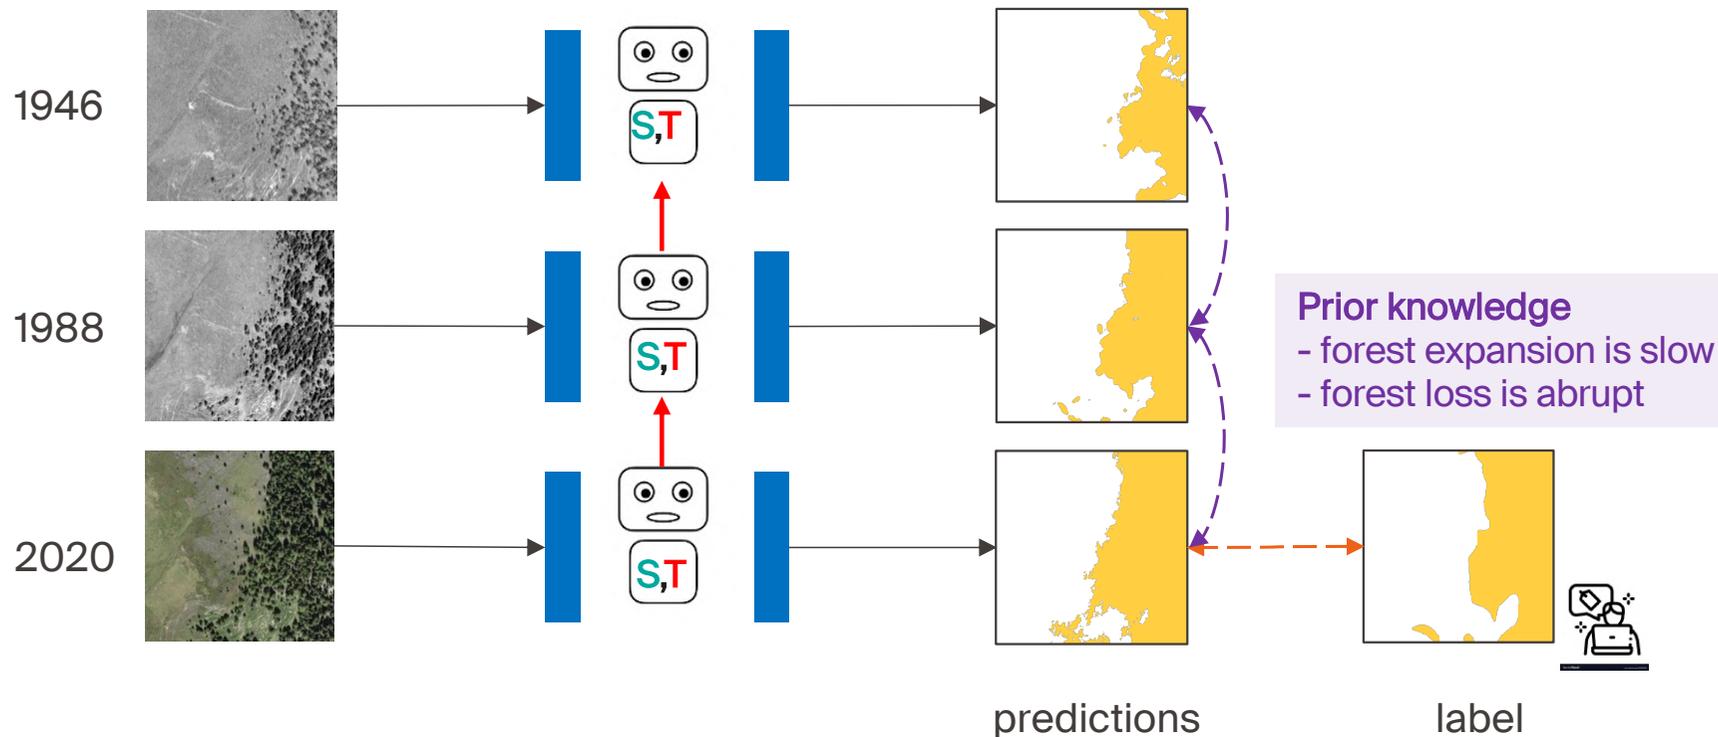
A multitemporal classifier

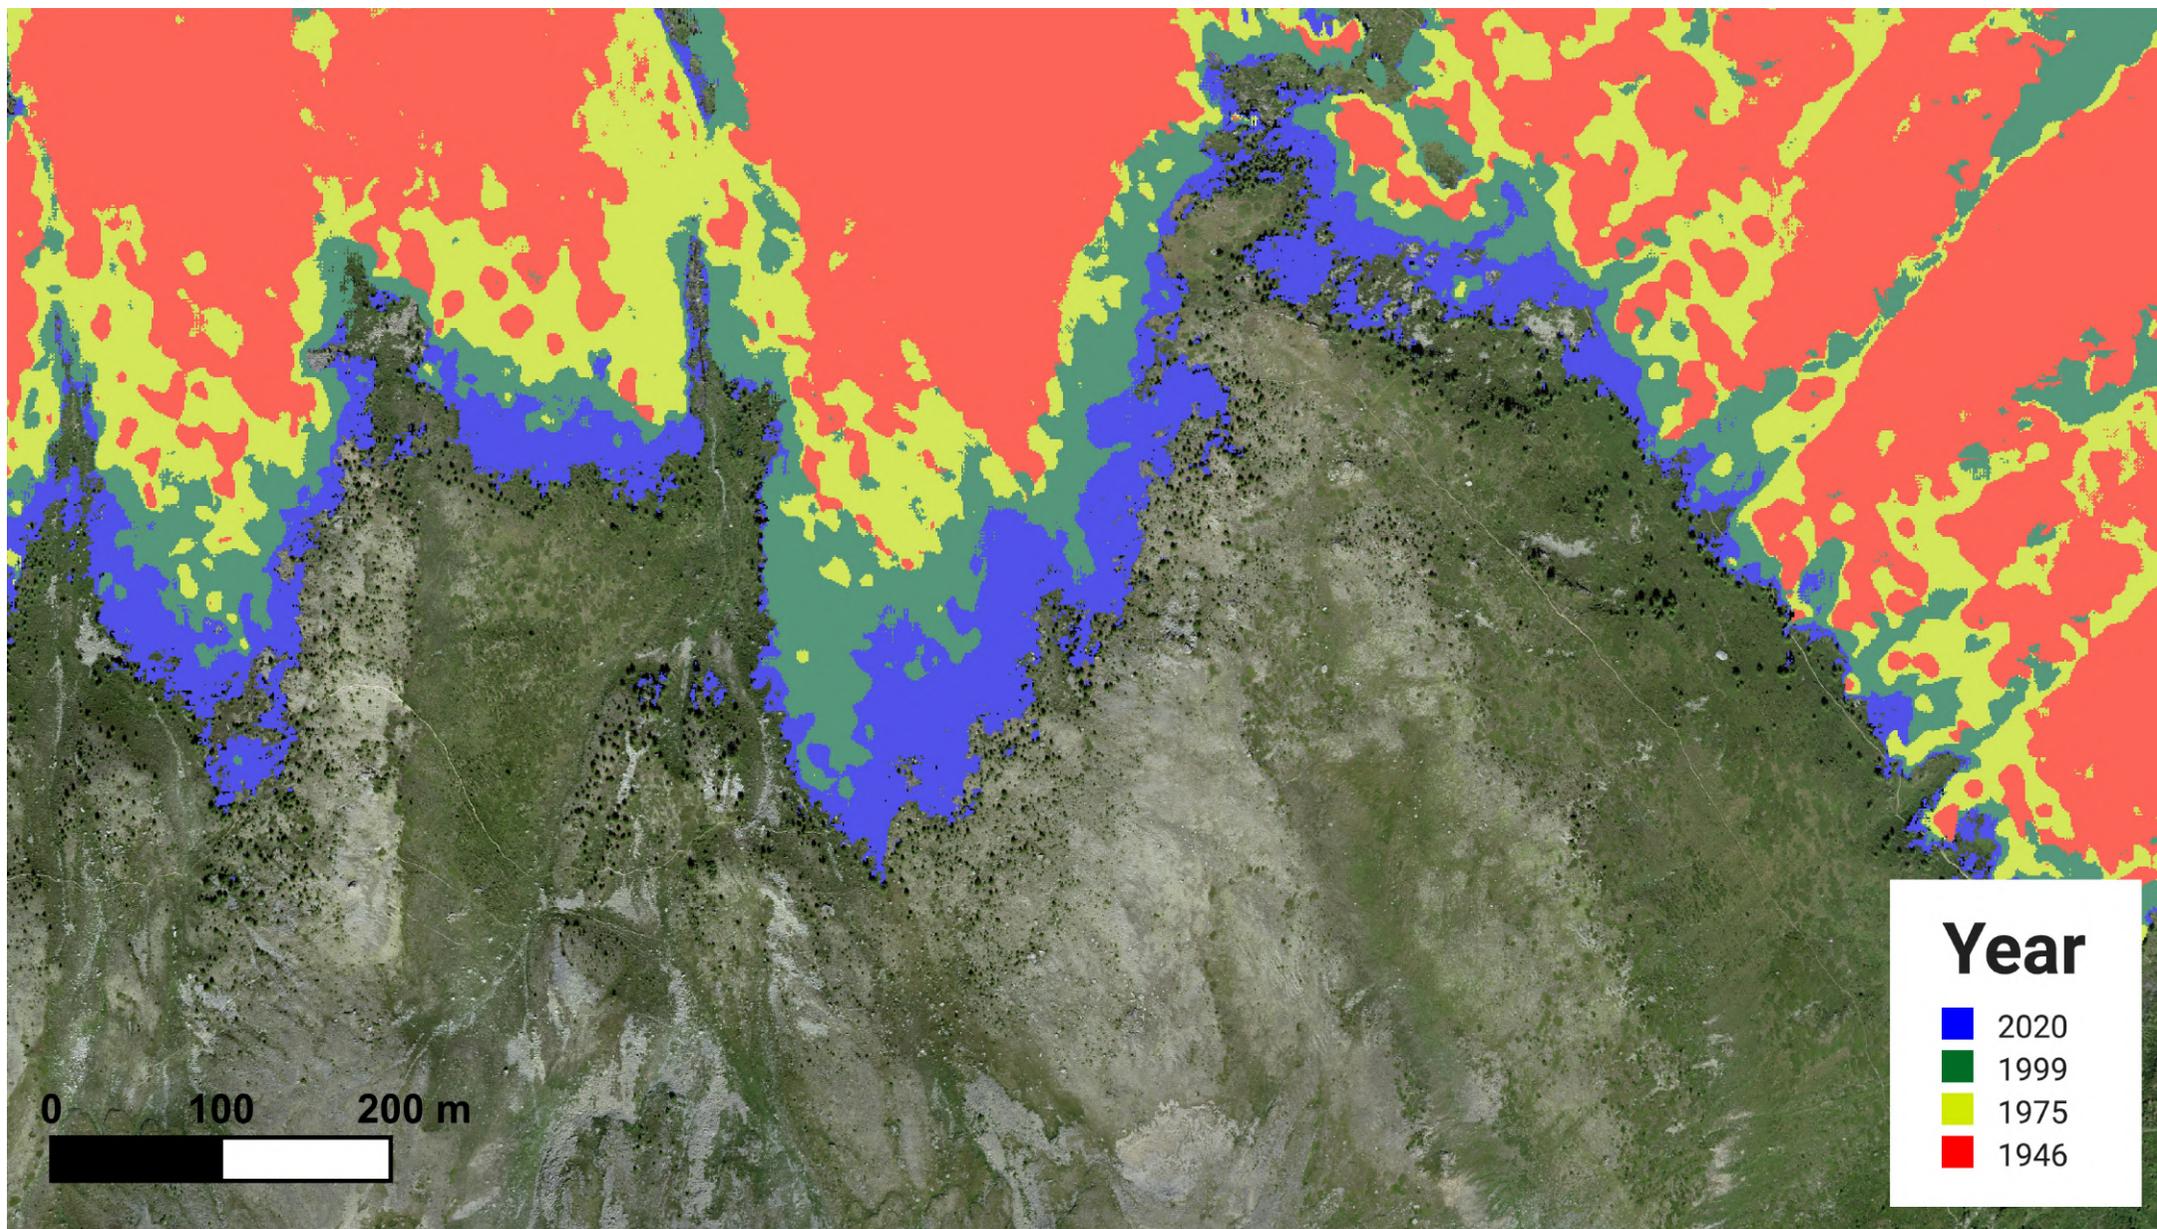


A multitemporal classifier enforcing forest dynamic knowledge

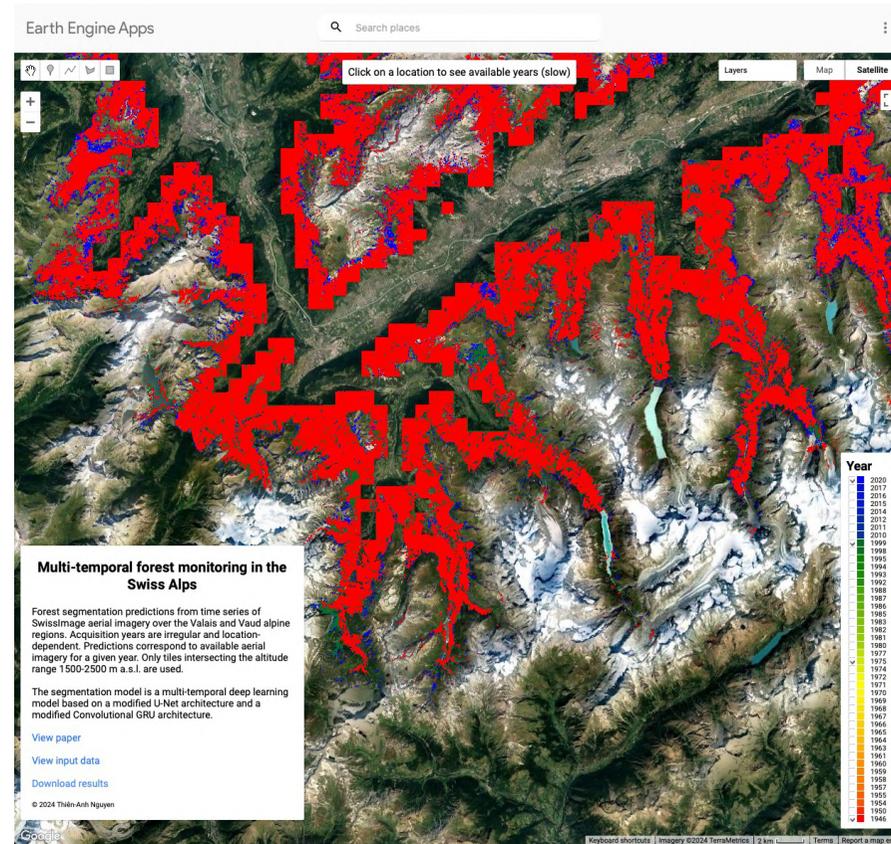


A multitemporal classifier enforcing forest dynamic knowledge

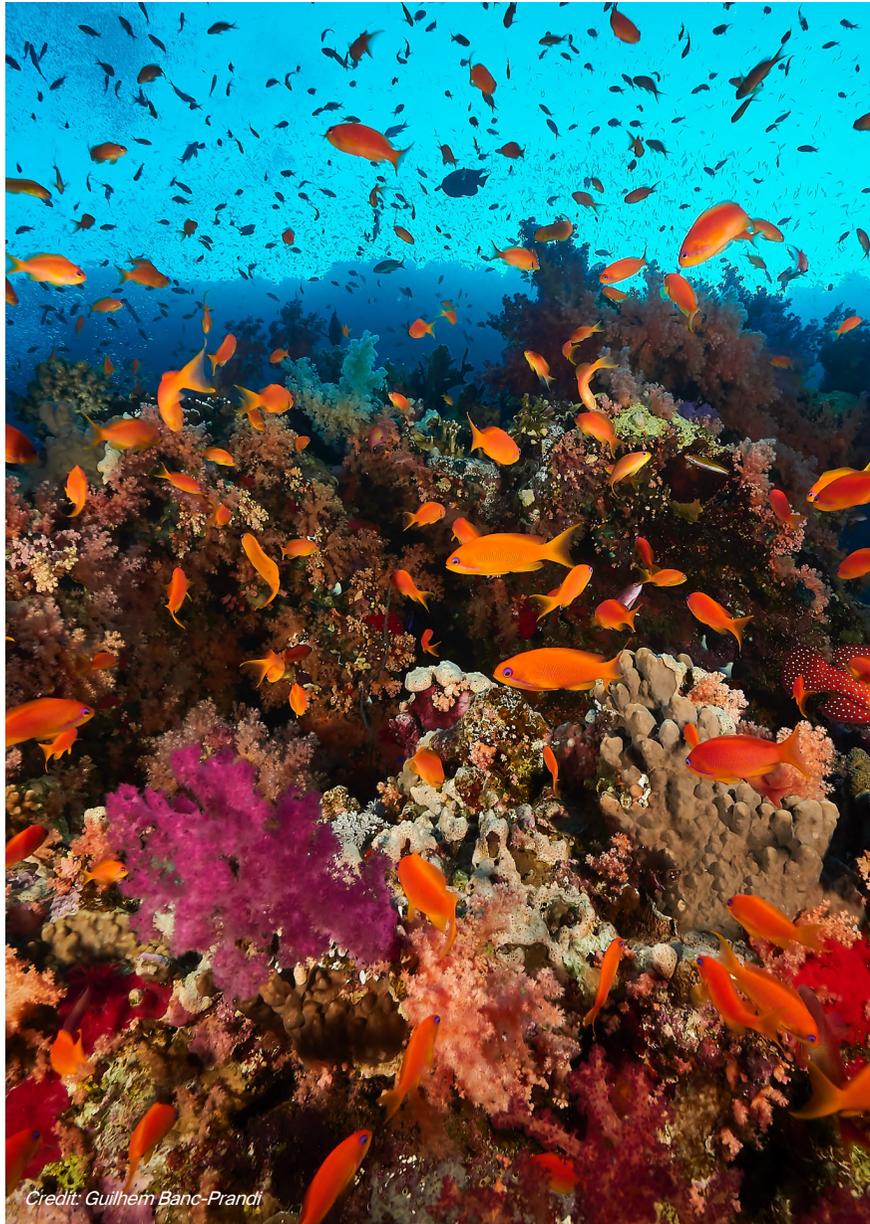




Want to explore?



- <https://temp-forest-mapping.projects.earthengine.app/view/multitempforestmap>



Credit: Guilhem Banc-Prandi

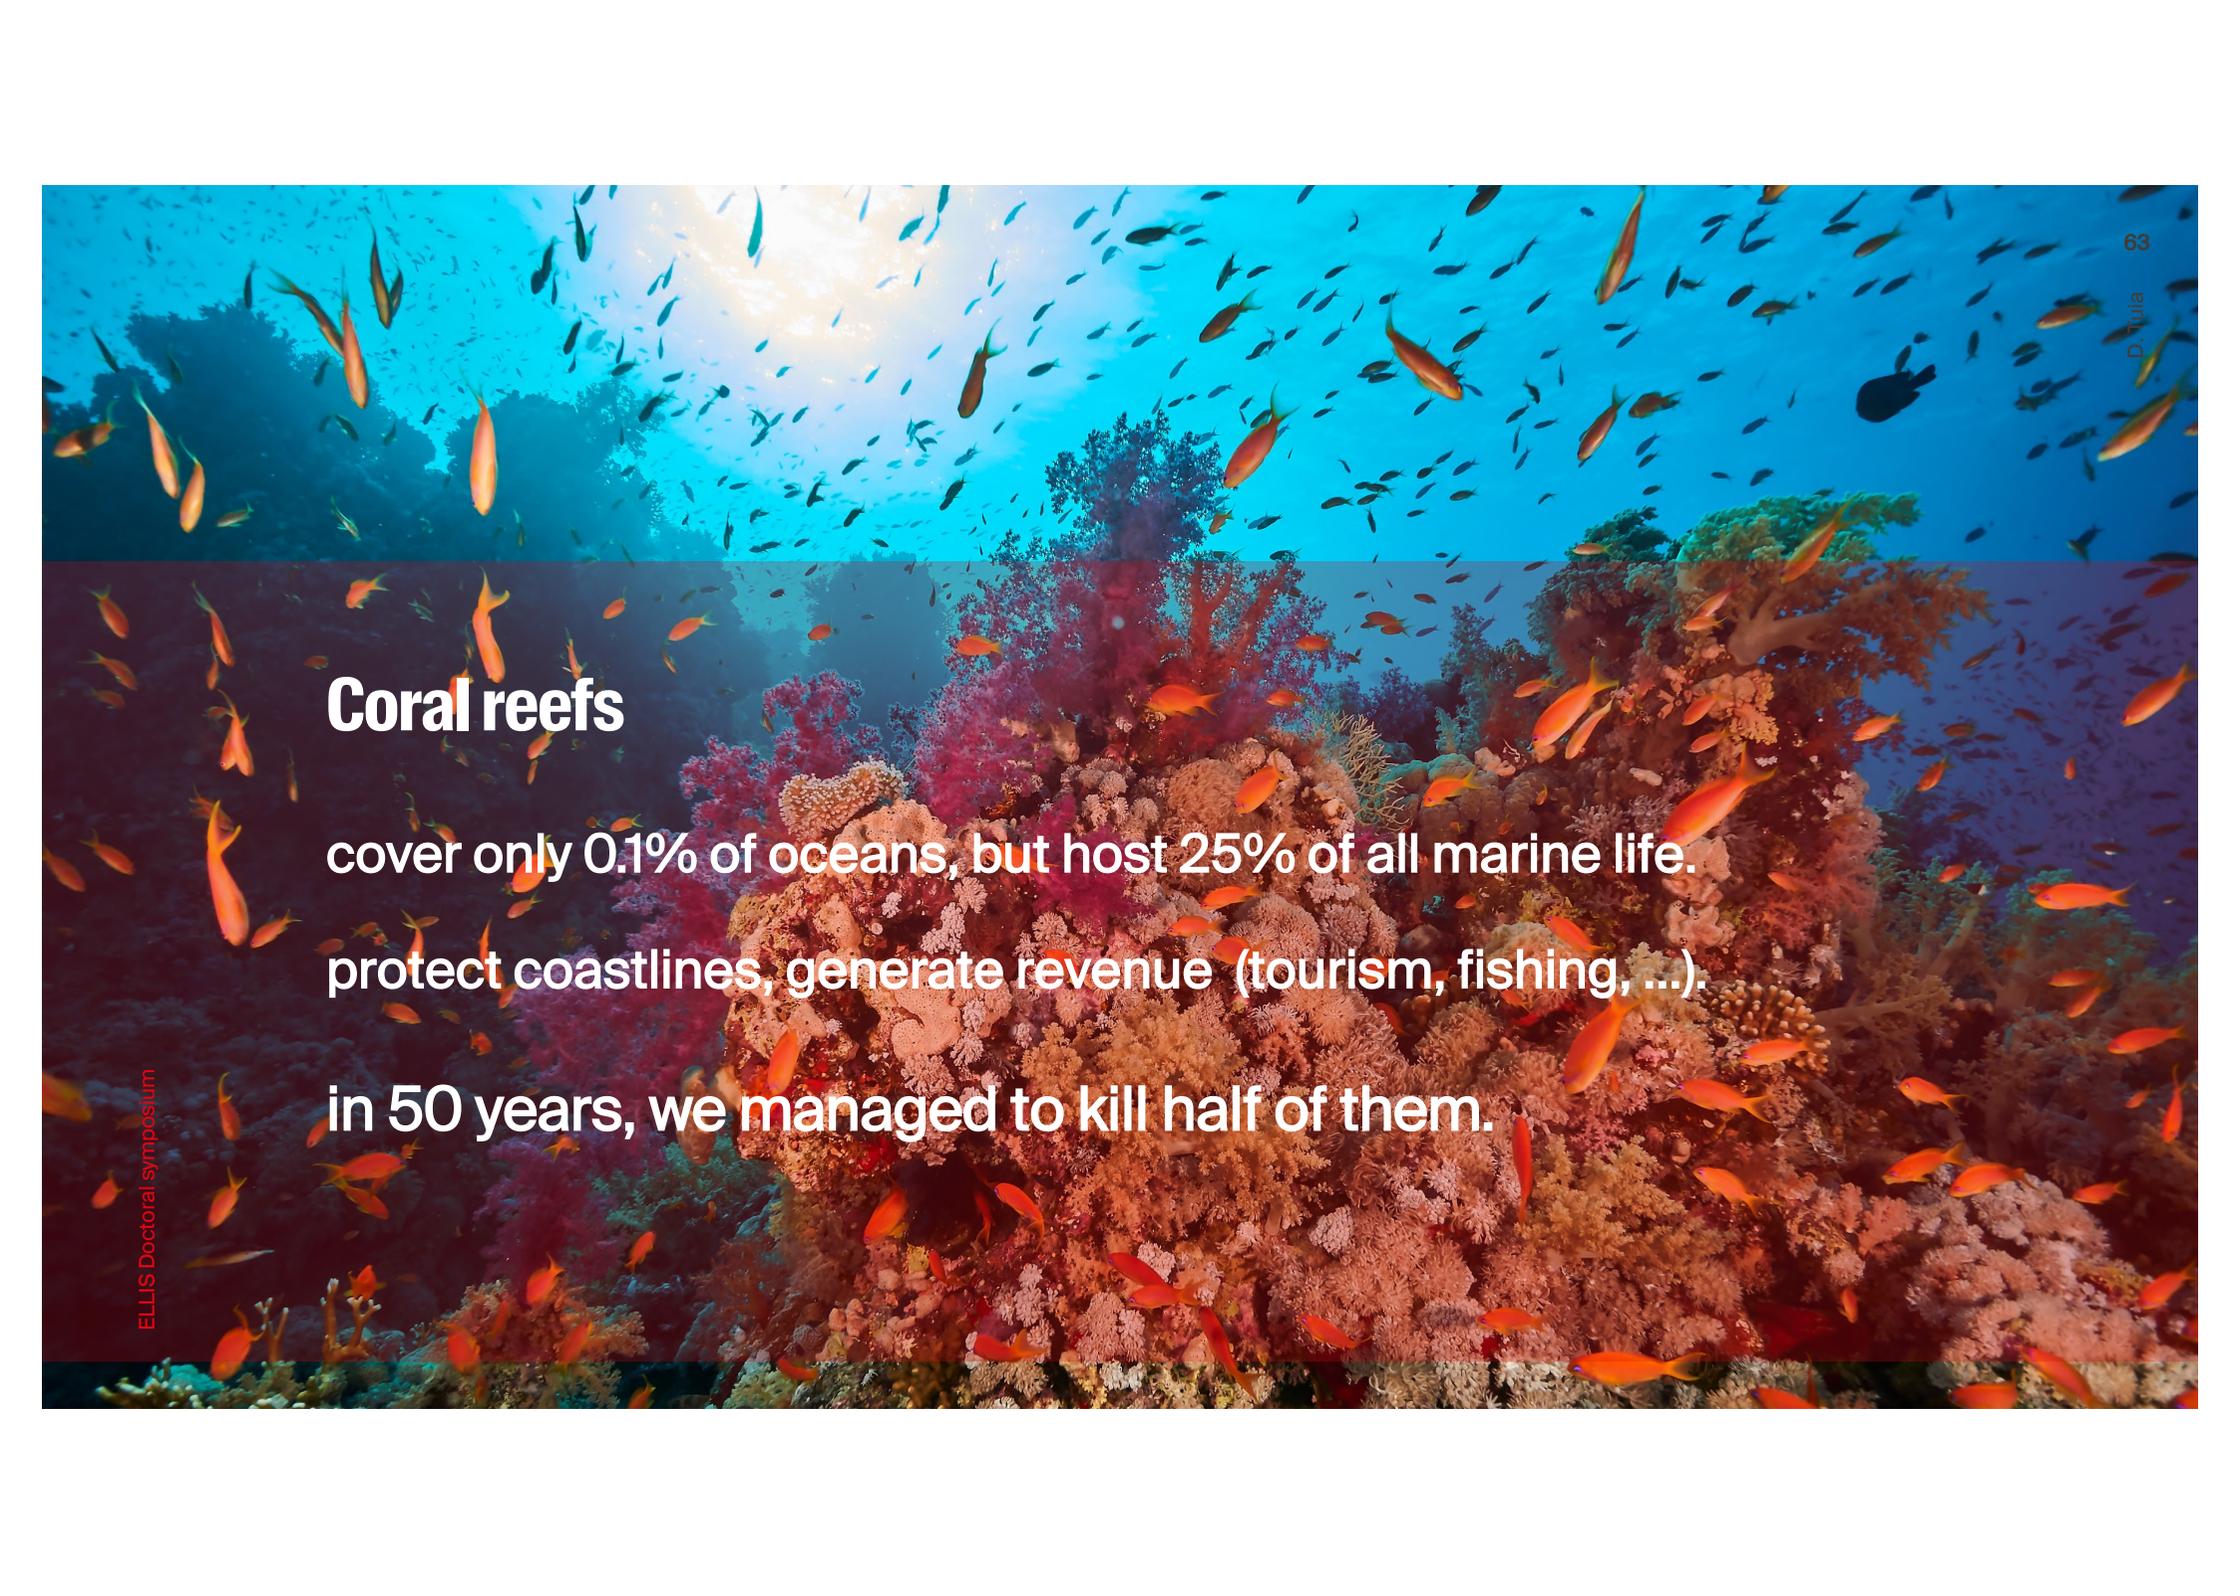
Towards environmental machine learning that is

Scalable

Knowledge-driven

Accessible to anyone





Coral reefs

cover only 0.1% of oceans, but host 25% of all marine life.

protect coastlines, generate revenue (tourism, fishing, ...).

in 50 years, we managed to kill half of them.

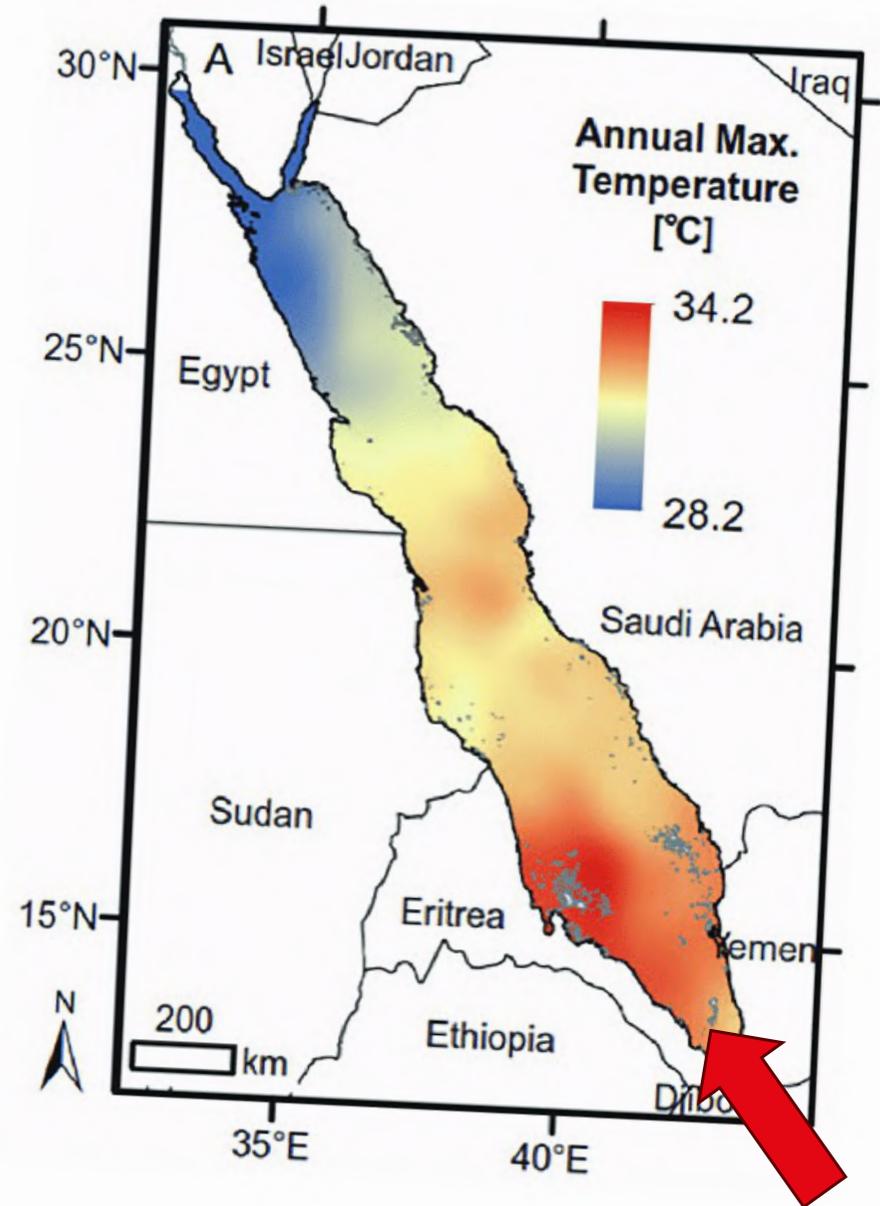
Still, in some places corals resist.

- Red sea corals are much more resistant to heat
- We need to understand *why*
- We need to map and monitor, to better follow the evolution of reefs' health and protect them



Still, in some places corals resist.

- Red sea corals are much more resistant to heat
- We need to understand *why*
- We need to map and monitor, to better follow the evolution of reefs' health and protect them





How do we monitor a large ecosystem like that?

SOTA - Photo quadrats

- Current method used in international efforts
- Easy to deploy
- Large interpretation effort
- Recognition can be scaled with ML, but still it remains snapshots of the diversity of reefs (and needs technical know-how)

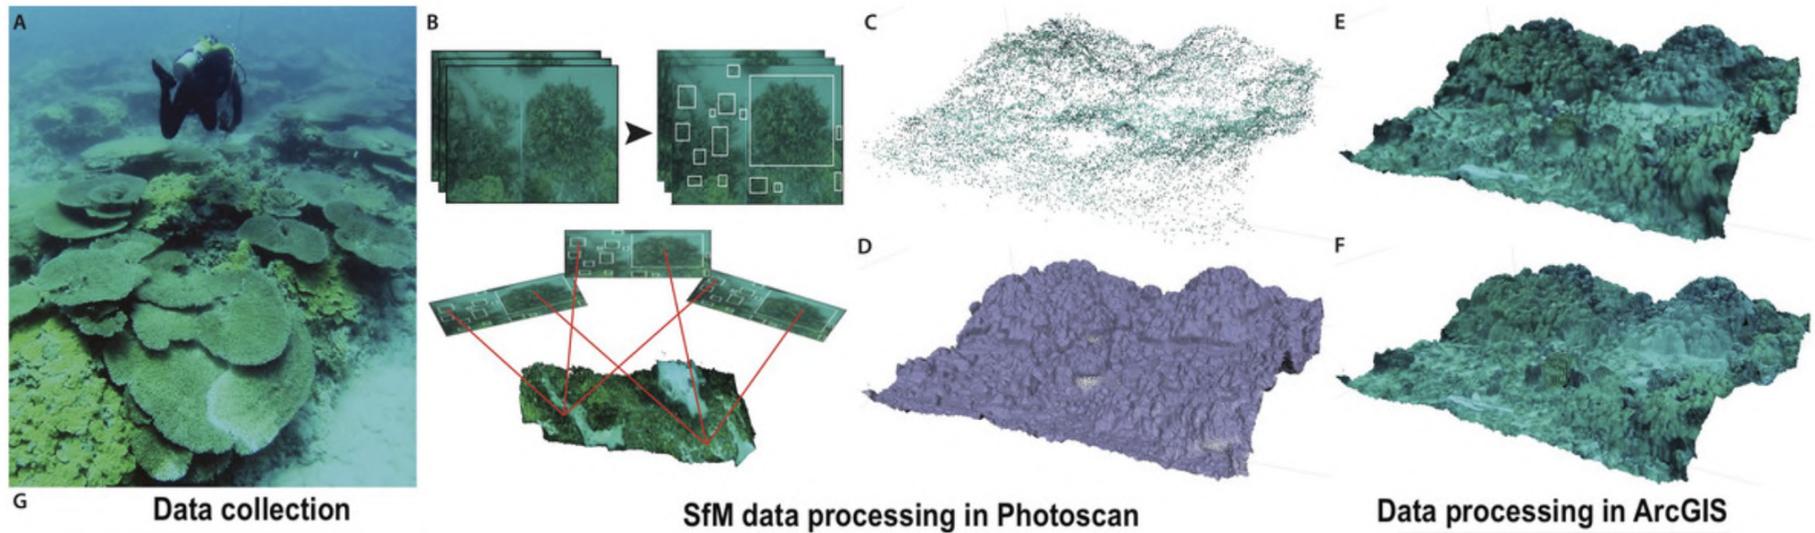


First attempts at scaling

28 x 6 m plot

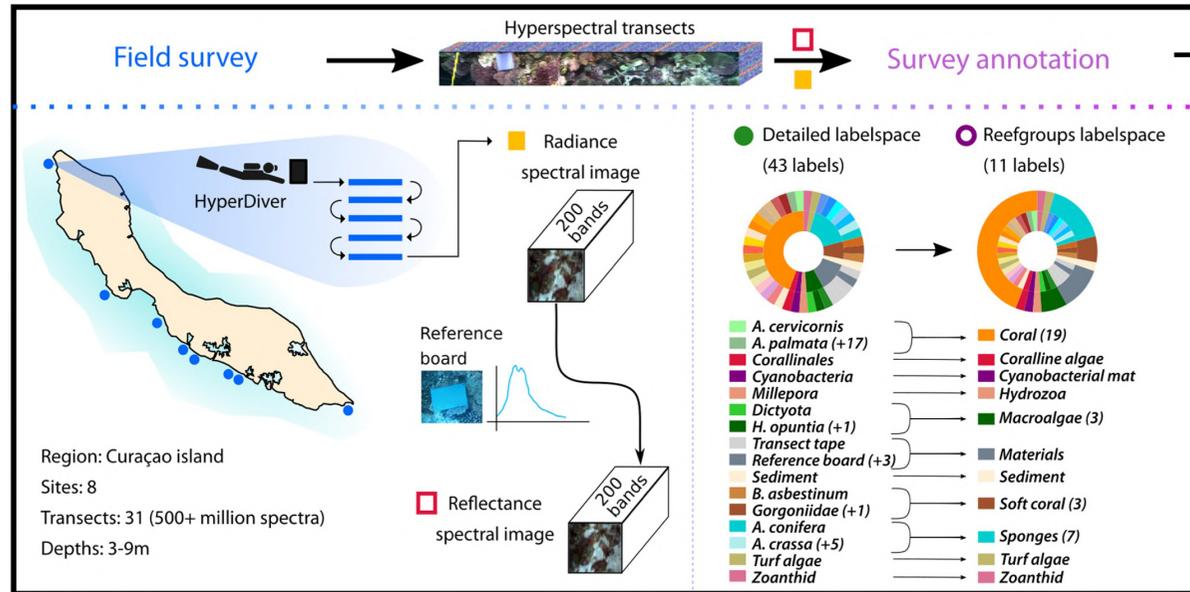
1h with proprietary software

20h manual work to extract information



Unique setups: great results, but difficult to apply elsewhere

- Published models often rely on complex setups, very expensive
- E.g. hyperspectral sensors



Schürholz and Chennu, Methods in Ecology and Evolution, 2022

Our bet: affordable setups



- Scalable to other reefs
- Easy to acquire / replace
- Can train local communities



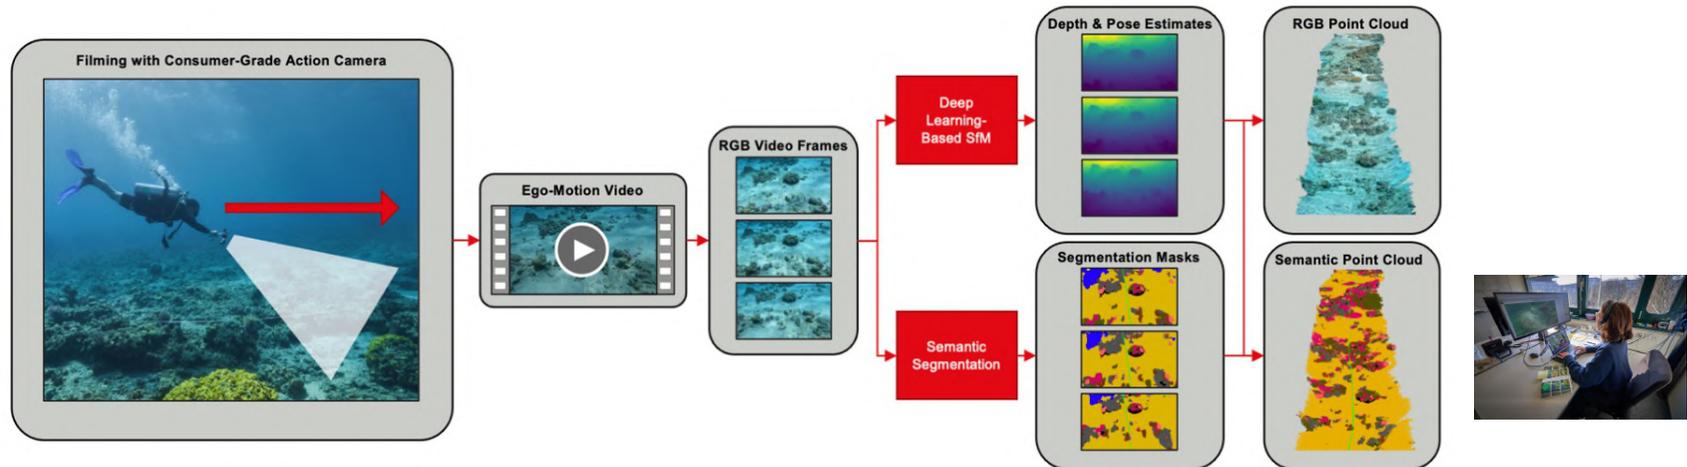
Mark I: March 2022 – Isreal / Jordan

Mark II: August 2022 – Djibouti
November 2023 – Djibouti
April 2024 – Jordan
November 2024 - Eritrea



Enabling scalable reef monitoring: Open source, fast, large scale.

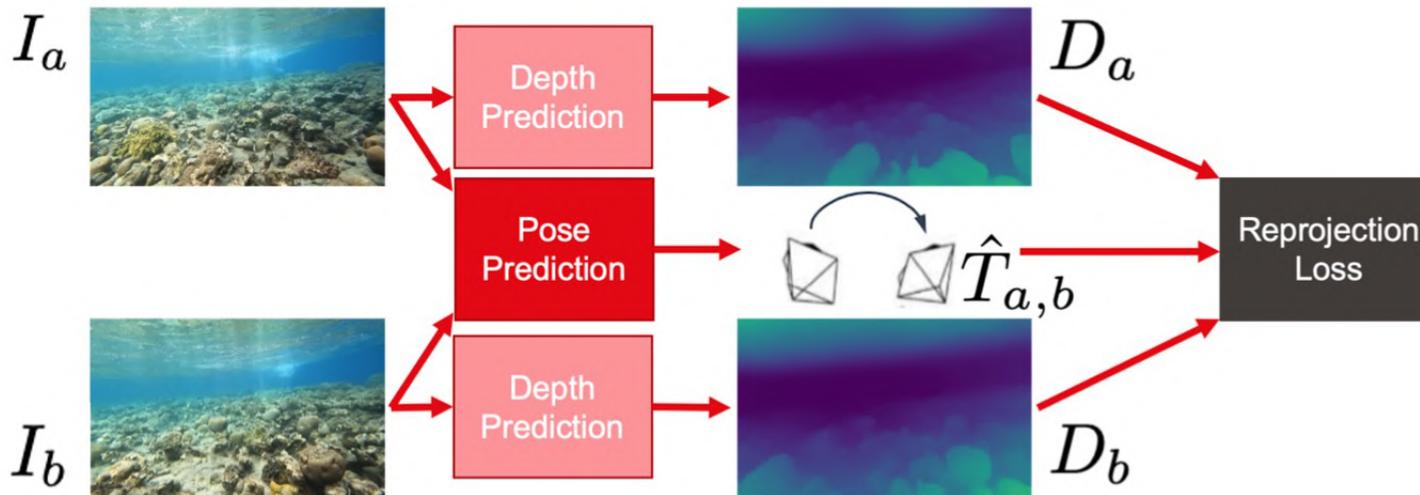
- With custom-built, affordable imaging setup
- A model that works on videos, leveraging 2 tasks
- Once trained, 3D reconstructs a 100m transect in playing video time
- Tested in Isreal, Jordan and Djibouti in 2022



J. Sauder, G. Banc-Prandi, A. Meibom, and D. Tuia. Scalable semantic 3d mapping of coral reefs with deep learning. *Methods in Ecology and Evolution*, 2024. <https://arxiv.org/abs/2309.12804>

Pose and depth estimation

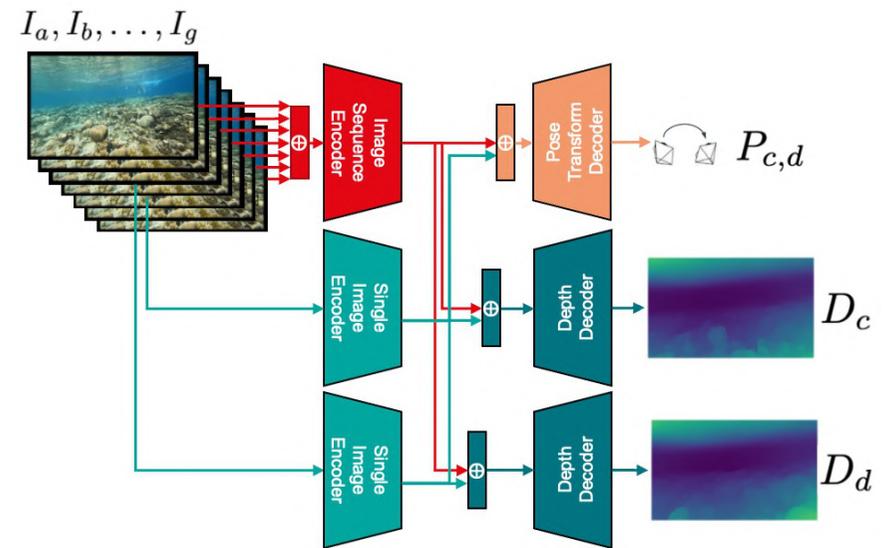
- Encoders based on ResNet-34
- Can create the 3D map at 18 frames per second



J. Sauder, G. Banc-Prandi, A. Meibom, and D. Tuia. Scalable semantic 3d mapping of coral reefs with deep learning. *Methods in Ecology and Evolution*, 2024. <https://arxiv.org/abs/2309.12804>

Pose and depth estimation

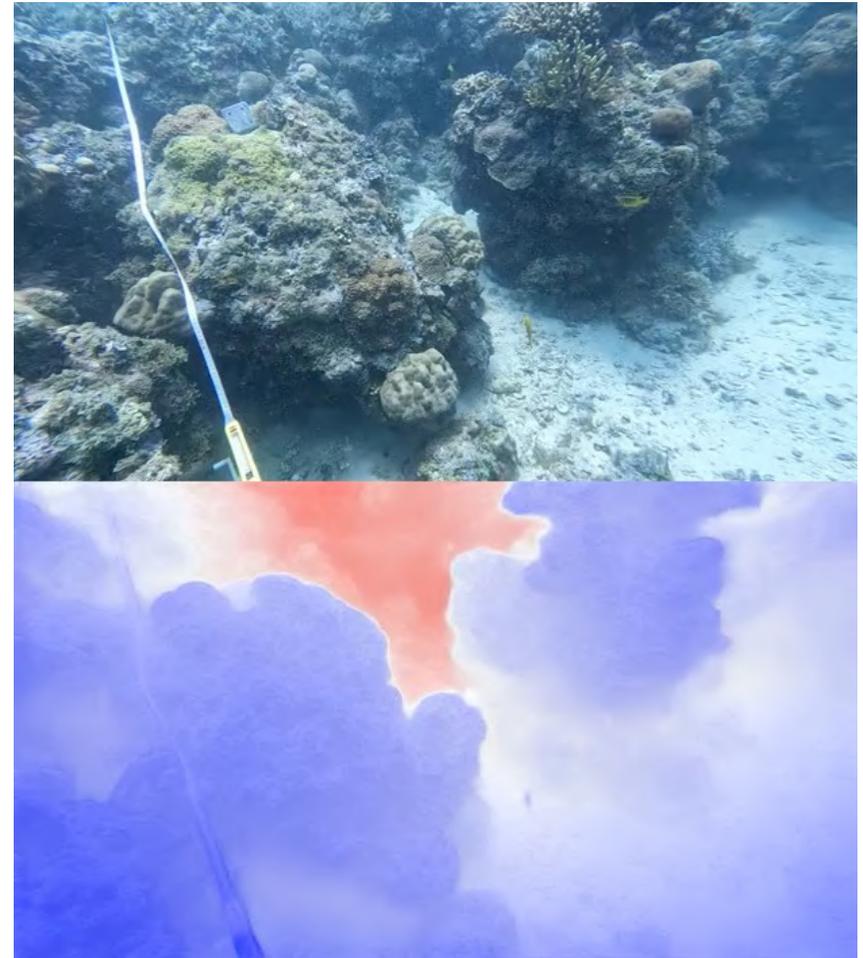
- We enforce temporal consistency by encoding 7 frames at a time
- 3 before and 3 after
- Each decoder uses the frame + the sequence features



J. Sauder, G. Banc-Prandi, A. Meibom, and D. Tuia. Scalable semantic 3d mapping of coral reefs with deep learning. *Methods in Ecology and Evolution*, 2024. <https://arxiv.org/abs/2309.12804>

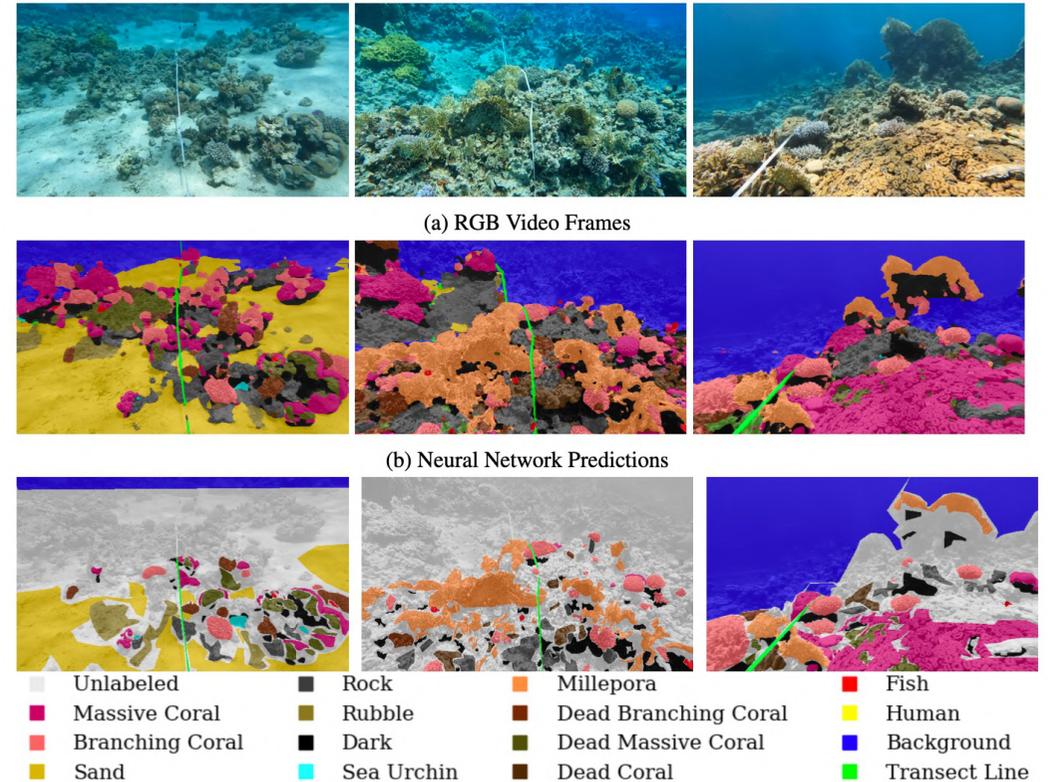
Pose and depth estimation

- We enforce temporal consistency by encoding 7 frames at a time
- 3 before and 3 after
- Each decoder uses the frame + the sequence features



Semantic segmentation

- Unet with ResNeXt backbone
- ~85% accurate in Jordanian and Israeli reefs

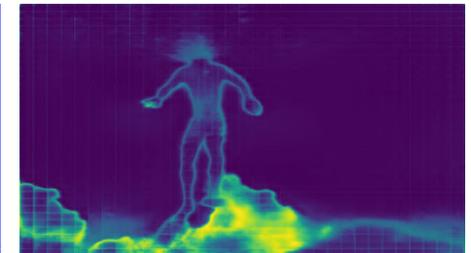
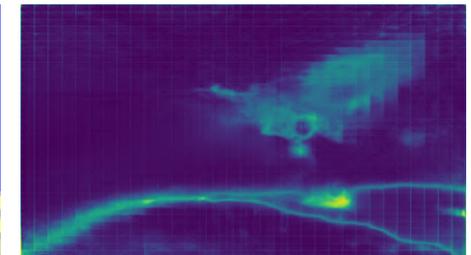
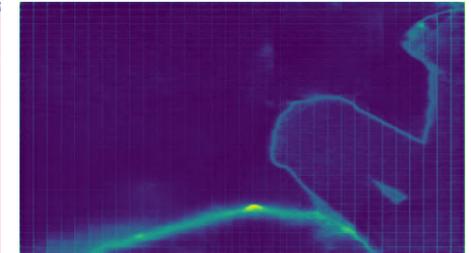
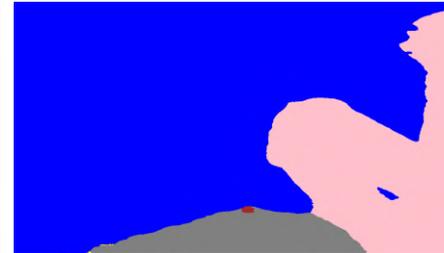


J. Sauder, G. Banc-Prandi, A. Meibom, and **D. Tuia**. Scalable semantic 3d mapping of coral reefs with deep learning. *Methods in Ecology and Evolution*, 2024. <https://arxiv.org/abs/2309.12804>

Learning to detect unwanted classes

Used to remove unwanted classes prior to 3D reconstruction

- Diver body
- Fishes
- Far away pixels



Video frame

Segmentation

Uncertainty

Learning to detect unwanted classes

Used to remove unwanted classes prior to 3D reconstruction

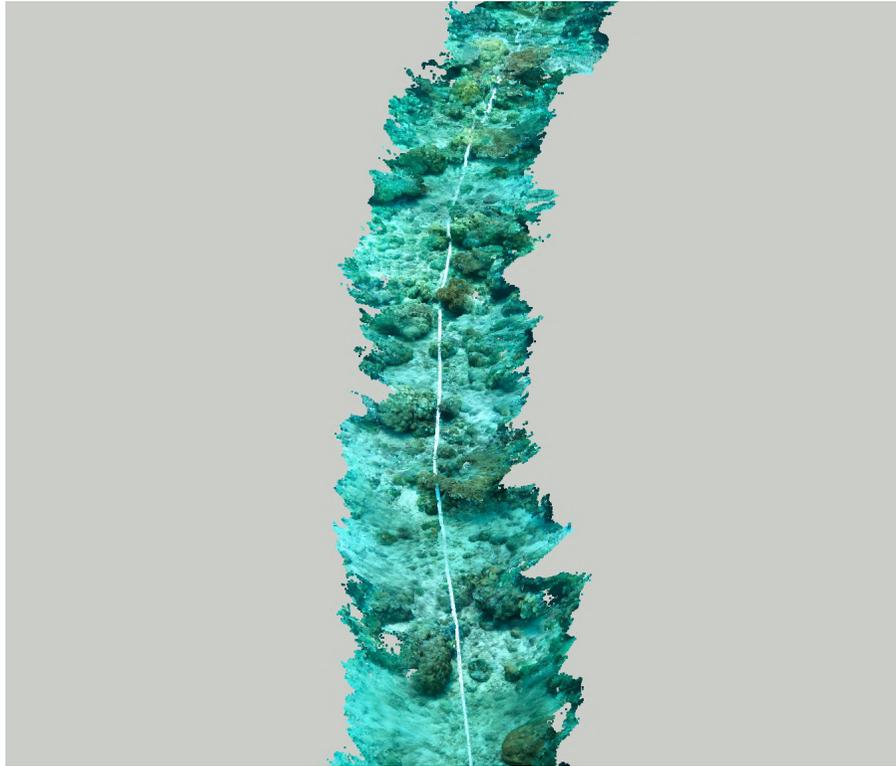
- Diver body
- Fishes
- Far away pixels



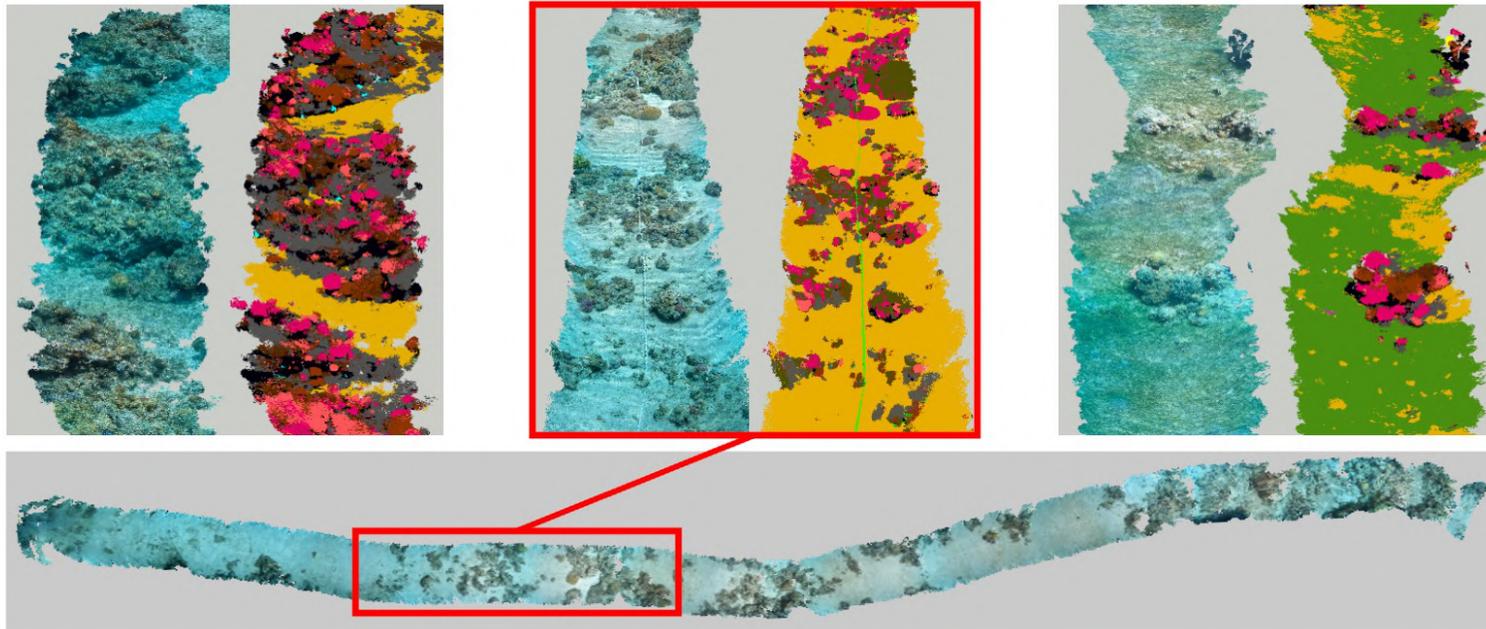


- | | | | | |
|--------|--------------|------------------|---------------------|-----------------------------|
| fish | anemone | transect line | other coral dead | table acropora alive |
| sand | seagrass | branching dead | meandering alive | other coral bleached |
| dark | millepora | acropora alive | pocillopora alive | unknown hard substrate |
| clam | background | transect tools | other coral alive | algae covered substrate |
| human | turbinaria | branching alive | branching bleached | massive/meandering dead |
| trash | sea urchin | meandering dead | table acropora dead | massive/meandering alive |
| rubble | sea cucumber | stylophora alive | meandering bleached | massive/meandering bleached |
| sponge | other animal | | | |

The multitask model allow us to create reliable 3D reconstructions of the reef



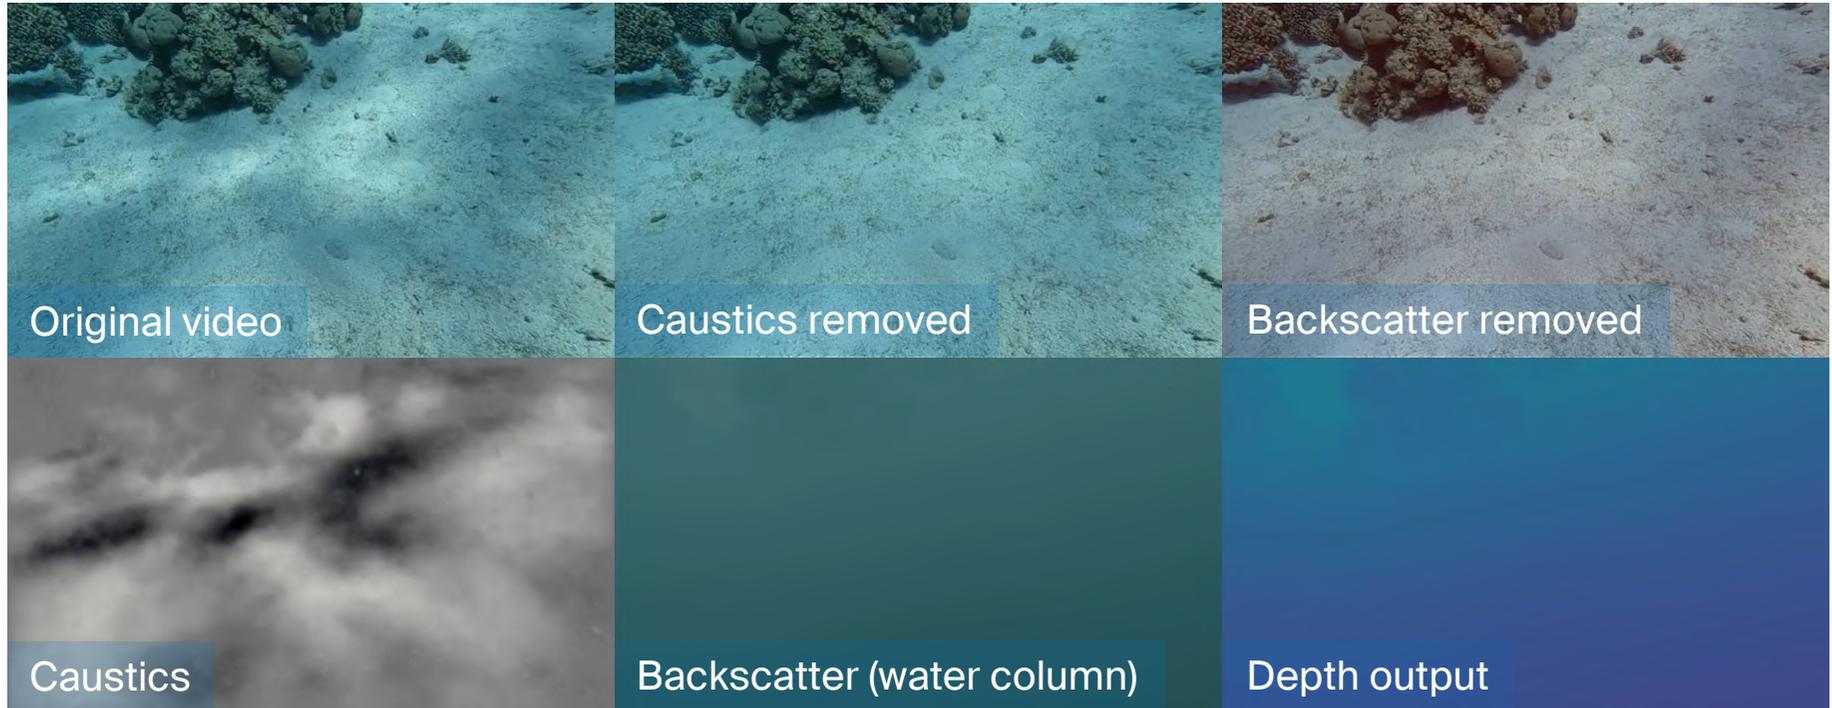
Mapping entire dive sites (here: 100m long)



- Dead Branching Coral
- Dead Massive Coral
- Dead Coral
- Branching Coral
- Massive Coral
- Sand
- Rock
- Seagrass
- Dark
- Sea Urchin
- Transect Line
- Rubble

Moving forward

New techniques



Moving forward

Toward real monitoring



Jordan 2024



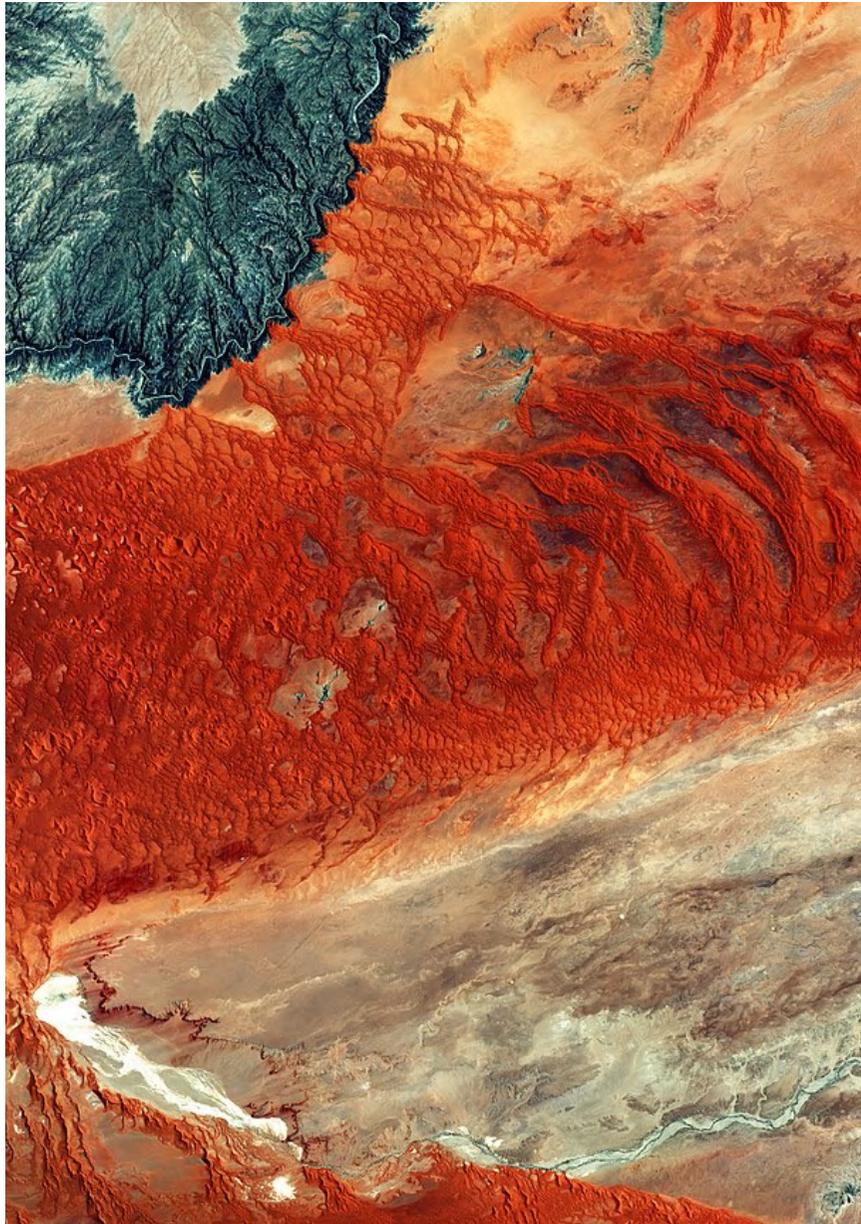
Djibouti 2023



Photo Credits: Lwimages, G. Banc-Prandi

EPFL

eceo
EPFL



My view on Remote sensing and AI

Advance remote sensing science to
monitor and protect Earth

Interface disciplines and approaches

Bring new, open tools making EO science
accessible to anyone



But more importantly

Fall in love with the problem, not the technology!

Work with domain specialists, be interdisciplinary!

Earth is under pressure. Models use energy, emit CO₂. Think before computing!

