Deep learning for environmental sciences

Jan Dirk Wegner
EcoVision Lab
Universität Zürich & ETH Zürich

Data-driven methods

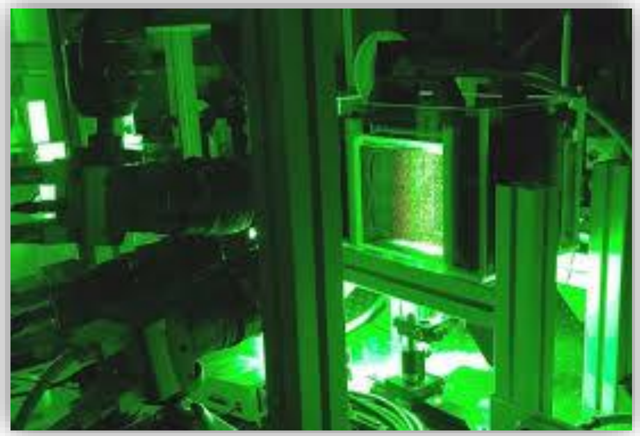

Particle tracking, fluid flow and turbulence

Glacier retreat and discharge

Landslide risk assessment

- ✓ Provide additional evidence if modelling all functional relationships for solving a task is hard and a lot of data are available
- ✓ Complementary to physics-based modelling and simulation

Deep learning: what's new about it?

Most of the "heavy lifting" in here.
Final performance only as good as the feature set.

Traditional machine learning:

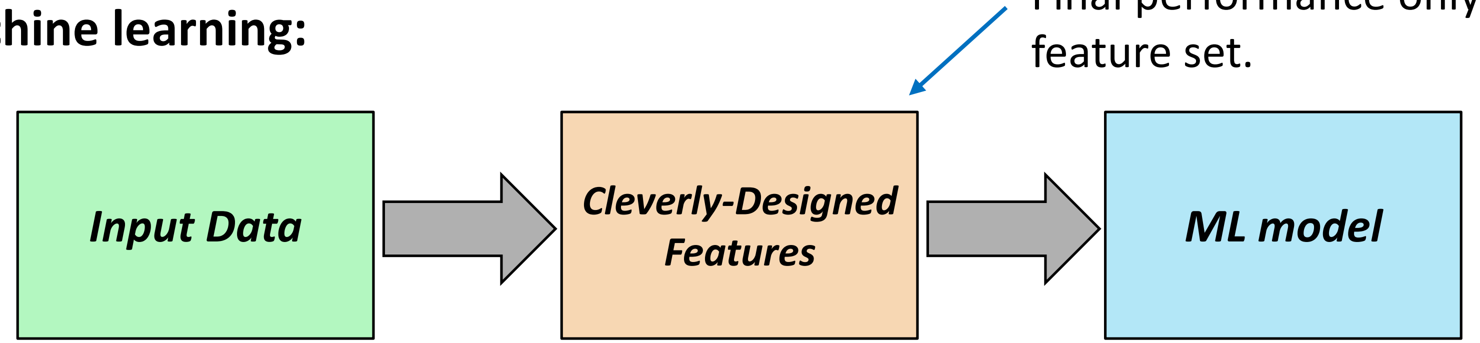

Deep learning:

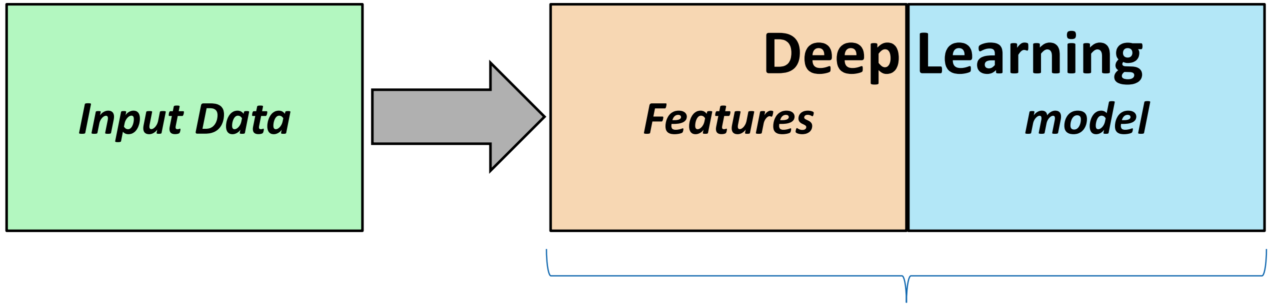

Features and model learned together, mutually reinforcing

slide credit:
Michele Catasta

a. Biodiversity

b. Vegetation analysis

a. Biodiversity

b. Vegetation analysis

Biodiversity

- **Goal:** Automated species mapping

- **Idee:** Combine crowd-sourcing, deep learning and remote sensing

crowd-sourced Biodiversität

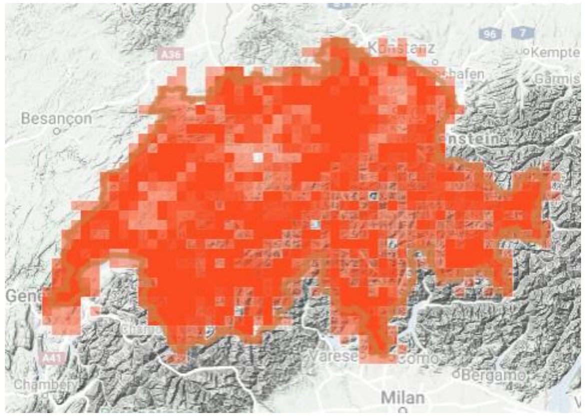

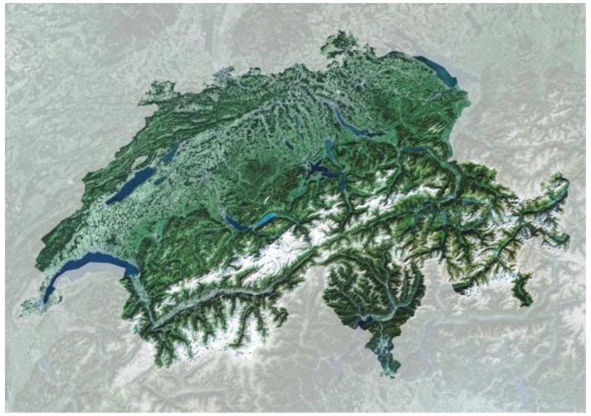

Satellitenbilder

App-based observation of species

User (i.e., citizen scientists) map species via mobile phone app:

- Experts recognize species and enter into app directly (minority of app-users)
- Unexperienced user (majority of users) acquires photos and contextual evidence, which are analysed automatically via deep learning

Idea: Simultaneous reasoning across photos and additional evidence:

- spatial and seasonal context of acquisition place
- hierarchical relationship of species

florapp
iNaturalist
Pl@ntNet

(Image credits: iNaturalist, Freepik.com)

Properties of app-based crowd-sourcing

- Few species reported very often, many species only rarely ➔ caused by systemic effects (visually attractive and easily accessible species over-represented,…)

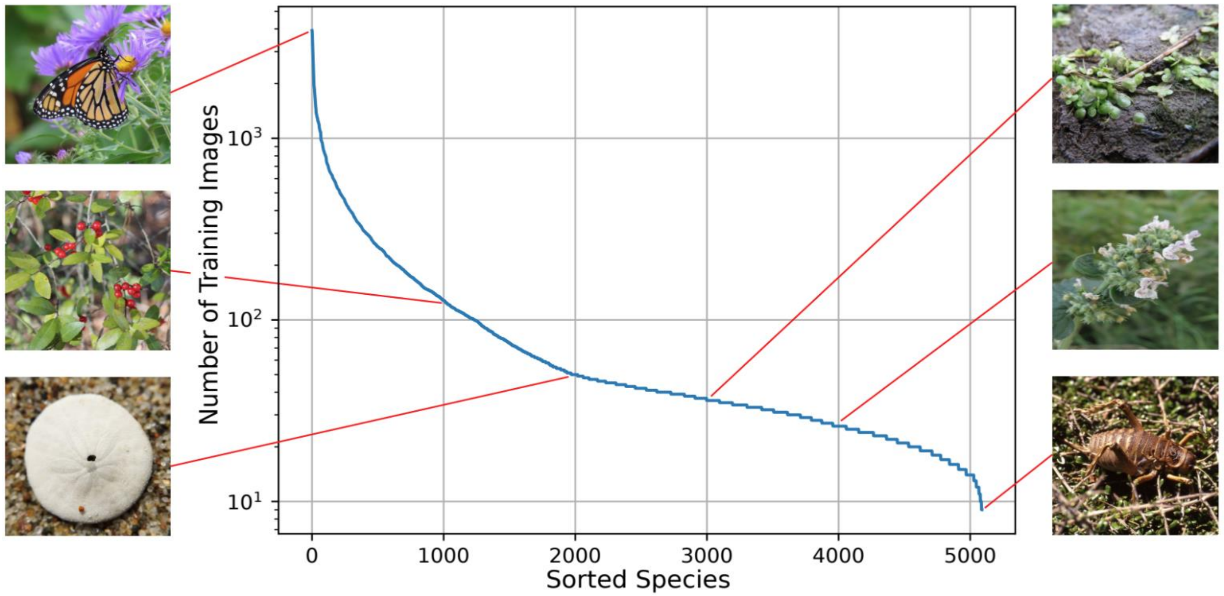

- Usually only one species reported per location although many others exist

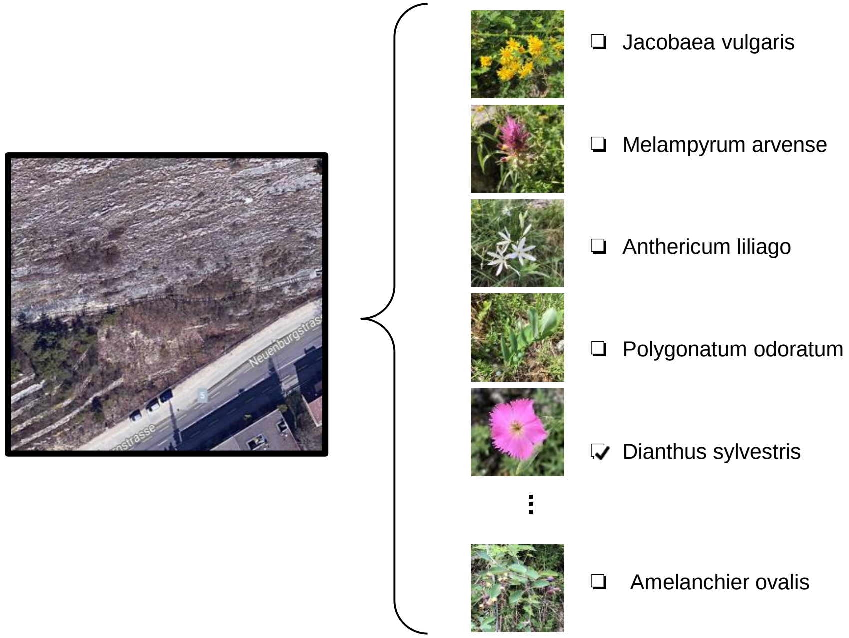

Method: contextual evidence

Contextual evidence of photo is recorded in App:

- Geographic coordinates and altitude
- Date

Joint reasoning across photos and contextual evidence:

- Product of image interpretation results from ResNet and context-classifier
- Approximation of species probability distribution

Identification of species in photo

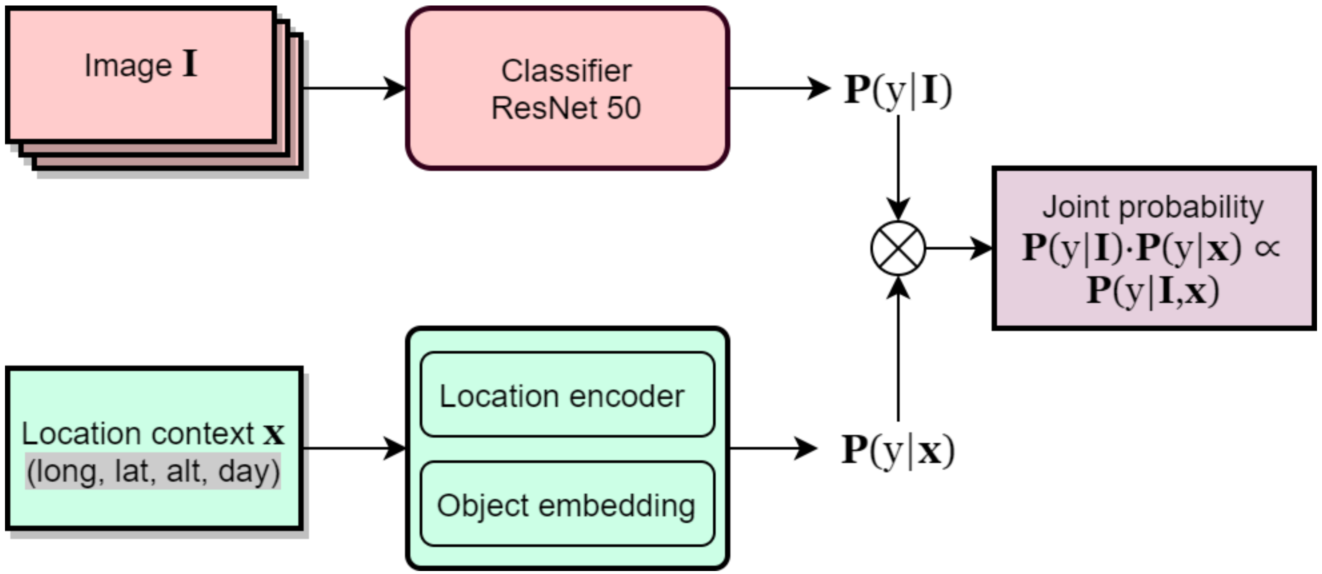

Information about seasonality, geographic distribution, altitude

Method: Species hierarchy

Identification of species in photo

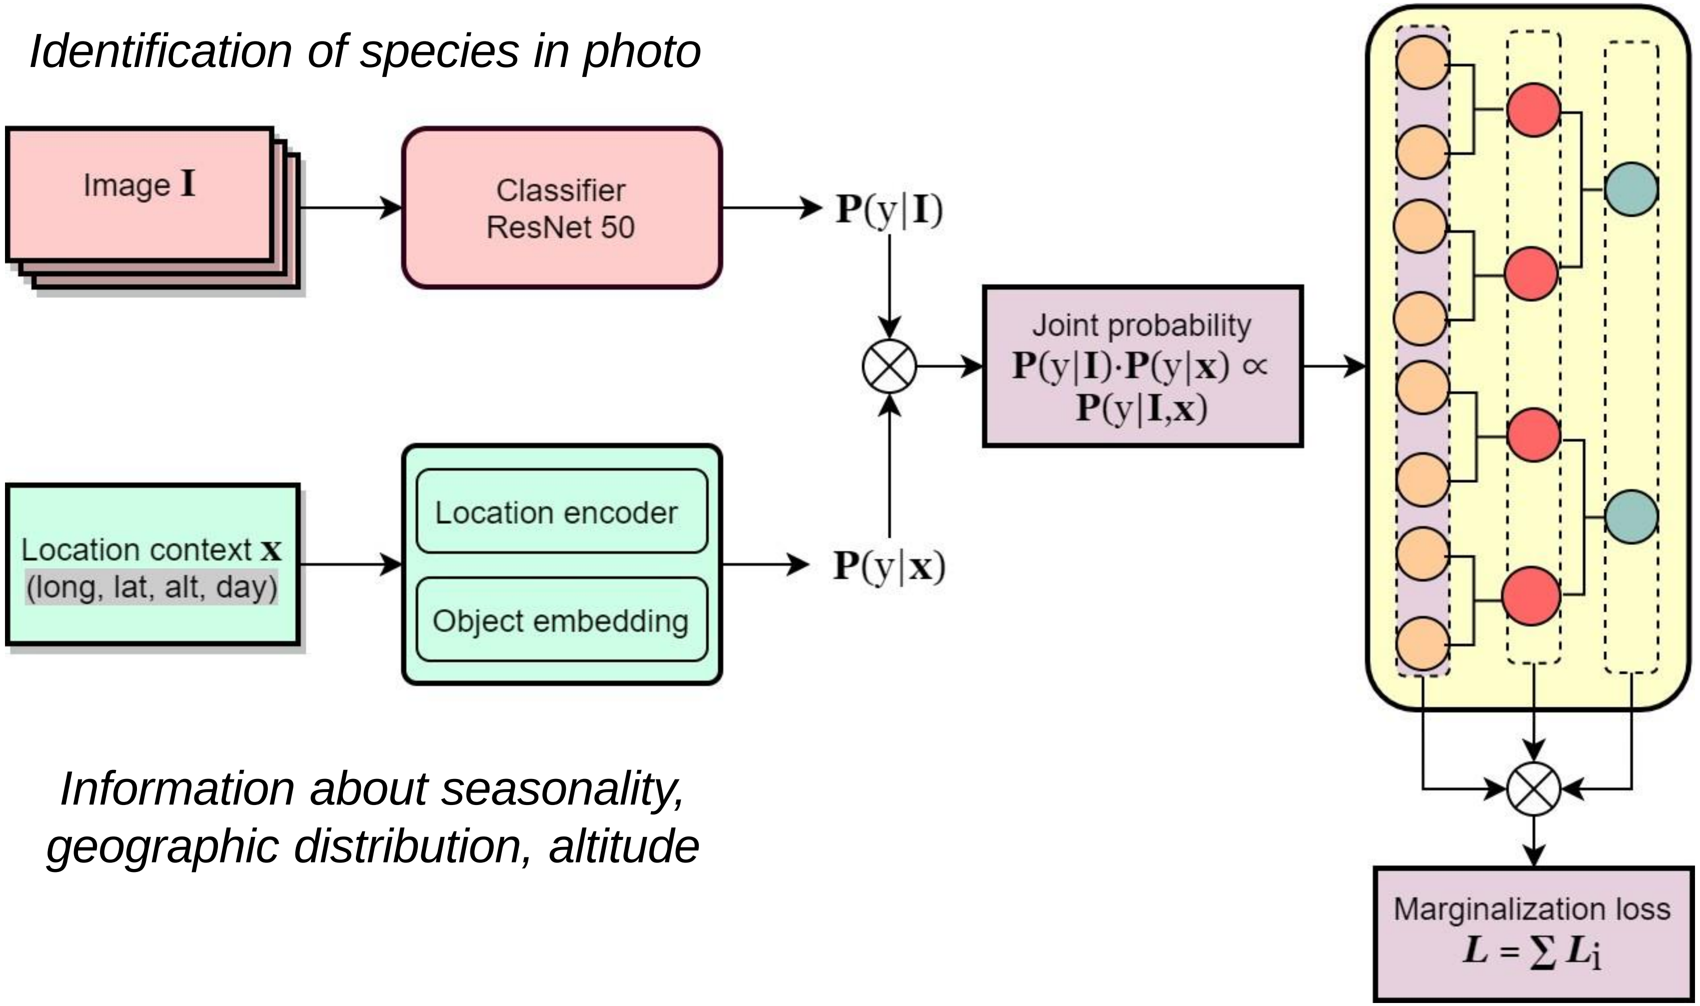

Information about seasonality, geographic distribution, altitude

Rare species benefit from other species with similarities according to visual appearance and contextual evidence

Preliminary results: context information and species hierarchies

***without** context and hierarchies*

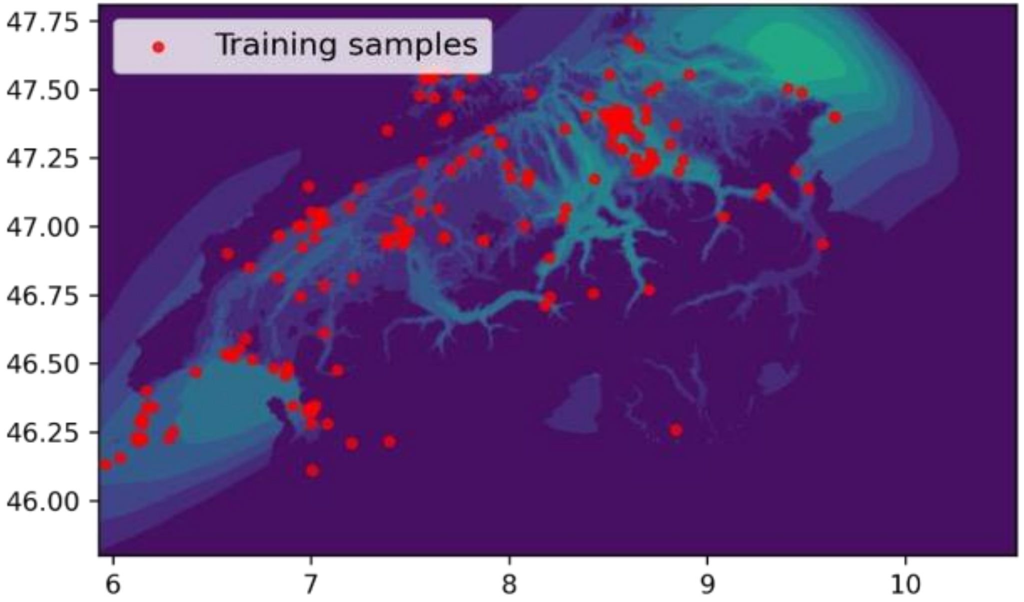

***with** context and hierarchies*

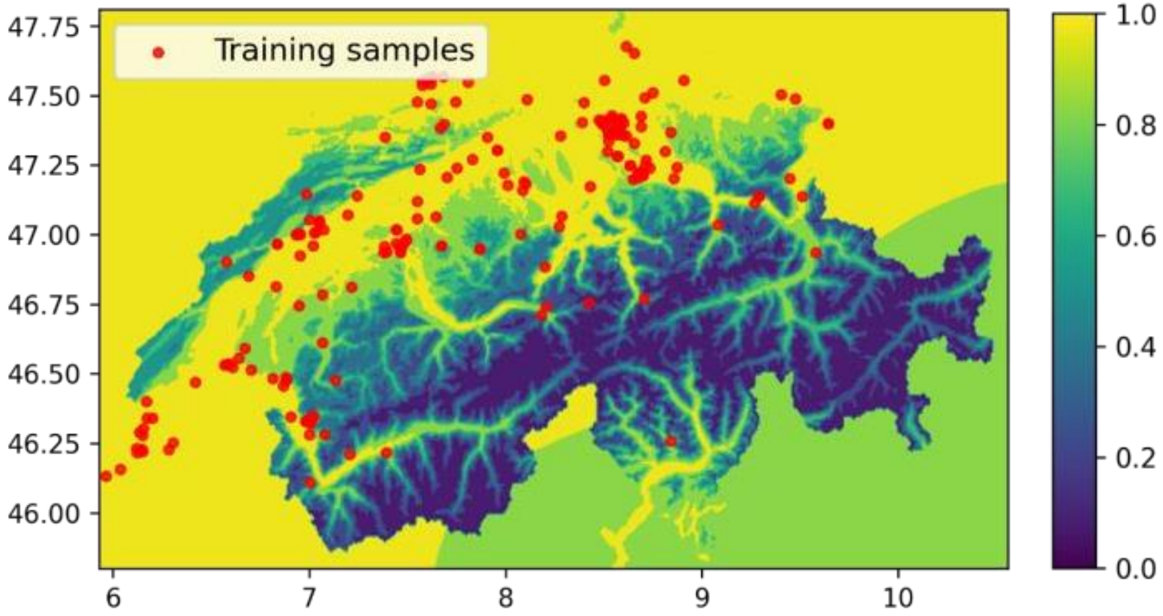

Level	Species	Genus	Family	Order	Class	Phylum
Top-1	**68.09%**	**74.90%**	**80.86%**	**82.68%**	**95.58%**	**99.73%**
Top-3	82.78%	86.73%	91.89%	93.59%	99.86%	--*
Top-5	87.47%	90.35%	94.76%	96.42%	99.98%	--

a. Biodiversity

b. Vegetation analysis

Vegetation analysis at global scale

- **Goal**: Dense, 20 meter resolution vegetation parameters at global scale that are frequently updated

- **Idea**: use single satellite images to predict vegetation height (and later more variables like biomass)

© Mighty Earth

ESA satellites Sentinel-2

- Constellation of 2 satellites
- Revisit cycle 3-5 days
- 290 km swath width
- 13 spectral bands tailored for vegetation analysis
- Ground sampling distance 10, 20 and 60 meters depending on spectral channel

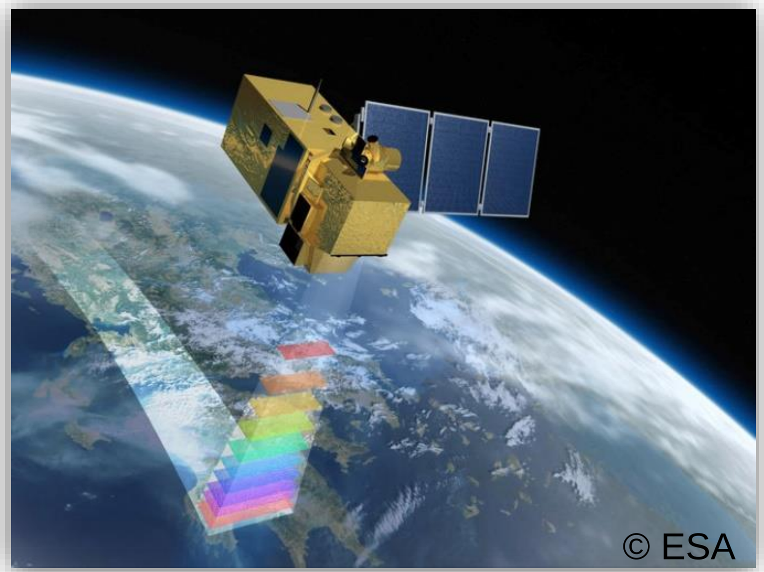

Method: network architecture

- Avoid down- or up-sampling: stride 1, no max-pooling
- 18 identical, separable convolution (SepConv) blocks do not only learn spectral features that correlate with canopy height, but also spatial context and texture features.

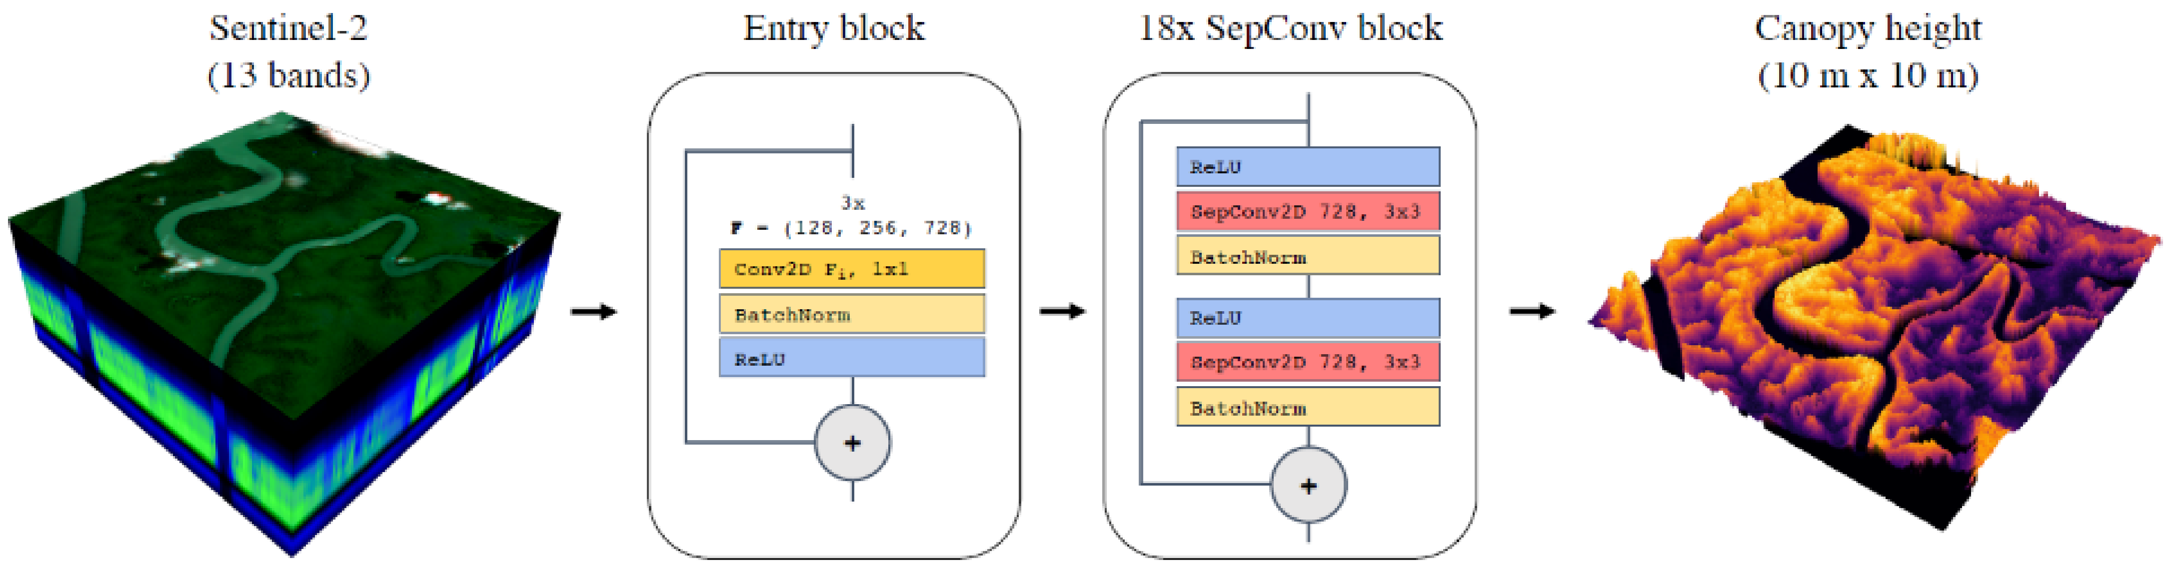

Lang, N., Schindler, K., Wegner, J.D.: Country-wide high-resolution vegetation height mapping with Sentinel-2, Remote Sensing of Environment, 2019, vol. 233, article 111347.

Method: loss function

- Euclidean loss function for regression of continuous height values
- L2-penalty term on model parameters («weight decay») as regularizer

ground truth at pixel *i*

model parameters

$$Loss = \frac{1}{N}\sum_{i=1}^{N}\left(f(x_i) - y_i\right)^2 + \lambda\frac{1}{W}\sum_{j=1}^{W}(w_j)^2$$

number of samples (pixels)

pixel intensities

Vegetation height predictions at pixel *i*

number of model parameters

Experiments

Switzerland

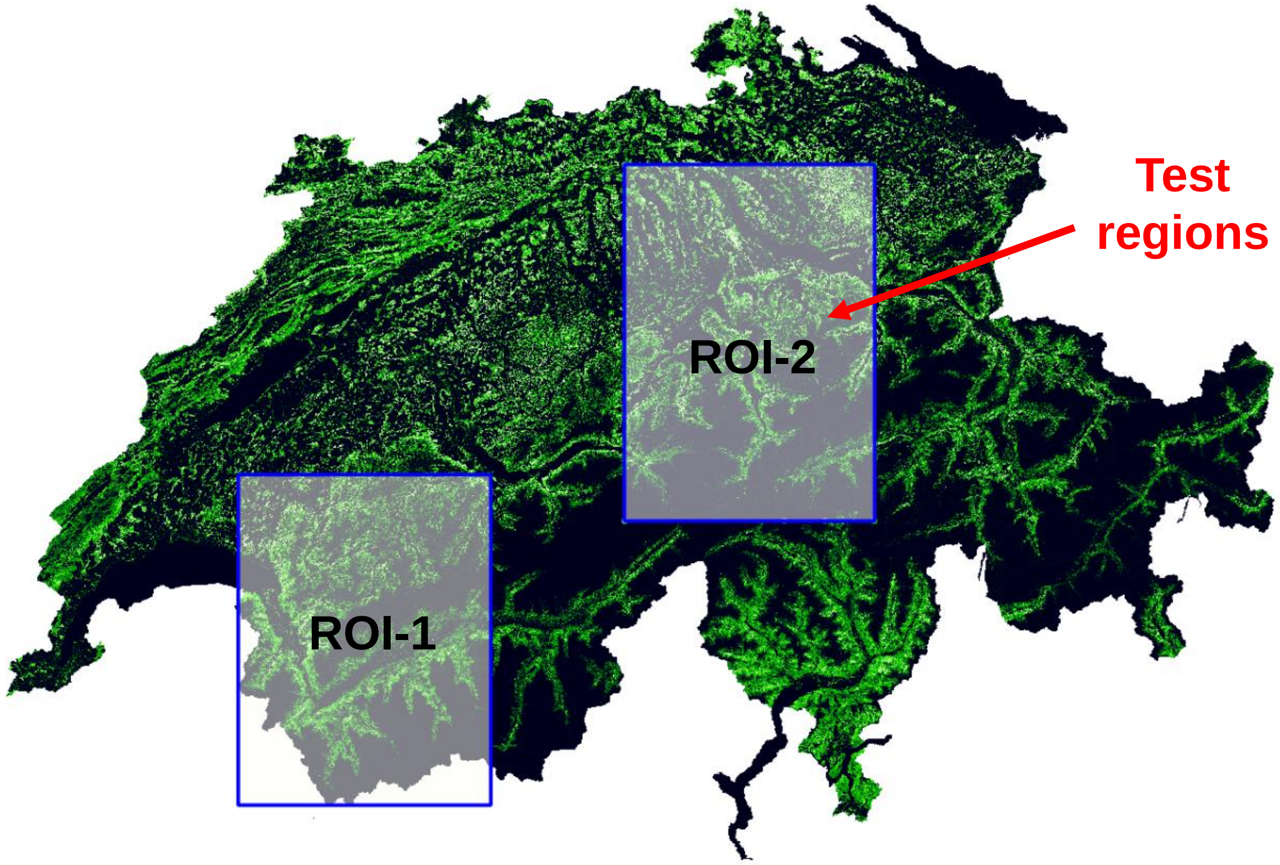

Test regions

Gabon

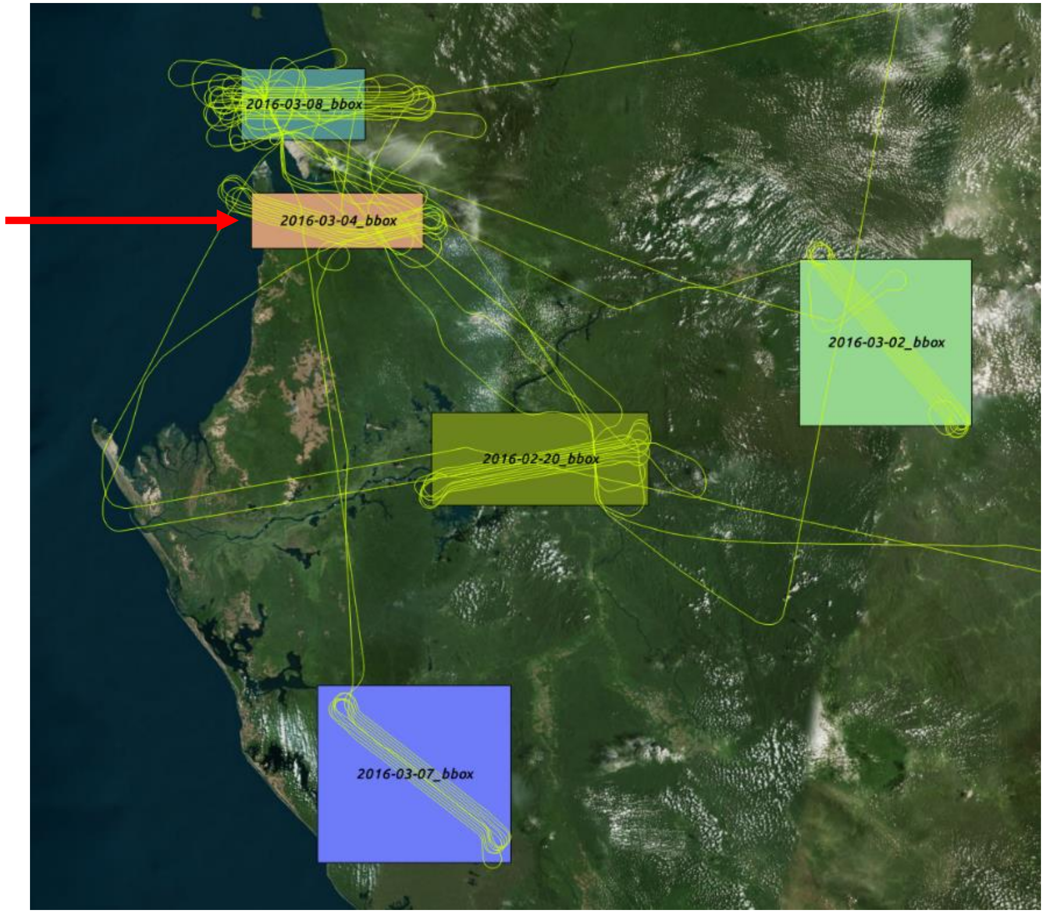

Result: Gabon

NASA: 1km

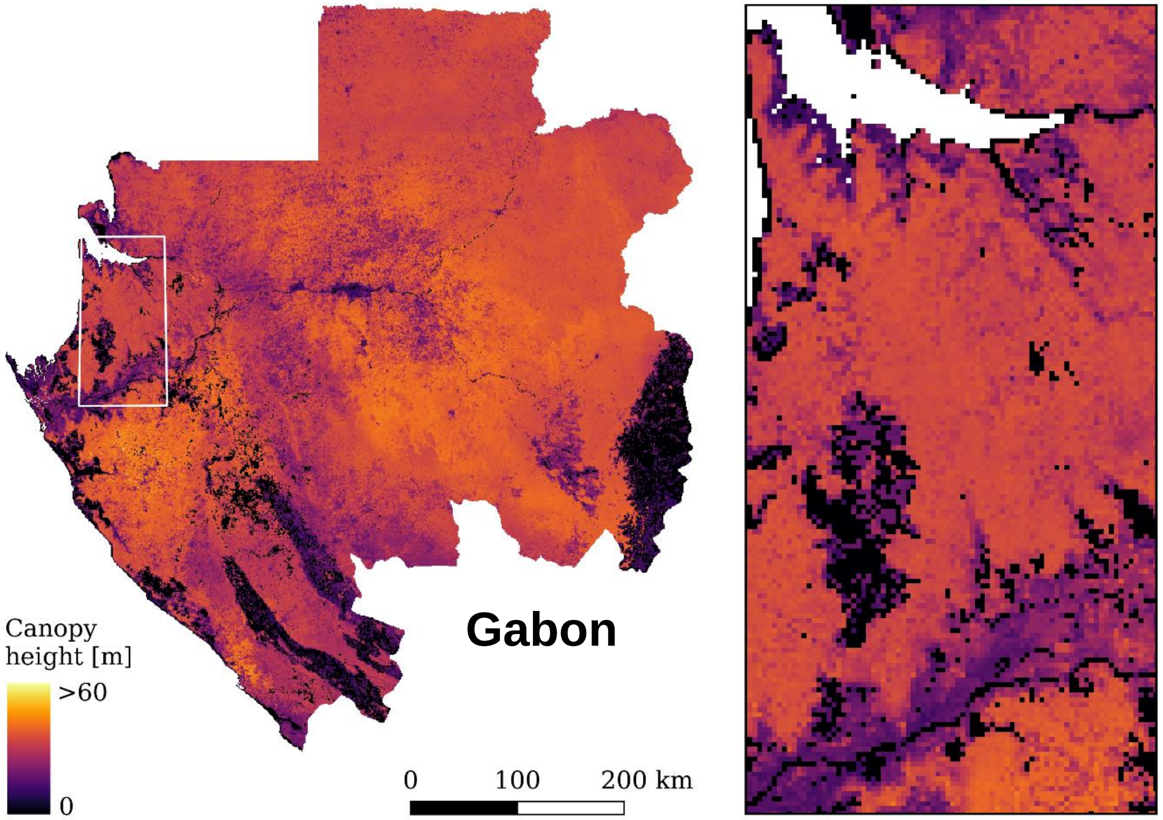

(Simard et al., 2011)

Ours: 10m

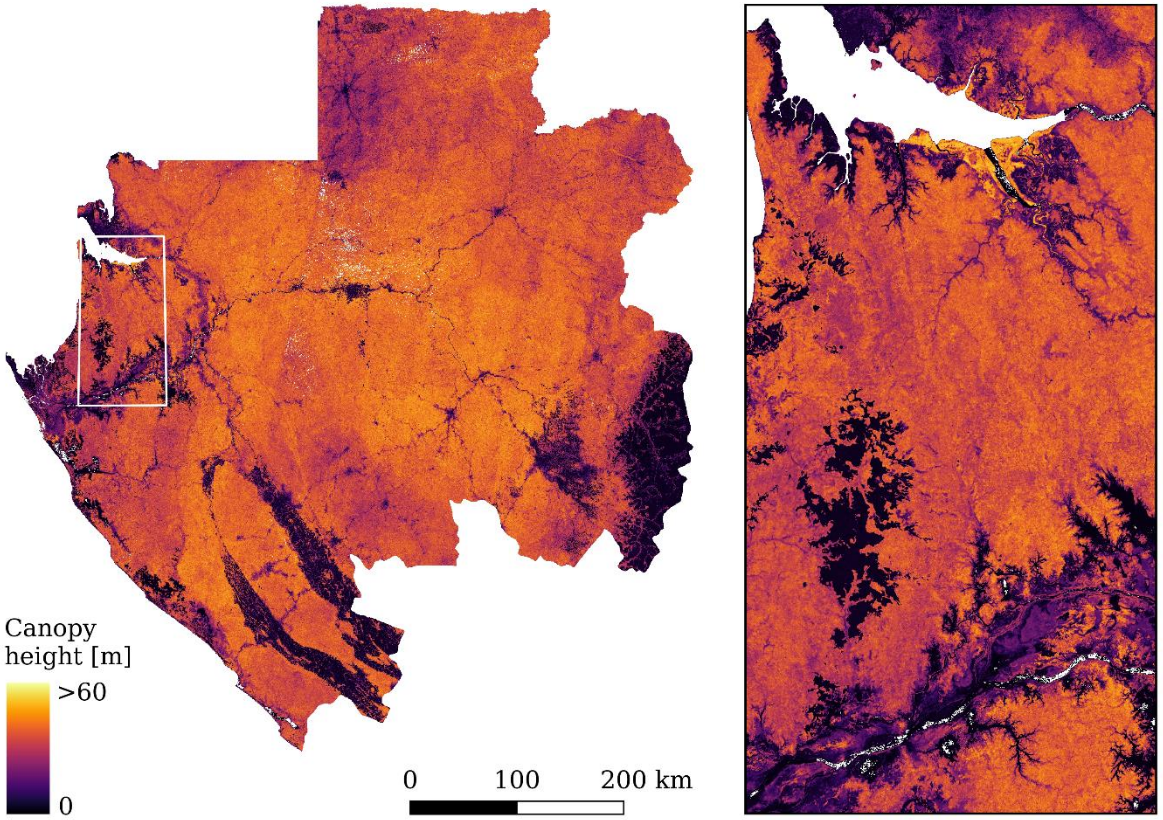

*MAE **±4.3 m** (for vegetation heights 0 to 60m)*

Result: Switzerland

NASA: 1km

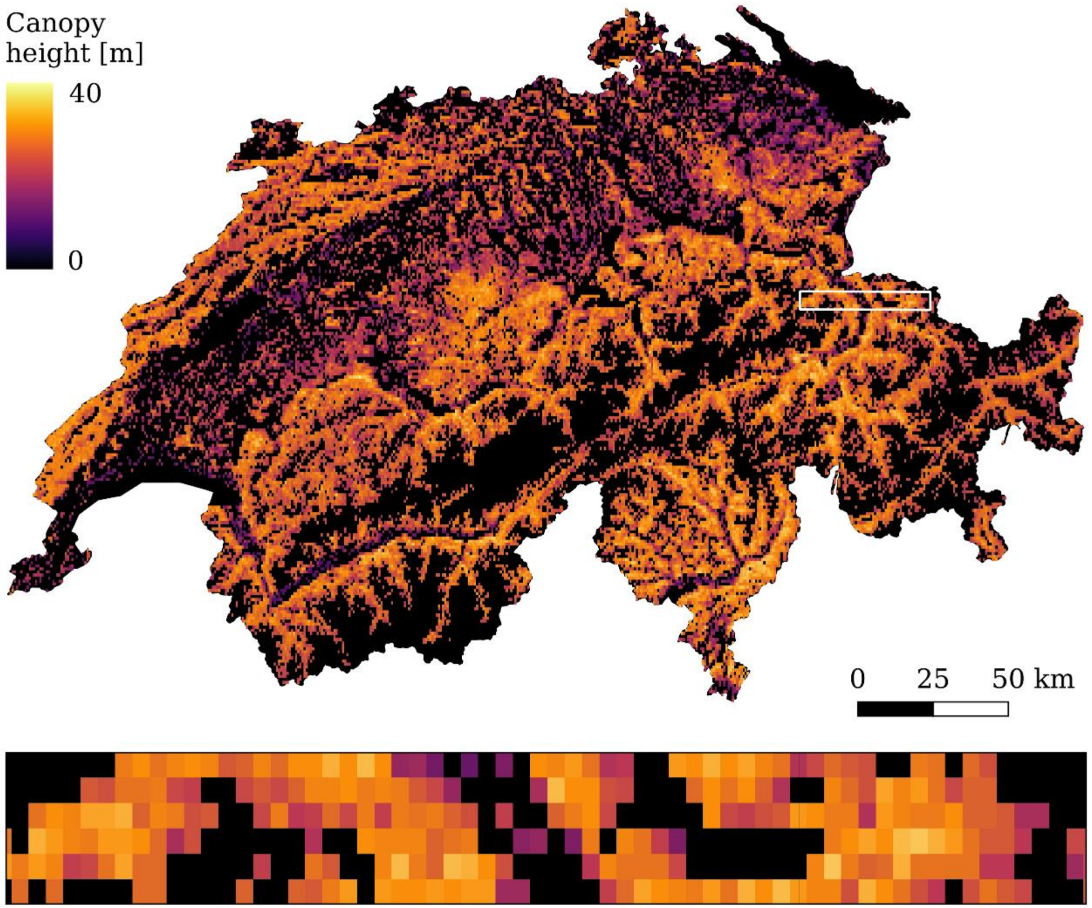

(Simard et al., 2011)

Ours: 10m

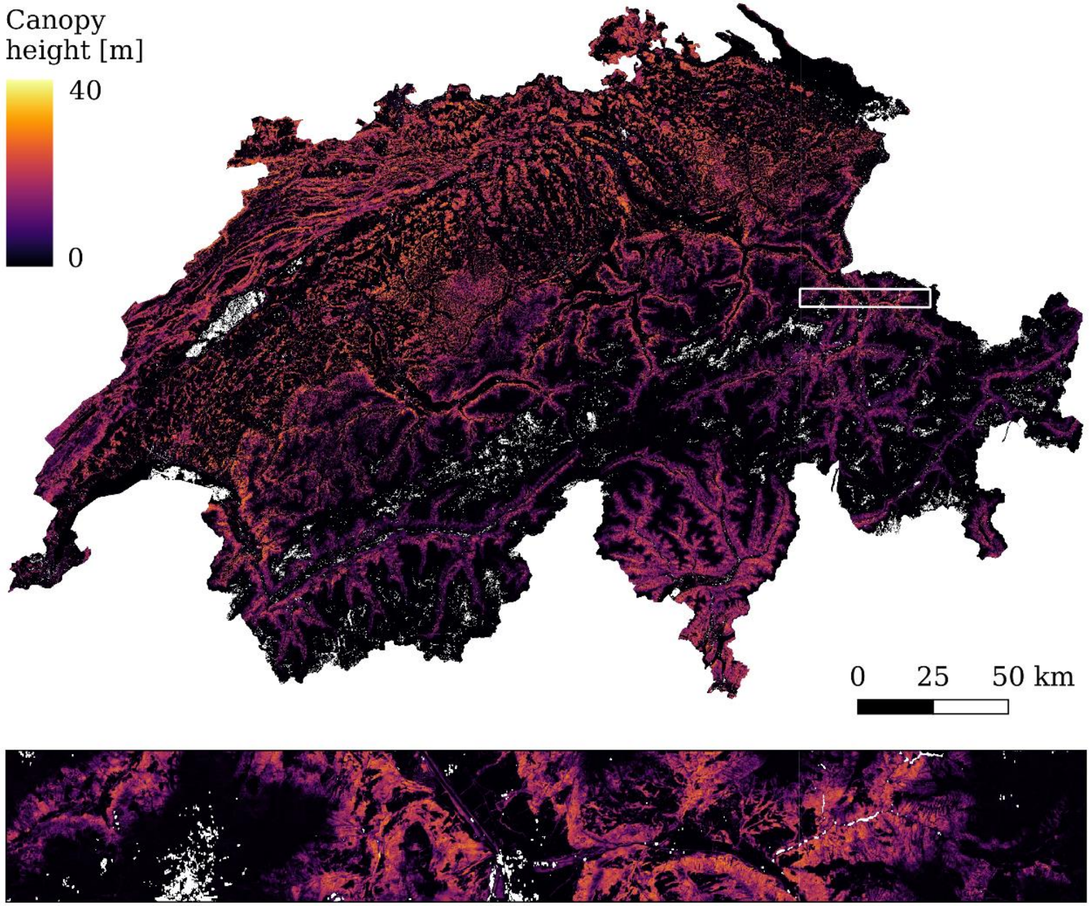

MAE ***±1.7 m*** *(for vegetation heights 0 to 40m)*

Universität Zürich[UZH] ETH zürich EcoVision

Ongoing work: scale globally

- **Goal**: global vegetation height, biomass and HCS map with 20 meter resolution and near-realtime updates

- Collaboration with **NASA GEDI** team und **Amazon Research**

- **Idea:** train deep learning model on *full-waveform spaceborne Laserscanning* points of NASA GEDI mission und interpolate with satellite data

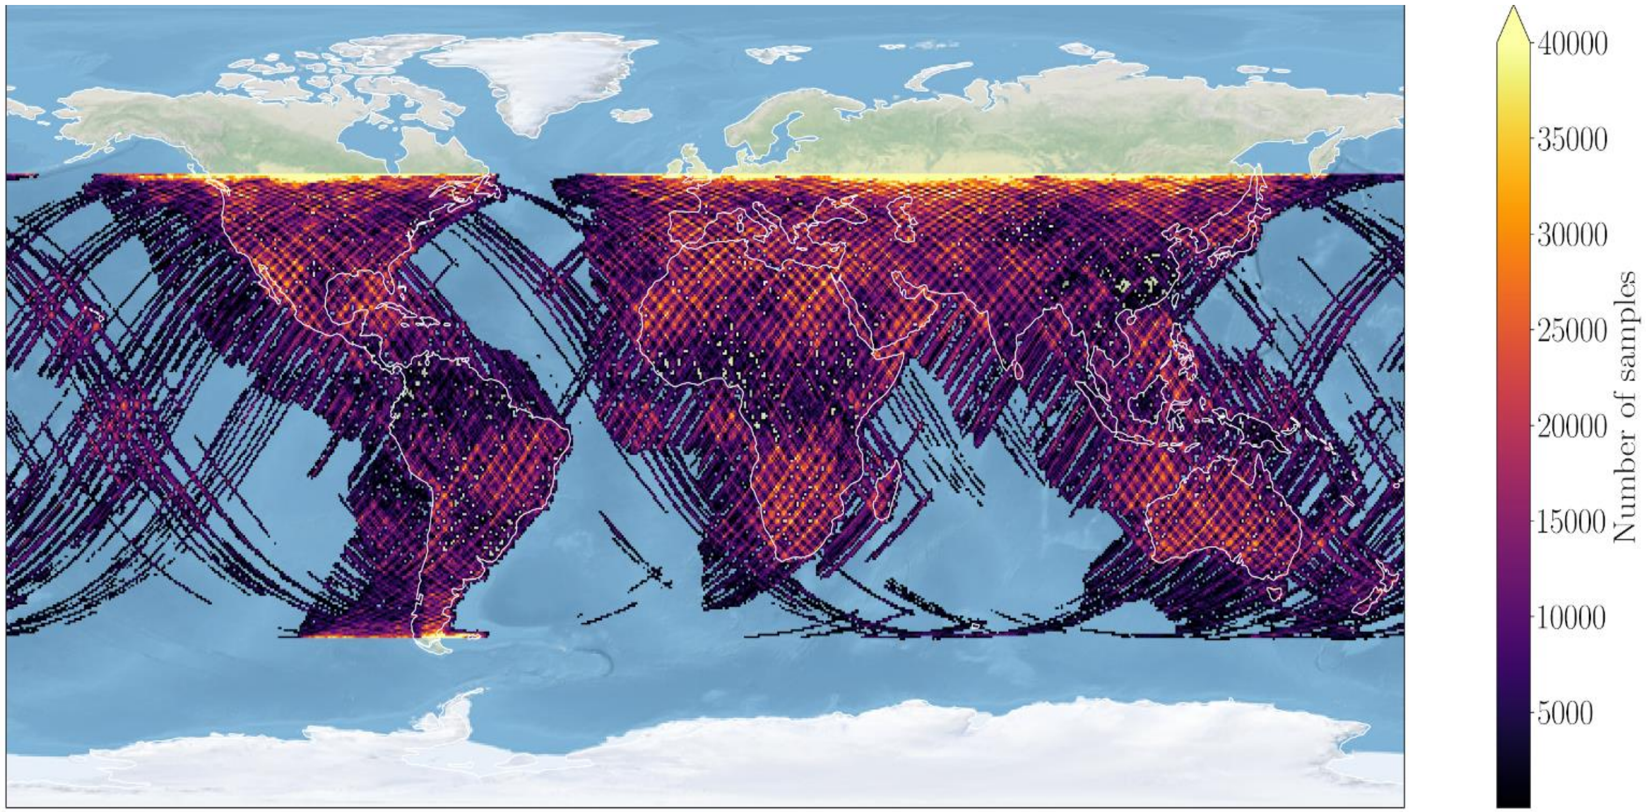

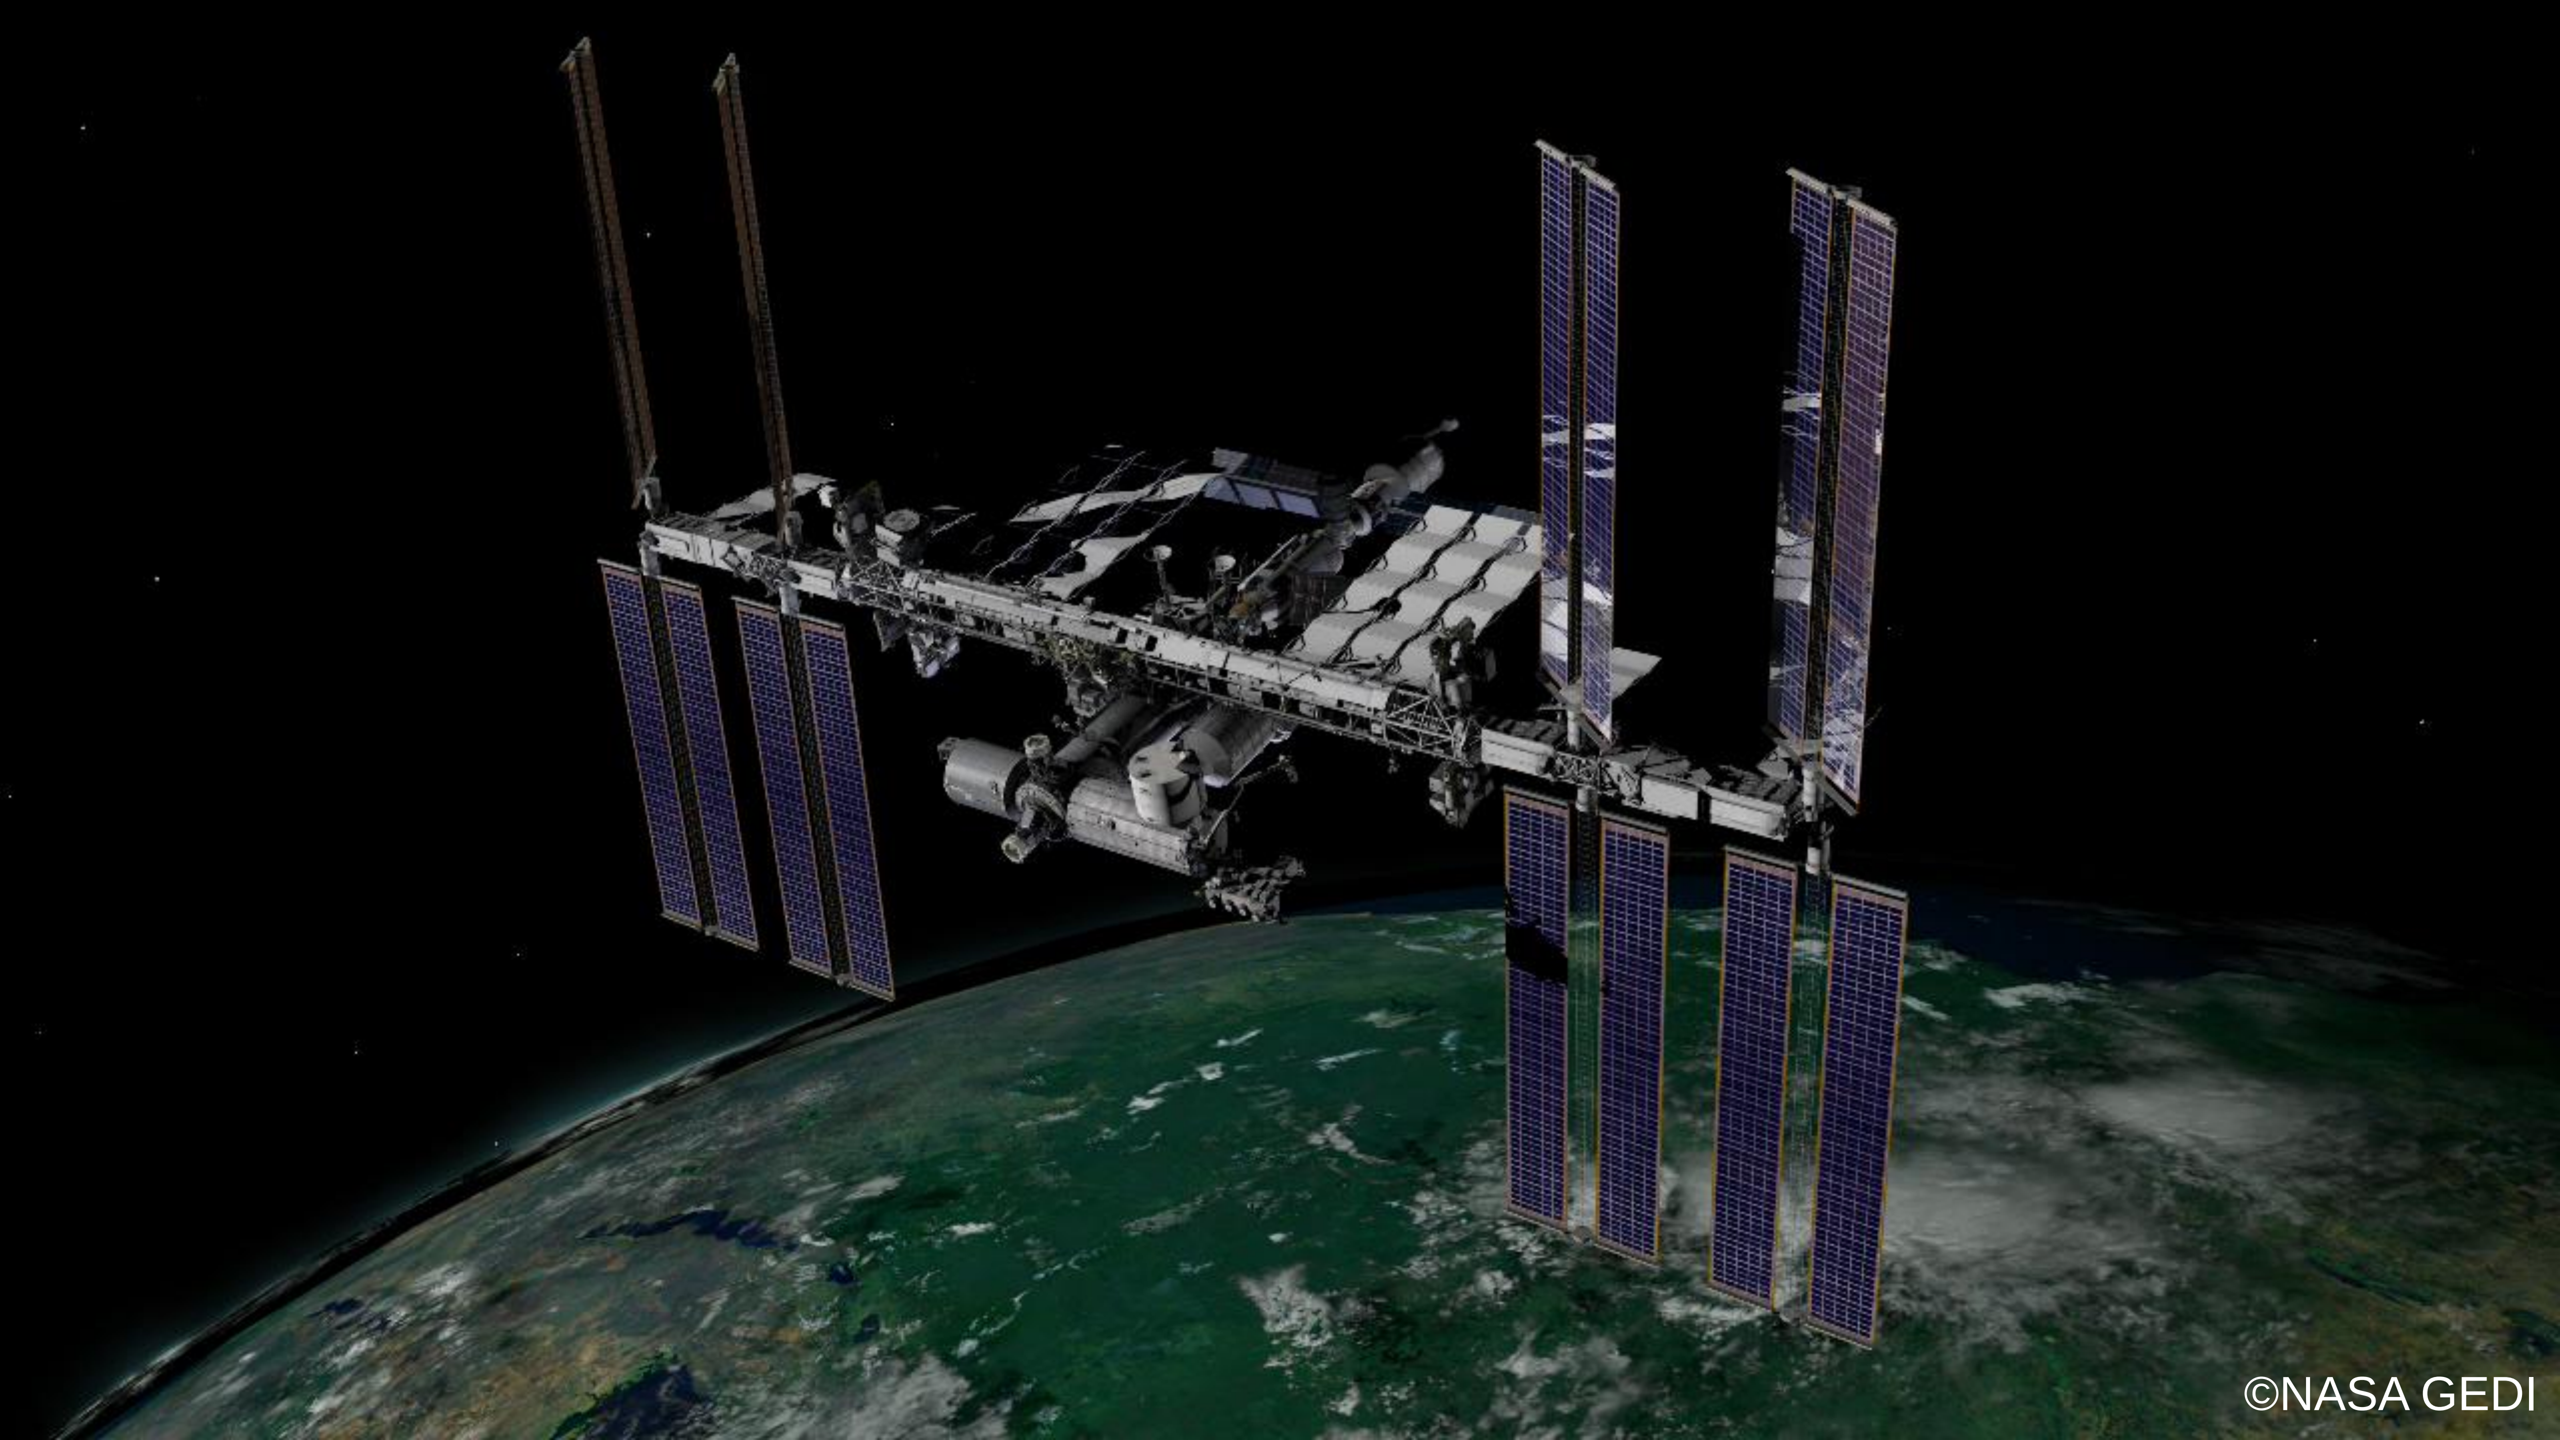
©NASA GEDI

Kalibrierung der GEDI-Daten

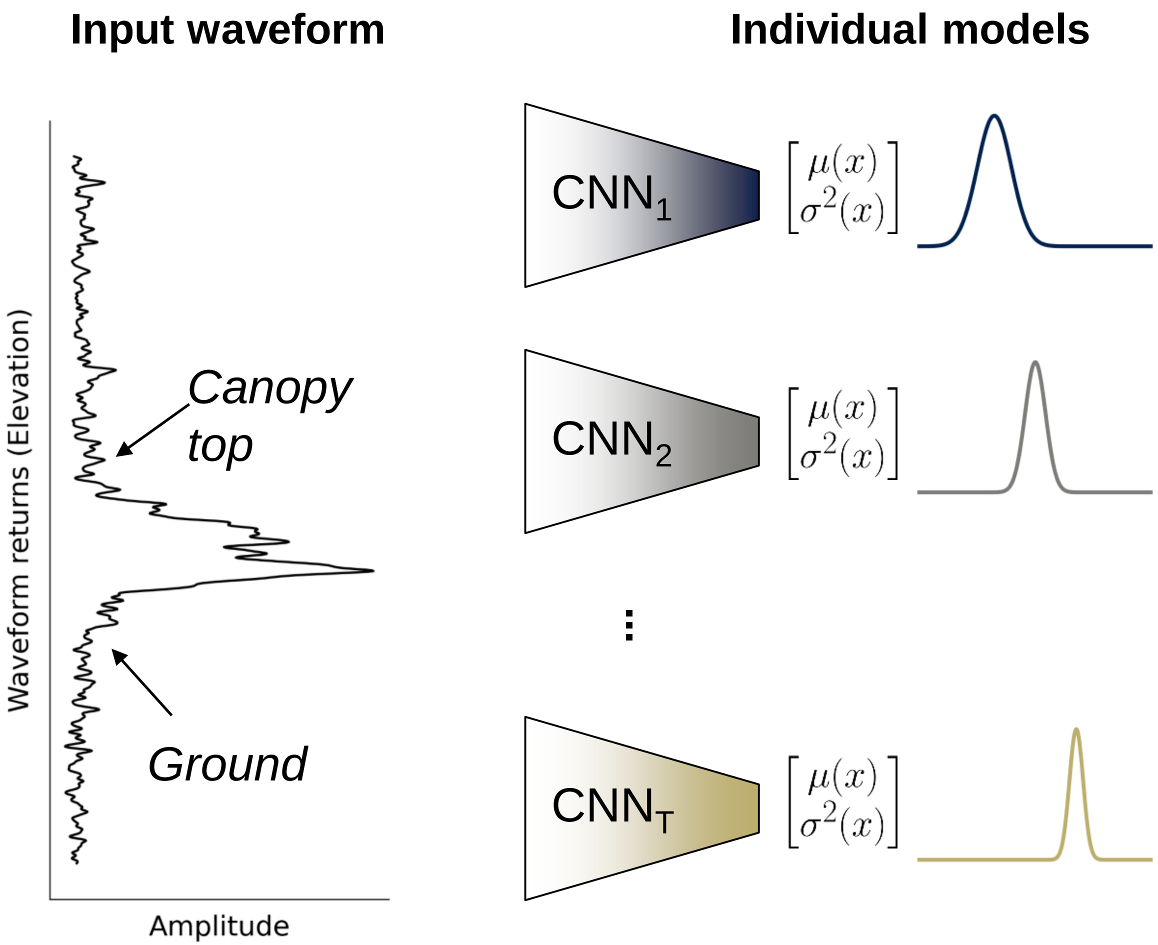

- Each CNN model is trained separately starting from random initializations
- Two outputs per model to approximate the conditional distribution p(y|x)
- Minimize the Gaussian negative log likelihood
- Optimize CNN parameters with stochastic gradient descent (SGD)

Training loss function

$$\mathcal{L}_{NLL} = \frac{1}{N}\sum_{i=1}^{N} \frac{1}{2\sigma(x_i)^2}\left(\mu(x_i) - y_i\right)^2 + \frac{1}{2}\log \sigma(x_i)^2$$

Kalibrierung der GEDI-Daten

Individual models | **Ensemble averaging**

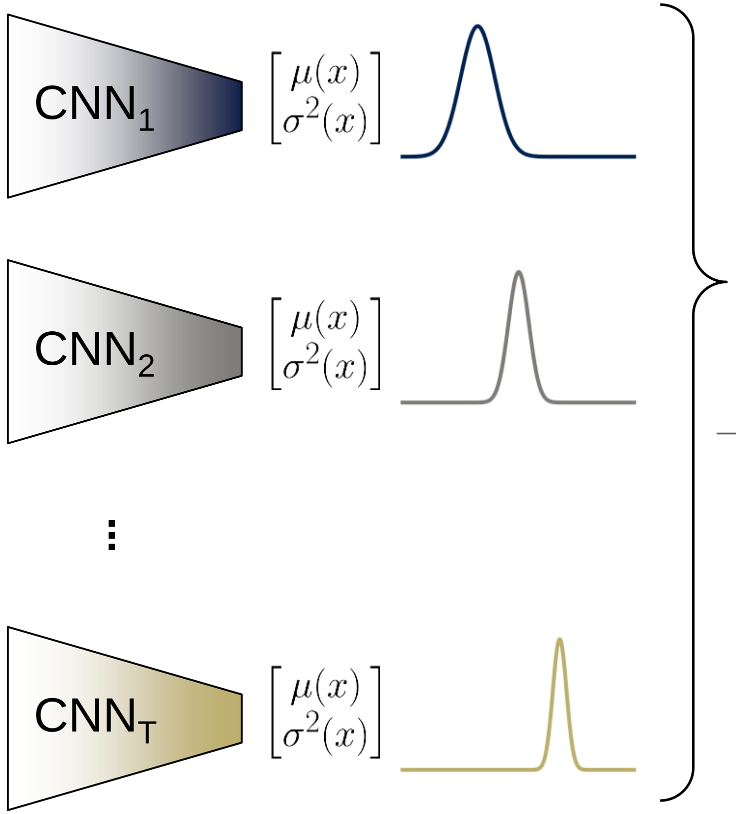

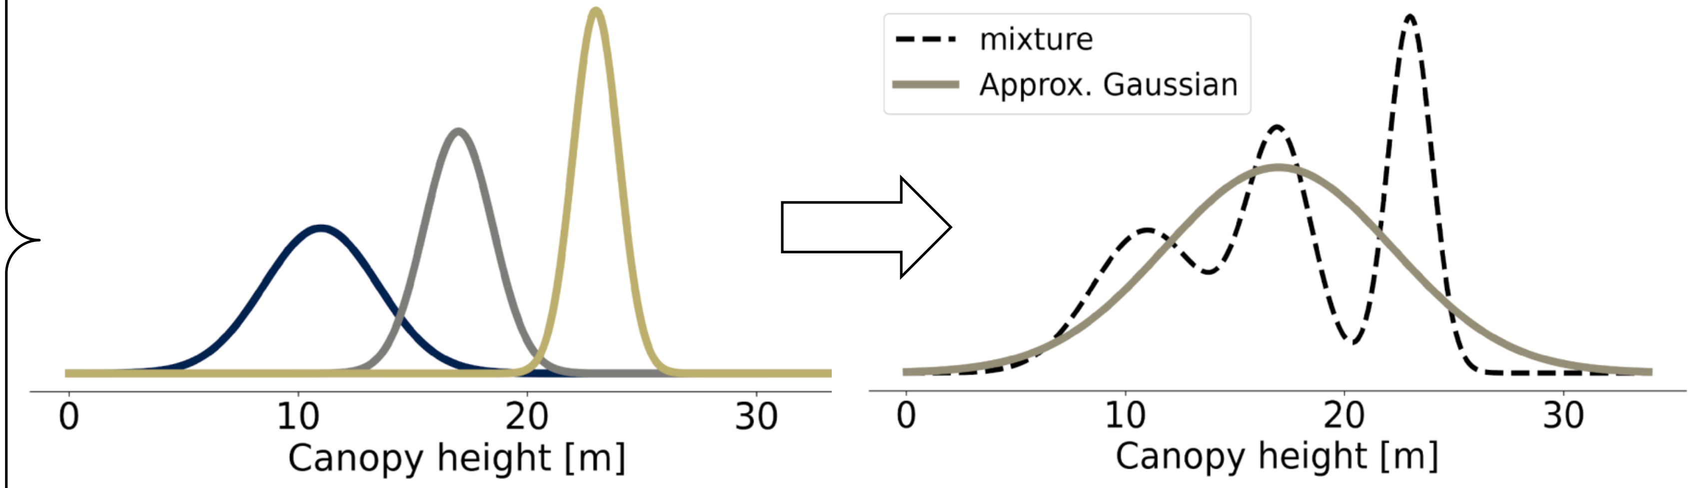

$$\hat{y} = \frac{1}{T}\sum_{t=1}^{T}\hat{\mu}_t \qquad Var(\hat{y}) = \underbrace{\frac{1}{T}\sum_{t=1}^{T}\hat{\mu}_t{}^2 - \left(\frac{1}{T}\sum_{t=1}^{T}\hat{\mu}_t\right)^2}_{\text{epistemic}} + \underbrace{\frac{1}{T}\sum_{t=1}^{T}\hat{\sigma}_t{}^2}_{\text{aleatoric}}$$

epistemic
(model uncertainty)
"knowledge"

aleatoric
(data uncertainty)
"rolling a dice"

Preliminary results: global canopy height with 2.7 m RMSE

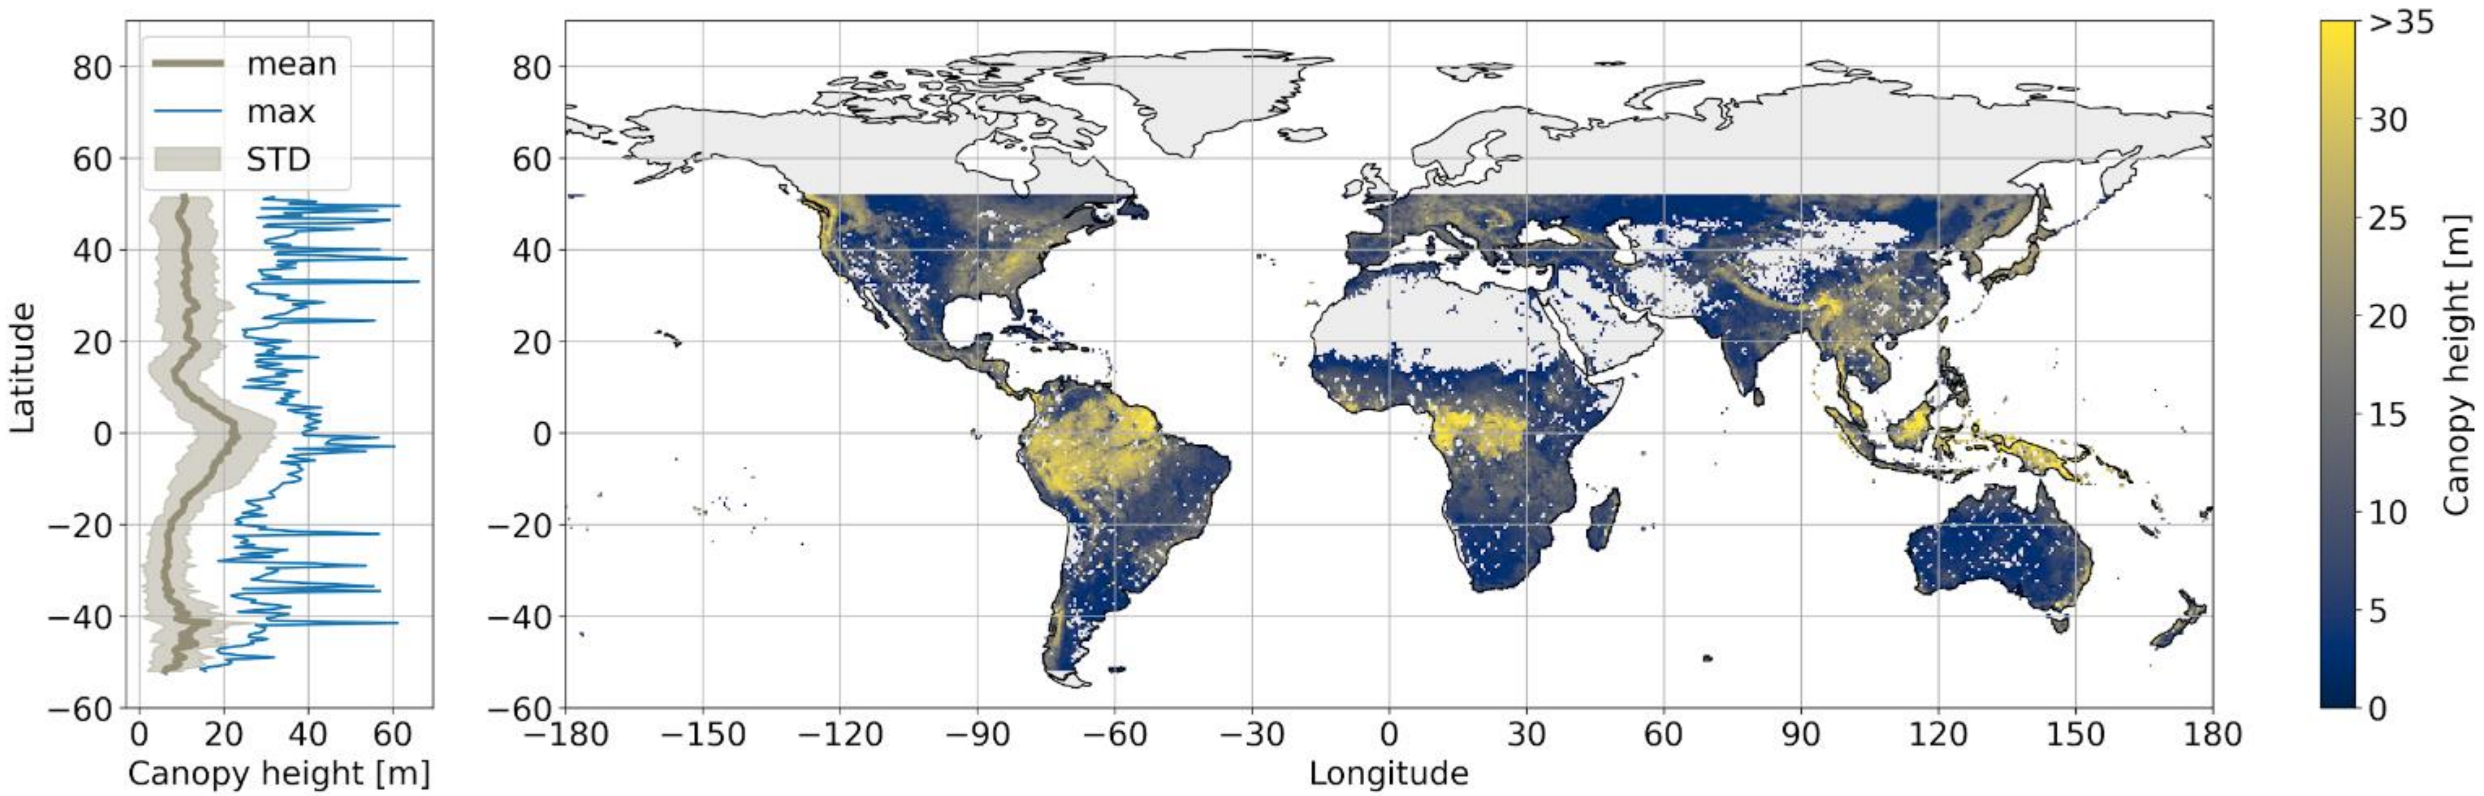

waveforms in non-vegetated areas are filtered out based on MODIS Vegetation Continuous Fields (MOD44B), waveforms filtered based on predictive uncertainty according to the 70% recall setting (i.e., 30% with highest uncertainty filtered out) and values below 0m height suppressed

Preliminary results: global canopy height uncertainties

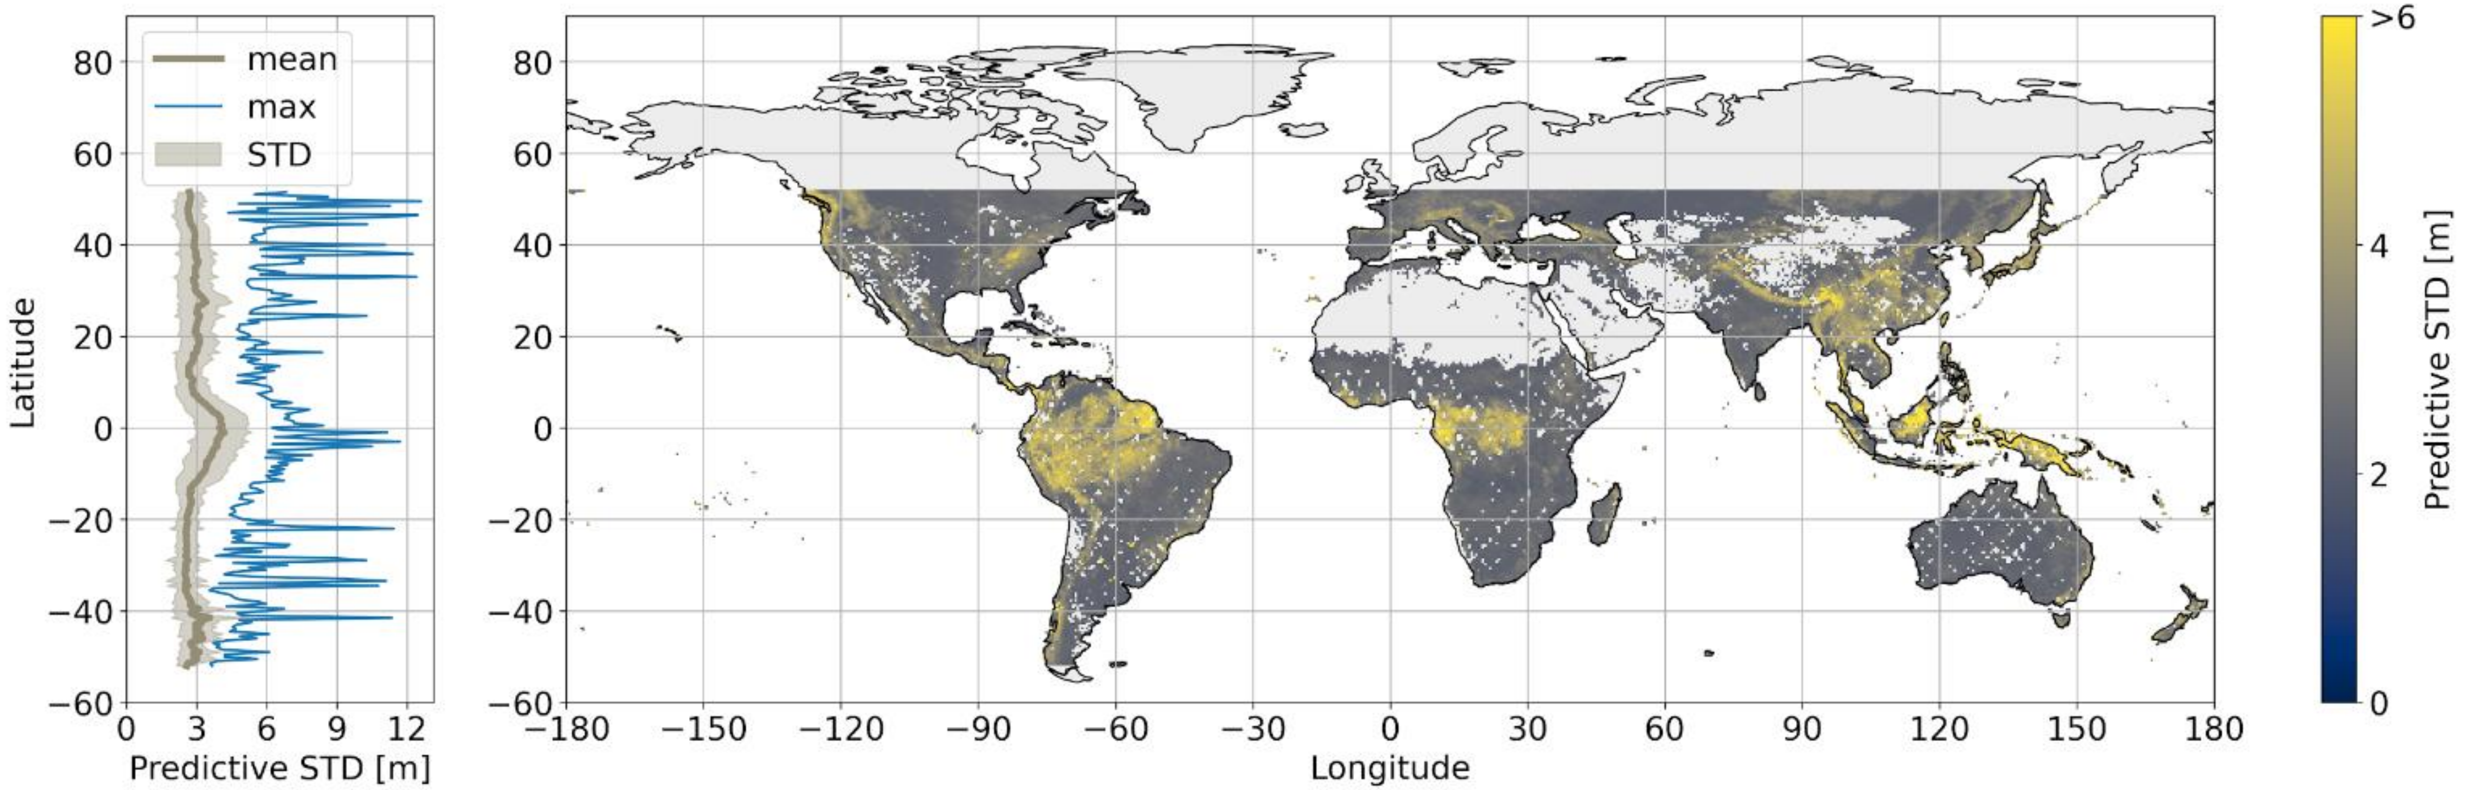

waveforms in non-vegetated areas are filtered out based on MODIS Vegetation Continuous Fields (MOD44B), waveforms filtered based on predictive uncertainty according to the 70% recall setting (i.e., 30% with highest uncertainty filtered out) and values below 0m height suppressed

Combination of sparse LiDAR footprints and dense Sentinel-2 predictions

GEDI

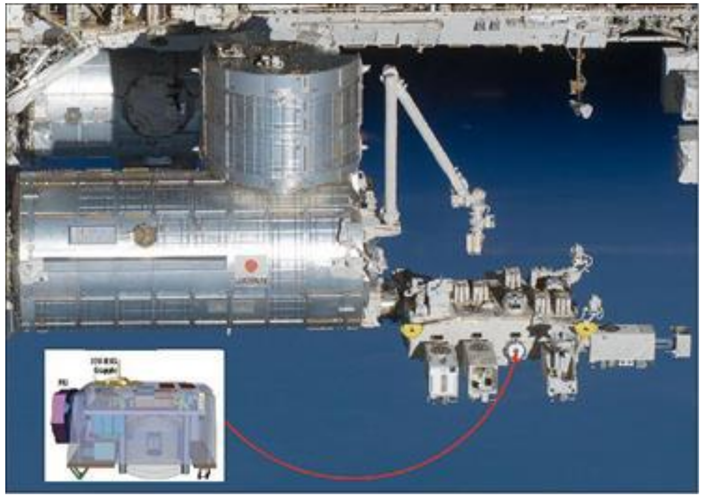

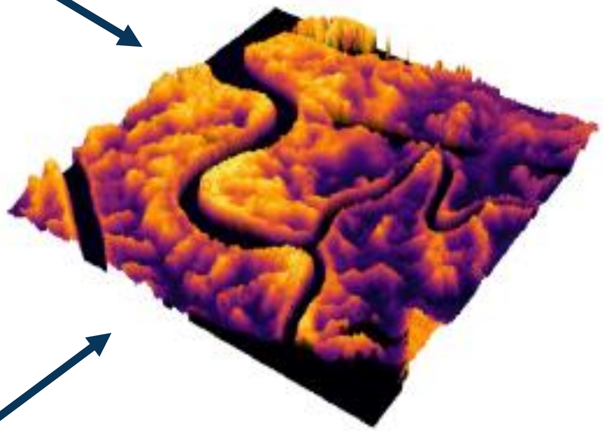

Sentinel-2

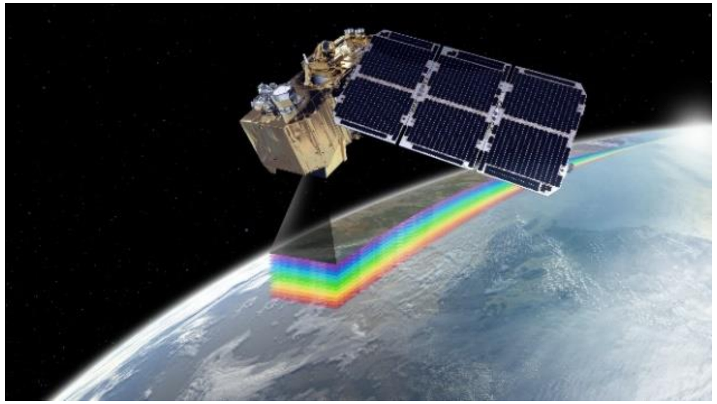

Exciting future directions

- ✓ *combine machine learning and physics-based modelling*
- ✓ *geometric deep learning for non-grid structured data*

Thanks to a great team @ EcoVision Lab !