Machine Learning-Based Reconstruction and Projection of Global Land Use at High Resolution

Marina Castaño[1], Amirpasha Mozaffari[1], Stefano Materia[1], Amanda Duarte[1]

[1] Barcelona Supercomputing Center (BSC), Barcelona, Spain

Background

Land use (LU), land cover (LC), and leaf area index (LAI) are crucial variables for modeling the terrestrial carbon cycle, which in turn influences climate projections. Studies have shown that having LC change data at a resolution higher than 1km can introduce significant biases in the estimation of terrestrial carbon sequestration, its interannual variability, and spatial patterns. Although Earth system observations give us global coverage for land surface boundary conditions at a high resolution, we are lacking historical data and future scenarios.

We propose a machine learning framework to reconstruct LU data in the pre-observation period and create projections for the future. The final product is a continuous LU dataset from 1850 to 2100 with 1km resolution that accounts for different future prospects. The final goal of this project is to expand the framework to extend the temporal coverage of improved-resolution LC and LAI datasets, and then incorporate it into weather-climate frameworks to emulate land surface parameters for Digital Twins.

Data

Our framework is based on two LU products:

- **LUH2h:** LU data from 1850 to 2100 at ~30 km resolution.
- **HILDA+:** LU data from 1900 to 2020 at 1 km resolution.

We use coarse resolution LU (LUH2h) as well as other dynamic predictors (population, climatic zones), and static predictors (geophysical, soil characteristics) to predict fine resolution LU (HILDA+). This framework allows us to train a model that will consistently generate LU datasets from as far back as 1850 (the start of industrialization) to future projections up to the year 2100.

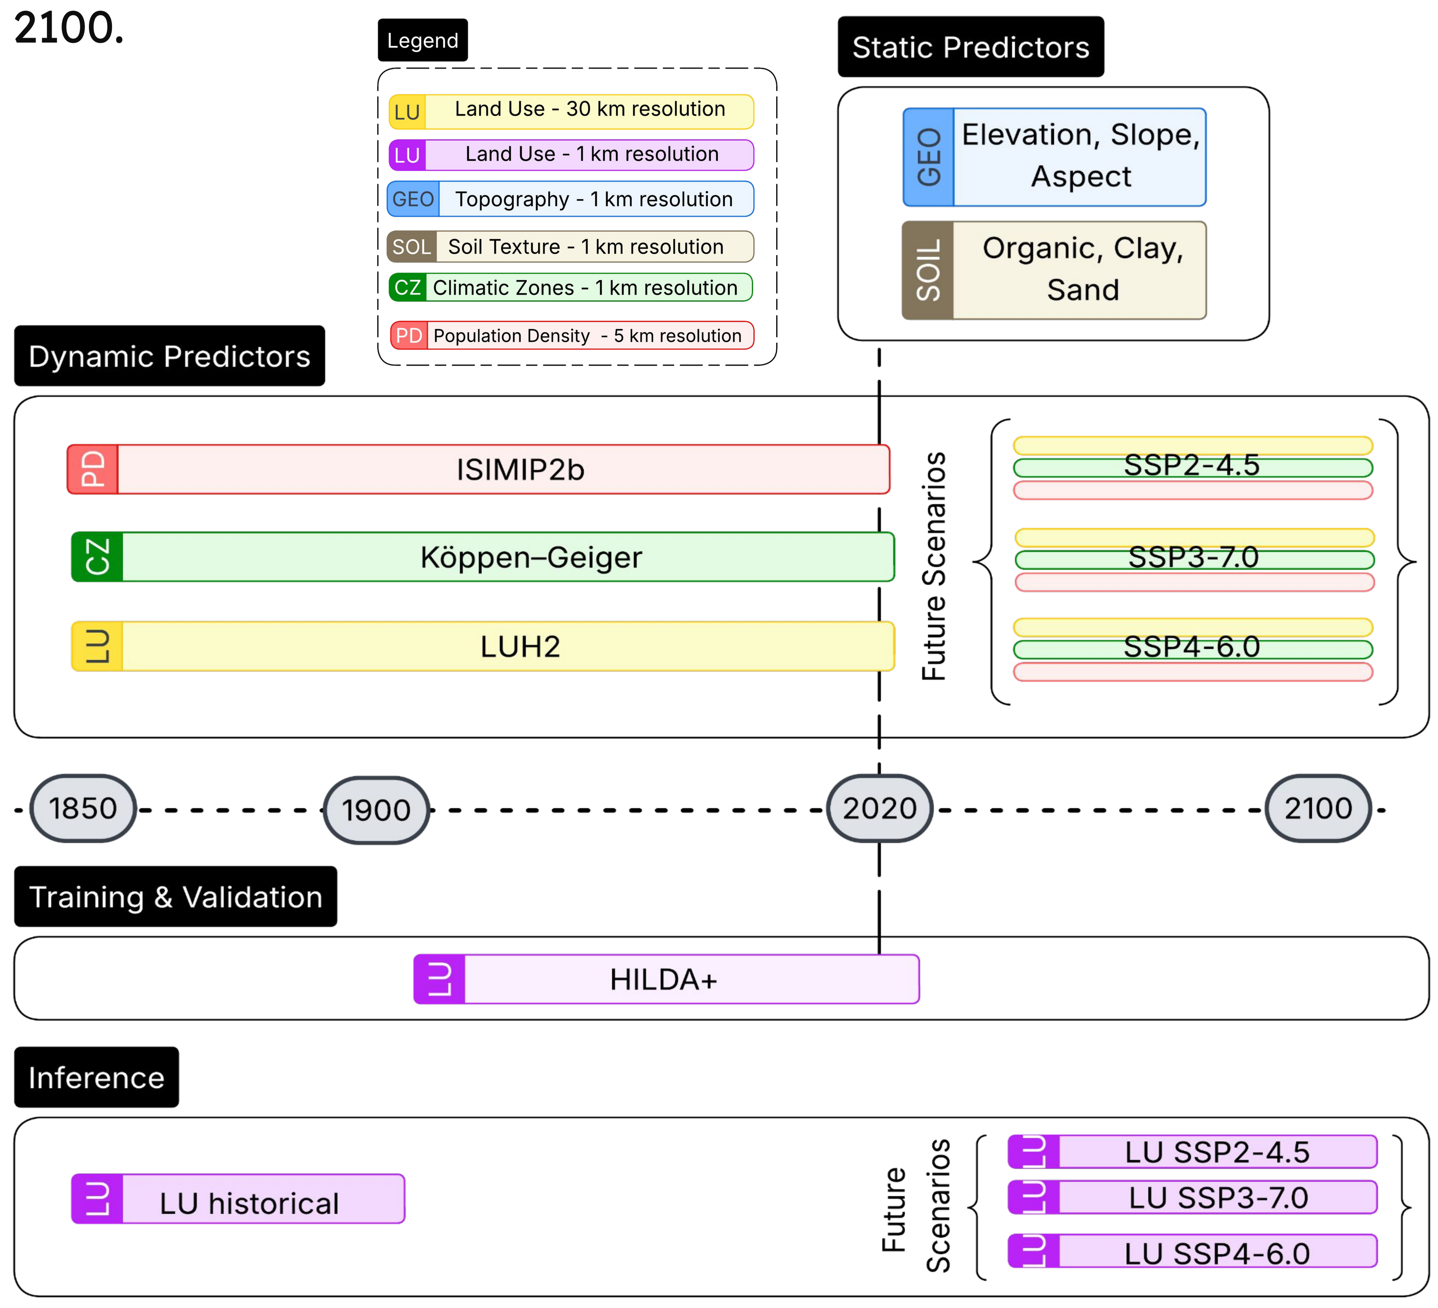

Figure 1: Temporal coverage of the framework's data. This includes the dynamic predictors (population ISIMIP2b, climatic zones Köppen-Geiger, and 30 km land use LUH2h), static predictors (elevation GEBCO, topography and soil PNNL), high-resolution training target (1 km land use HILDA+), and final inferred 1 km LU output for 1850-2100.

Model Architecture

We use a U-Net model that receives images of our static and dynamic channels. Additionally, we pass a prior time step of the high resolution LU target with noise injection of 10%. During training, the model is made to predict for 2 years recursively, using its own predictions as consecutive prior. The loss function is a mixture between the overall performance of the model, and its performance on pixels that have changed from the prior to the target. This ensures that the model doesn't limit itself to copying the prior, due to the high similarity between LU from one year to another. We implement this framework both forward (forecasting) and backward (hindcasting) to be able to expand our data in both directions.

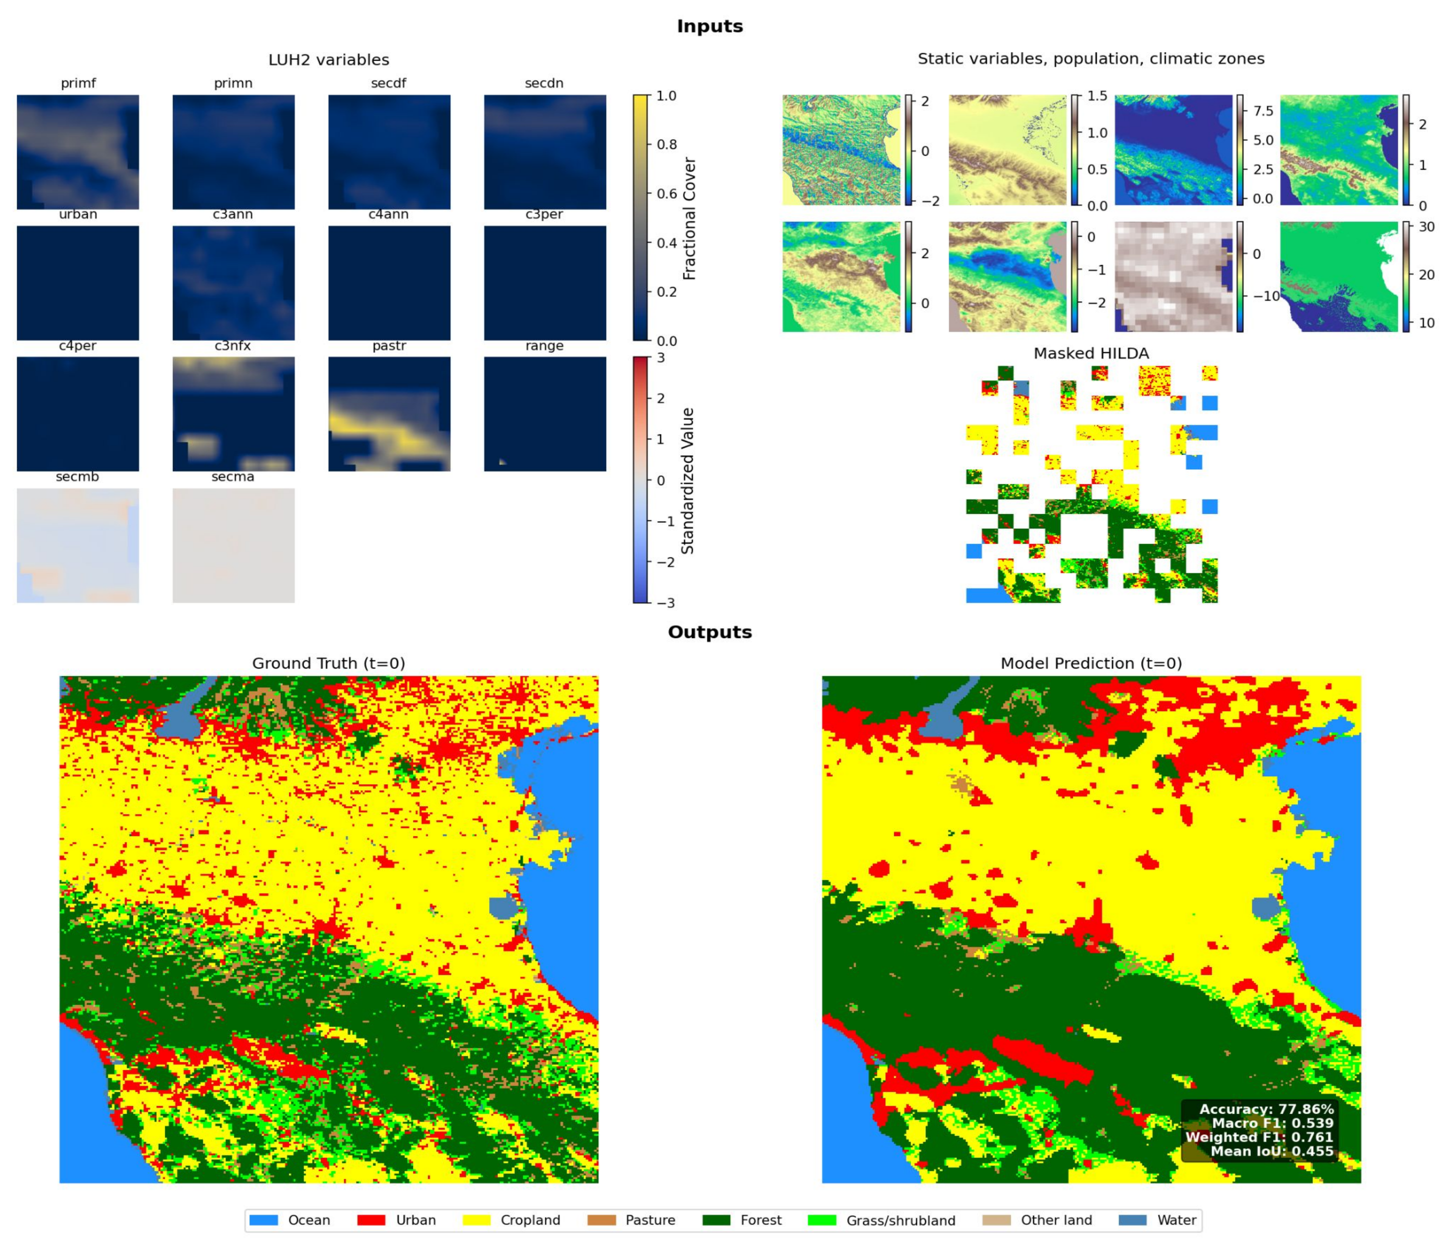

Figure 2: Overview of the LU reconstruction pipeline for Bologna, year 2014. The U-Net model combines coarse-resolution LU, population, climatic zones, elevation, soil, and a masked high-resolution prior LU to predict the final high-resolution LU output.

Results

Trained on 80k samples from 1960 to 2000 for 8 epochs, our U-Net achieved 0.8826 accuracy and 0.6983 IoU overall. We aim to learn the differences from the prior, where the model obtains 0.4796 accuracy and 0.2089 IoU. Table 1 provides further detail into the metrics. Since our goal is to provide inference for several years consecutively, we show the results of predicting for a single year, using a ground truth prior, versus predicting for 15 consecutive years where the priors are taken from previous predictions.

		Accuracy	Macro F1	Weighted F1	Mean IoU
First Year	Overall	0.8826	0.8072	0.8802	0.6983
	LU Change	0.4796	0.3182	0.4818	0.2089
15-Year Average	Overall	0.8270	0.7110	0.8156	0.5894
	LU Change	0.4513	0.2661	0.4175	0.1748

Table 1: Evaluation metrics from a UNet model trained on data from 1960 to 2000 and evaluated on 2001 to 2015. We show little performance degradation when predicting for a single year with a ground truth prior ("First Year"), compared to the average of predicting for 15 years with its own predictions as priors ("15-Year Average"). We also show the performance of the model on the overall image ("Overall"), as opposed to the metrics it obtains on the pixels that differ from the prior ("LU Change"). We continue experimenting to improve the performance of the model on LU change.

Inference

The final product is a worldwide map of LU for each year, so we predict with a sliding window approach, and average the probability of each LU class on overlapping areas. This ensures smooth transitions between image boundaries without blocky artifacts.

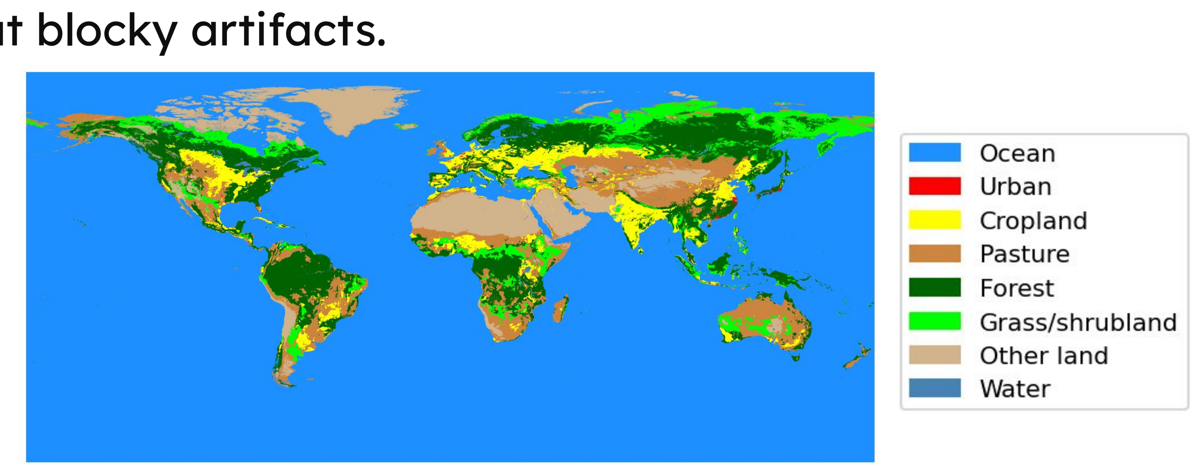

Figure 3: Worldwide inference map for the year 2014.

Future Work

- Extend the framework to predict annual LC and high-frequency (weekly or monthly) LAI.
- Couple the validated LU, LC, and LAI models into weather-climate frameworks to serve as real-time land surface parameters emulators for Digital Twins.

This work is funded by Grant JDC2023-051208-I (MICIU/AEI/10.13039/501100011033 and ERDF "A way of making Europe"), the European Union's Horizon Europe Framework Programme through CONCERTO (Grant Agreement 101185000), TerraDT - Digital Twin of Earth System for Cryosphere, Land Surface and related interactions (Grant Agreement 101187992). AD and SM acknowledge their AI4S fellowship within the "Generación D" initiative by Red.es, Ministerio para la Transformación Digital y de la Función Pública, for talent attraction (C005/24-ED CV1), funded by NextGenerationEU through PRTR".

Funded by the European Union