CON CERTO

Deliverable D2.1 – Creation of a high-resolution global map of land use change covering the historical and the scenario period (1850-2100)

Lead Beneficiary: BSC

Authors: Amirpasha Mozaffari, Stefano Materia, Marina Castaño, Etienne Tourigny, Amanda Duarte

23/12/2025

Funded by
the European Union

D2.1 [Creation of a high-resolution global map of land use change covering the historical and the scenario period (1850-2100)]
#101185000

Prepared under contract from the European Commission

Grant agreement No. 101185000

Project acronym	**CONCERTO**
Project full title	**Improved CarbOn cycle represeNtation through multi-sCale models and Earth obseRvation for Terrestrial ecOsystems**
Start of the project	January 2025
Duration	48 months
Project coordinator	Manuela Balzarolo
Deliverable title	Report describing the global land use change dataset at 1km resolution
Deliverable number	D2.1
Nature of the deliverable	[Report]
Dissemination level	[Public]
WP responsible	WP2
Lead beneficiary	BSC
Author/reviewers	Authors: Amirpasha Mozaffari, Stefano Materia, Marina Castaño, Etienne Tourigny, Amanda Duarte Reviewers: Simone Mantovani (MEEO); Sabrina Outmani (SISTEMA)
Citation	Mozafarri, A. et al. (2025) CONCERTO ***D2.1. Creation of a high-resolution global map of land use change covering the historical and the scenario period (1850-2100).***
Due date of deliverable	31-22-2025
Actual submission date	23-12-2025
Deliverable status	Submitted

Version	Status	Date	Author(s)/Reviewer
0.1	Draft	11 Dec 2025	Amirpasha Mozaffari - BSC
0.2	Review	15 Dec. 2025	Simone Mantovani - MEEO
0.3	Review	15 Dec. 2025	Sabrina Outmani - SISTEMA
0.4	Draft	17 Dec. 2025	Amirpasha Mozaffari - BSC
0.5	Review	22 Dec. 25	Manuela Balzarolo, Sara Totaro - CMCC
Final	Submitted	23 Dec. 25	Sara Totaro – CMCC

Contents

Summary

This deliverable presents the results of the research activities carried out under CONCERTO Task 2.1, which aims to develop a high-resolution global map of land-use change, including historical scenarios.

Uncertainty in the terrestrial carbon cycle remains a major challenge for climate projections and is often exacerbated by coarse or outdated land boundary data that misrepresent land–atmosphere exchanges. Existing datasets frequently suffer from a scaling mismatch, as they lack the spatial resolution, temporal coverage, or logical consistency required for accurate modelling.

This deliverable addresses these gaps by presenting a **deep-learning framework** capable of generating **high-resolution (1 km) land-surface boundary datasets**. The framework aims to provide a **temporally continuous and spatially detailed record of land use** (LU) spanning **from 1850 to 2100**, covering both historical reconstructions and future climate scenarios.

This activity aligns with CONCERTO's objectives by supporting advanced land-management strategies and contributing to carbon-neutrality efforts through improved land-cover representation, the integration of new data sources, and enhanced modelling of carbon fluxes.

The framework employs a U-Net architecture, a deep learning model designed for semantic segmentation, to downscale coarse climate scenarios using high-resolution static features. The model integrates dynamic coarse-resolution forcing data from the Land-Use Harmonization 2 (LUH2) dataset with high-resolution static predictors, including topography, soil composition, slope, and aspect. It is trained against the HILDA+ dataset, which serves as the high-resolution ground truth consolidated into eight distinct land use classes. To ensure data consistency and robust evaluation, the study employed a standardized preprocessing pipeline using Climate Data Operators (CDO) to transform inputs into Analysis-Ready Cloud Optimized (ARCO) Zarr format, alongside Grid-Based Partitioning (GBP) and Farthest Point Sampling (FPS) to eliminate spatial autocorrelation leakage between training and testing sets.

The U-Net model demonstrated high fidelity in reconstructing complex land-use patterns. During the inference validation for the year 2014, the framework achieved a global accuracy of 94.88% and a mean Intersection over Union (mIoU) of 0.8569. While the model performed exceptionally well on dominant classes such as forests, crops, drylands and other land, performance dropped for underrepresented classes like

urban areas, highlighting a sensitivity to class imbalance. Qualitative analysis of consecutive years confirmed the model's ability to maintain temporal consistency and smooth transitions across patch boundaries without edge artifacts.

The execution of the inference pipeline produced the high-resolution global map of land use change, a comprehensive dataset stored in both Zarr and NetCDF formats. The output includes a historical dataset spanning 1850–2020, which integrates model predictions for the pre-observation era with consolidated HILDA+ observational records, as well as three distinct future projection cubes corresponding to the SSP2-4.5, SSP3-7.0, and SSP4-6.0 climate scenarios. This work successfully provides the high-resolution land use dataset necessary for coupled carbon-water-energy flux simulations.

List of Abbreviations

ANNs	Artificial neural networks
ARCO	Analysis-ready cloud optimised
BSC	Barcelona Supercomputing Center
CA	Cellular automata
CDO	Climate data operators
CNNs	Convolutional neural networks
ConvRNN	Convolutional recurrent neural network
FAO	Food and Agriculture Organization
FLUS	Future land use simulation
FPS	Farthest point sampling
GEBCO	General bathymetric chart of the oceans
HILDA+	Historic land dynamics assessment +
IoU	Intersection over union
LC	Land cover
LSTM	Long short-term Mmemory
LU	Land use
LUH2	Land-use Harmonization 2
LUH2f	Land-Use Harmonization 2 Future
LUH2h	Land-Use Harmonization 2 Historical
mIoU / IoU	(Mean) Intersection over Union
ML	Machine learning
PNNL	Pacific Northwest National Laboratory
RF	Random Forests
RNN	Recurrent Neural Network
WP	Work Package
XGBoost	Extreme Gradient Boosting

1. Introduction

The terrestrial carbon cycle remains a major source of uncertainty in climate projections (Booth et al., 2012), partly because land surface processes are not fully resolved. Land-use and land-cover (LU and LC) changes are significant sources of uncertainty in estimating carbon fluxes (Gier et al., 2024; Brasika et al., 2024). However, relying on coarse or outdated land boundary data can lead to a critical misrepresentation of land–atmosphere exchanges. Studies have shown that coarse land cover spatial resolutions can introduce substantial biases in simulated terrestrial carbon sequestration, affecting its magnitude, interannual variability, and spatial distribution (Zhao et al., 2010). Given that land-use change has consistently driven substantial carbon emissions over the last century, as shown in Figure 1 (Friedlingstein et al., 2025), the accurate representation of high-resolution land boundary conditions is essential for reducing uncertainties in coupled carbon–water–energy fluxes (Gier et al., 2024).

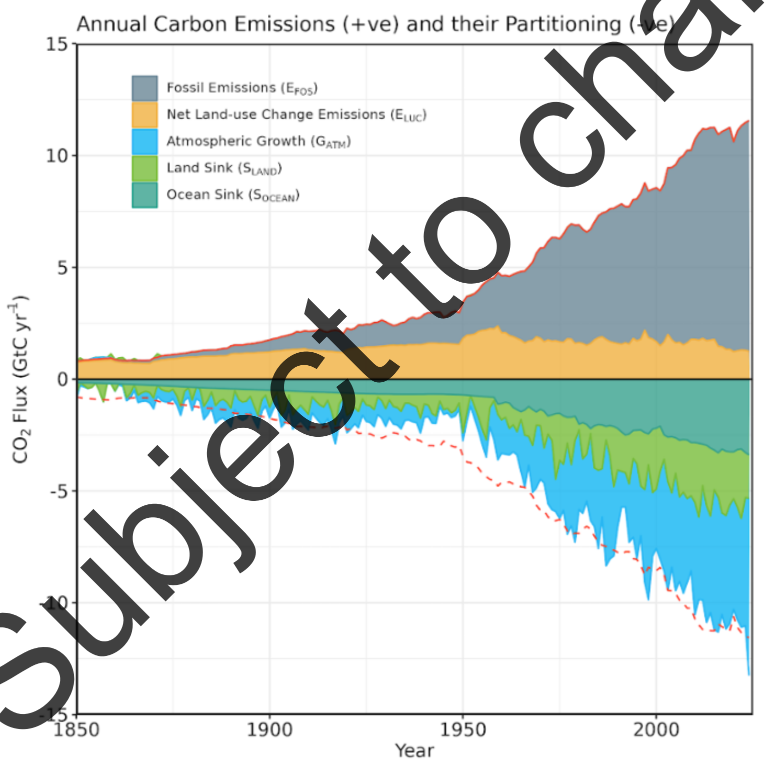

Figure 1. Annual global carbon emissions and their partitioning from 1850 to present. The data highlights the significant and persistent contribution of net land-use change emissions (E_{LUC}, orange) to the atmospheric carbon burden, alongside fossil emissions. Source: Global Carbon Budget 2025 (Friedlingstein et al., 2025).

While several datasets offer valuable insights into land use, each fall short in terms of spatial resolution, temporal coverage (historical or future), or logical continuity (Hurtt et al., 2020; Winkler et al., 2020; Lv et al., 2025; Chen et al., 2020). To address this scaling mismatch, our work proposes a deep learning framework to develop high-resolution land surface boundary datasets that span historical, contemporary, and future periods. By integrating satellite-era observations with reconstructions for pre-observation times and projections for observation-limited futures, we provide a temporally continuous, spatially detailed land use dataset for use in climate and reanalysis models. Specifically, our approach leverages static geophysical features to downscale coarse scenarios, improving the representation of land surface heterogeneity and enabling more accurate simulations of carbon, water, and energy fluxes over time.

2. Related approach

The reconstruction and projection of land use and land cover have historically relied on rule-based harmonisation (Gao et al., 2023), such as the standard LUH2 framework (Hurtt et al., 2020). In addition, methods like land-use translators are created to downscale historical and future changes to the regional scale (Hoffmann et al., 2022). However, a broad spectrum of machine learning (ML) techniques is increasingly employed to overcome the 'scaling mismatch' of these coarse baselines and to generate high-resolution, spatially continuous datasets.

Pixel-wise and tree-based reconstruction, particularly Random Forests (RF) and Extreme Gradient Boosting (XGBoost), have become the standard for high-resolution historical reconstruction due to their ability to handle non-linear relationships in heterogeneous data, and have been widely adopted for diverse tasks, including cropland reconstruction (Yang et al., 2020) and LU/LC change estimation (Aboh et al., 2020). Similarly, the HILDA+ dataset employs data-driven fusion, utilising RF to downscale coarse FAO statistics by learning relationships with fine-scale drivers like population density (Winkler et al., 2020). XGBoost, another ensemble-based method, has shown promising results in the historical reconstruction of LC and LAI (Mozaffari et al., 2025b).

Dynamic simulation and hybrid systems capture the temporal evolution and competition between land-use classes. Researchers have developed hybrid systems combining ML with cellular automata (CA) (Wolfram, 1983). The future land use simulation (FLUS) model integrates artificial neural networks (ANNs) to estimate suitability probabilities, coupled with CA to simulate complex local competition and inertia (Xu et al., 2019) (Chen et al., 2021). Building on systems theory, the Modified CLUMondo model employs RF to dynamically calculate 'competitive advantage', allowing for the simulation of complex transitions in tipping-element regions under future climate pledges (Lv et al., 2025). Additionally, recurrent architectures like long short-term memory (LSTM) networks have been applied to forecast land use change by explicitly modelling temporal dependencies in historical time series (Zhao et al., 2023; Yang et al., 2021). Convolutional neural networks (CNNs) and all their variations have emerged as a powerful tool for spatial downscaling and super-resolution, addressing the limitations of pixel-wise and rule-based methods. Unlike RF or CA, CNNs utilise spatial context to resolve texture and hierarchy, making them ideal for bridging the gap between coarse climatic forcing and kilometre-scale boundaries (Truong et al., 2024) (Zhu et al., 2015) (Li et al., 2024), (Feng et al., 2025).

3. Methods

3.1. Data

Our study integrates multiple global datasets, distinguishing between the high-resolution target variable used for training and the diverse set of predictor variables used as model inputs.

Target Variable: **HILDA+ (high-resolution land use change)**: the ground truth for our model is derived from the HILDA+ dataset (Winkler et al., 2020). This dataset provides global land use/cover information at a spatial resolution of 1 km. While the original HILDA+ classification scheme contains 13 classes, we consolidated these into 8 distinct land use (LU) categories to reduce class imbalance and streamline the learning objectives for this study, as shown in Table 1.

Predictor Variables: to train our model, we defined the input feature space, that is the collection of variables used to make predictions. This consists of climatic, topographical, and anthropogenic variables. To ensure compatibility, all predictor datasets were resampled to the target 1~km grid during the preprocessing stage and are detailed below:

- **land-use harmonisation (LUH2)**: we utilised both historical (LUH2h) and future (LUH2f) forcing data (Hurtt et al., 2020). Data comprise 12 fractional LU variables (ranging from 0 to 1) and two land surface parameters standardised to unit variance on a 0.25° × 0.25° (~31 km) grid;

- **topography and bathymetry**: elevation and bathymetry data were obtained from the GEBCO 2024 Grid (GEBCO, 2024), providing essential vertical datum information to constrain land suitability;
- **geophysical parameters**: we incorporated high-resolution (1 km) geophysical surface parameters derived from PNNL global datasets (Li et al., 2024). Specifically, we utilised:
    - **soil composition:** to capture the heterogeneity of the root zone, we incorporated data on clay, sand, and organic content. These variables are provided at multiple vertical resolutions (6 and 10 layers) to account for depth-dependent substrate variability (Li et al., 2024).
    - **slope and aspect:** to capture terrain characteristics that influence local microclimates and land utility.

Table 1: Top) HILDA+ LU/LC classes code and class names, middle) LUH2 abbreviated and full class name, bottom) Concerto LU dataset

HILDA+ LU/LC Classes	
Code	**Class Name**
11	Urban
22	Cropland
33	Pasture
40	Forest (Unknown/Other)
41	Forest (Evergreen, needle leaf)
42	Forest (Evergreen, broad leaf)
43	Forest (Deciduous, needle leaf)
44	Forest (Deciduous, broad leaf)
45	Forest (Mixed)
55	Grass/shrubland
66	Other land
77	Water

LUH2	
Abbreviation	**Full Class Name**
primf	Primary Forest
primn	Primary Non-Forest
secdf	Secondary Forest
secdn	Secondary Non-Forest
pastr	Managed Pasture
range	Rangeland
urban	Urban Land
c3ann	C3 Annual crops
c4ann	C4 Annual crops
c3per	C3 Perennial crops
c4per	C4 Perennial crops
c3nfx	C3 Nitrogen-fixing crops
secma	Secondary Mean Age
secmb	Secondary Mean Biomass

Concerto LU dataset v0.1	
Code	**Class Name**
11	Urban
22	Cropland
33	Pasture
40	Forest
55	Grass/shrubland
66	Other land
77	Water

3.2. Preprocessing

To ensure consistency across diverse data sources, we established a standardised preprocessing pipeline centred on the Climate Data Operators (CDO) software (Schulzweida, 2019). We remapped these inputs using nearest-neighbour interpolation to align with the grid and projection of the target HILDA+ dataset (Winkler et al., 2020). This process involved spatial interpolation to generate a unified dataset at 1 km resolution in the WGS84 projection.

The PNNL six-layer vertical soil characteristic (Li et al., 2024a) was consolidated by generating 2D depth-weighted averages for each soil variable to serve as surface-representative features. A non-linear weighting scheme (w = [0.40, 0.25, 0.15, 0.10, 0.05, 0.05]) was applied to the vertical column, prioritising the upper soil horizons where land–atmosphere interactions and biological activity are most pronounced.

Finally, to facilitate scalable machine learning workflows, the processed data was transformed into Analysis-Ready Cloud Optimized (ARCO) Zarr format (Newman, 2024). The resulting datasets were stored as 3D cubes, chunked in 512 × 512-pixel blocks to ensure efficient data retrieval during training.

The designed data scheme, as illustrated in Figure 2, represents the dynamic and static predictors along with the target variable over time.

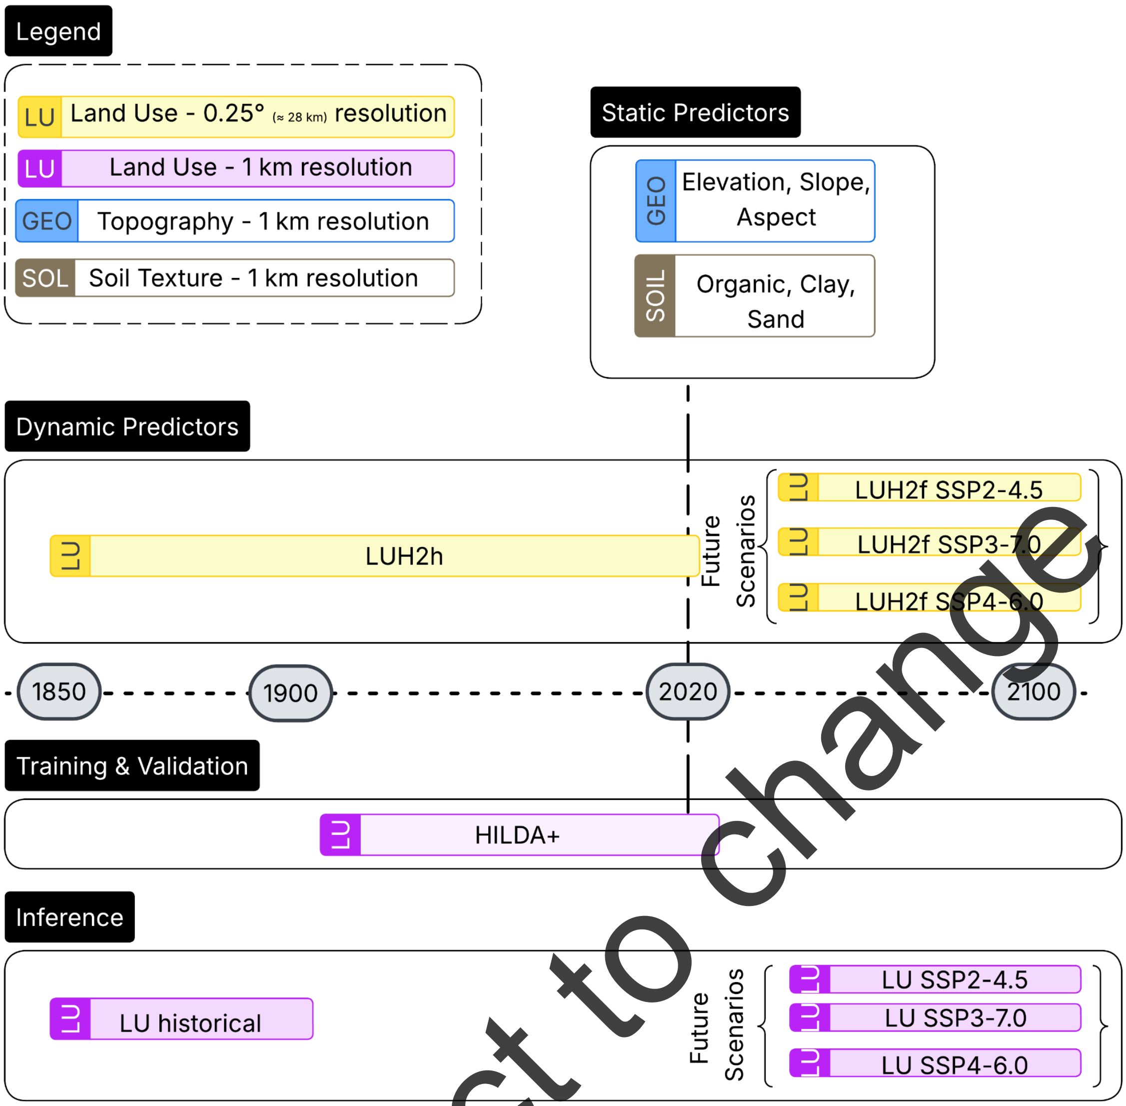

Figure 2. Temporal coverage of the framework's data. This includes the dynamic predictors (coarse 30 km LUH2h), static predictors (1 km topography, soil, and climate), the high-resolution training target (1 km HILDA+), and the final inferred 1 km LU output for 1850-2100.

3.3. Model

We developed the backbone for LU reconstruction and projection using a U-Net architecture (Ronneberger et al., 2015). Our model employed a standard U-Net with 35 base channels (14 LUH2 for 2 timesteps, 6 static, 1 HILDA prior), accepting 512 × 512-pixel inputs. Its encoder consisted of four max-pooling and double-convolution blocks, mirrored by a symmetric decoder with upsampling and skip connections, terminating in a 1 × 1 convolution for final segmentation (Mozaffari et al., 2025a). To encode the categorical target, we utilised entity embeddings (Guo &

Berkhahn, 2016). The dimensionality of these embeddings was determined by one-quarter of the number of unique categories (8 HILDA classes), specifically for the target (2). Figure 3 shows the overview of the LU reconstruction pipeline with the U-Net model.

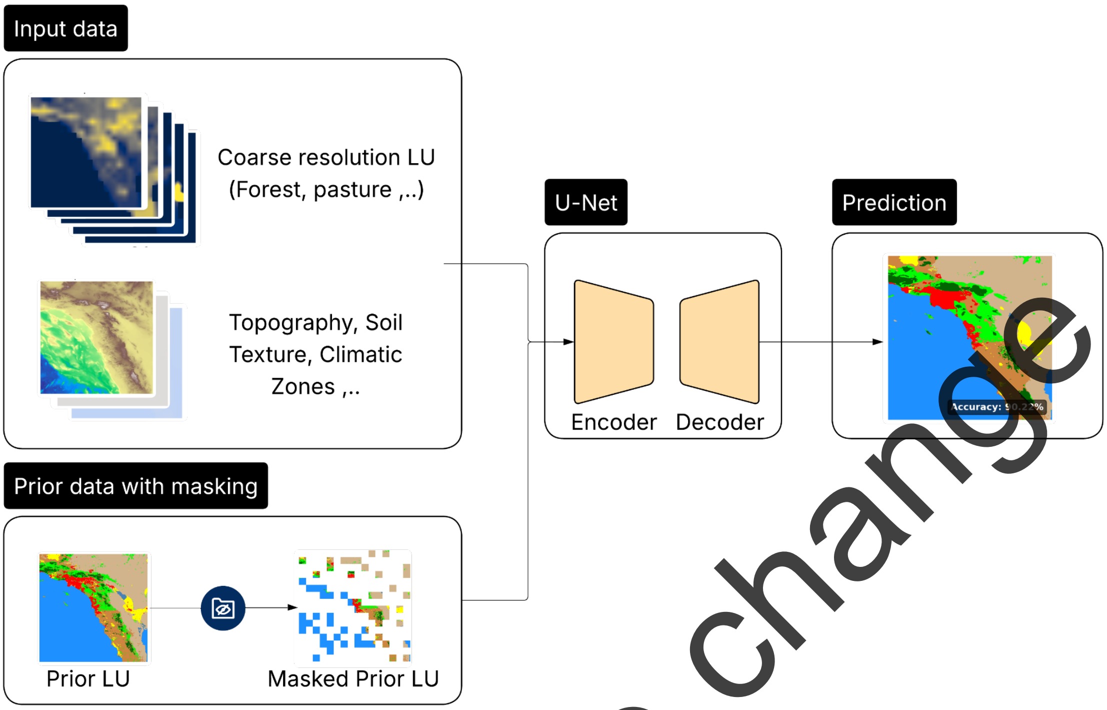

Figure 3. Overview of the LU reconstruction pipeline: A U-Net model takes as input coarse-resolution LUH2 land-use classes (e.g., forest, pasture), land surface parameters, high-resolution climatic zones, and static environmental variables (e.g., elevation, slope, soil). Masked land-use data from an adjacent year, preceding for reconstruction or following for forecasting, is added as auxiliary input. The standard U-Net then segments and predicts high-resolution land-use classes for the target year.

To establish a robust benchmark for our proposed U-Net architecture, we implemented three baseline models for performance comparison. While all models utilised patch-based inputs, we adapted the data representation to suit specific architectural requirements. We formulated the input as sequential data for the recurrent models, whereas for the gradient boosting model, we explicitly flattened the spatial dimensions to create vectorized inputs. The baseline models are defined as follows:

- recurrent neural network (RNN): a standard sequence-based network employed to capture temporal dependencies within the patch sequence (Rumelhart et al., 1986);

- convolutional RNN (ConvRNN): a hybrid architecture that integrates convolutional operations into the recurrent structure, allowing the model to process sequential data while preserving local spatial features (Shi et al., 2015).
- XGBoost: an efficient implementation of gradient tree boosting. Unlike the deep learning models, this required the input patches to be serialized (flattened) into 1D feature vectors for regression/classification (Chen & Guestrin, 2016).

We have built a data loader that provides all the necessary targets, priors, and inputs as separate channels, which is used for both our U-Net and RNN models and provides partial masking (60%) with randomly selected 16 × 16-pixel patches. We have replicated the same loader for XGB using Dask.

3.4. Evaluation

We have used a two-phase process for evaluation, designed to balance spatial independence with uniform coverage. First, we apply *Grid-Based Partitioning*, where the spatial domain is divided into a coarse grid (e.g., 30 × 30 bins). Instead of splitting individual points randomly, the algorithm assigns entire grid cells exclusively to the training, validation, or testing sets. This strategy creates physical separation between the datasets, ensuring that a patch in the test set never has an immediate neighbour in the training set, thereby effectively eliminating spatial autocorrelation leakage. The geographical distribution of the training, validation, and test sets used for model training is illustrated in Figure 4. Second, within these assigned regions, the method executes *Farthest Point Sampling* (FPS). This iterative algorithm selects a fixed number of points that are maximally distant from one another, forcing the dataset to be uniformly distributed across the available coordinates rather than clustered in high-density areas.

From a total pool of approximately 214 million land pixels available for a single year, we subsampled a final dataset of 448,000 points, stratified into 320,000 training, 64,000 validation, and 64,000 testing samples. Table 2 presents the percentage distribution of the HILDA+ classes across all patches. This distribution is also shown for the training, validation, and test sets, indicating their respective proportions of the full dataset.

Table 2: Percentage distribution of HILDA+ classes across the generated 512 × 512 patches for each data split compared to the full masked dataset.

Class	Train (%)	Val (%)	Test (%)	Fullset (%)
Urban	1.16	0.61	0.54	1.1
Cropland	10.19	9.22	12.65	10.29
Pasture	20.43	24.2	19.72	20.53
Forest	29.22	34.36	32.08	30.03
Grass/shrubland	16.1	8.76	11.7	14.67
Other land	20.85	20.06	21.96	21.34
Water	2.05	2.78	1.35	2.14

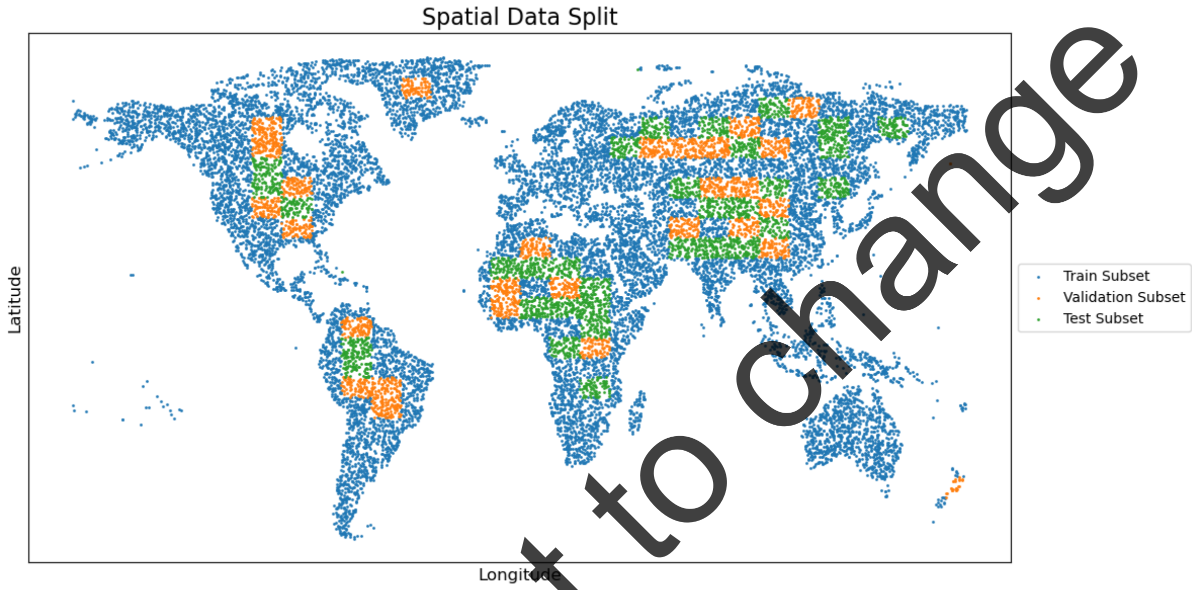

Figure 4. Spatial distribution of the training, validation, and test sets

3.5. Inference

To generate spatially consistent global land-use maps, the system employs a distributed inference pipeline that processes the globe in overlapping 512 × 512 patches using a U-Net model. To ensure spatial consistency and eliminate edge artifacts (blockiness), predictions are aggregated from overlapping patches Ω (x,y) using a Gaussian weighting function W that prioritises central pixels. The final probability map $\bar{P}(x, y)$ is computed by normalising the accumulated weighted probabilities by the overlap count N(x,y):

$$\bar{P}(x,y) = \left(\sum_{k \in \Omega(x,y)} p^{(k)}(x, y) \cdot W(u_k, v_k) \right) / N(x, y) \qquad (1)$$

The discrete land-use class is then derived via $L(x, y) = \mathrm{argmax}(\bar{P}c(x, y))$, effectively smoothing transitions across patch boundaries. These weighted probability maps are accumulated into global tensors across distributed processes, tracking both the sum of probabilities and the number of overlapping predictions per pixel. Once all patches are processed, the system normalises the accumulated sums by the overlap counts to compute an average weighted probability, determines the final class via an argmax operation, and serialises the resulting seamless mosaic into Zarr stores for both forward and backward temporal projections.

4. Results

4.1. Model

The U-Net architecture, a cornerstone of our LU classification methodology, was rigorously trained on a dataset comprising 800,000 carefully selected samples. As we have to produce both historical and future projections, we have tackled this effort with two models: one for predicting historical data and one for predicting future data. Since the performance metrics for the historical projection model and the future projection model are nearly identical, we use the average performance metrics for both models. Following this extensive training phase, the model's performance was assessed, yielding a Mean Intersection over Union (mIoU) score of 0.805. This metric corresponds to an overall classification accuracy of 94.67%, achieved under a specific configuration that involved a 60% masking strategy during the evaluation.

The training is carried out on the MareNostrum 5 accelerated partition. We used a single node with 4 Nvidia H100 GPUs. In each epoch, we trained on 10,000 batches of 8 samples and validated on 2,000 batches of 8 samples. Figure 5 shows the loss changes per epoch for training and validation for both forward and backward predictions. The training loss is much higher in comparison to validation, as we utilised regularization. Each epoch took roughly 2 hours and 40 minutes for training, validation, and I/O. We used Distributed Data Parallelism with the PyTorch framework and participated in the EuroCC Open Hackathon to optimise our code.

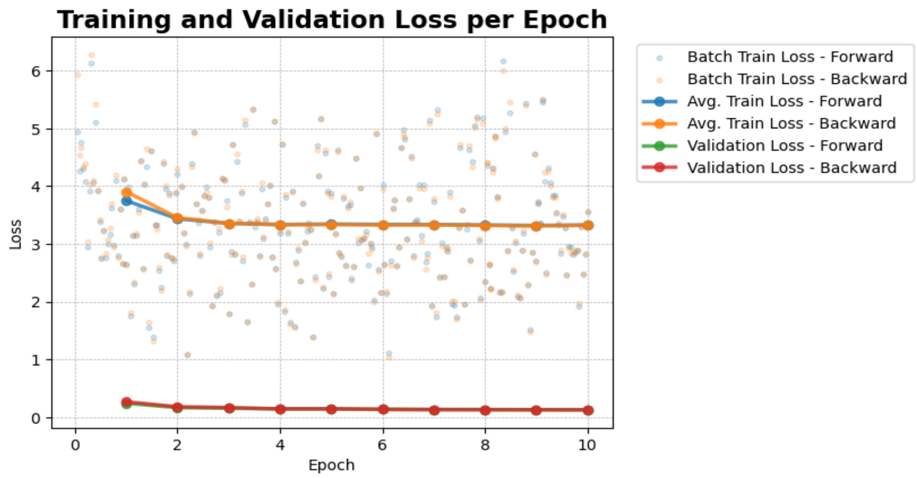

Figure 5. Training and validation loss for both forward and backward predictions

There is a performance disparity that was dramatically amplified for those land-use types that were underrepresented in the training dataset. For these minority classes, the model's ability to generalise and accurately predict decreased substantially, resulting in a notable drop in the IoU, in the urban class, to 0.46. This outcome underscores the model's sensitivity to class imbalance and highlights the need for further strategies to improve generalization across all LU types, particularly those with fewer training examples. Table 3 shows the accuracy (%), F1 score and IoU across the classes.

Table 3. Global Evaluation Results. Average loss: 0.1434

Class	Accuracy (%)	F1	IoU
Urban	48.68	0.632	0.463
Cropland	90.62	0.897	0.815
Pasture	89.52	0.897	0.814
Forest	94.93	0.941	0.889
Grass/shrubland	81.47	0.833	0.715
Other land	97.05	0.971	0.944
Water	99.24	0.995	0.991

To provide a concrete illustration of the model's output and the complexities of the classification task, Figure 6 presents a qualitative example. Figure 3 illustrates the multi-band input channels provided to the U-Net, the ground truth or ‘masked prior’ used for training and evaluation, and the final resulting model predictions. Importantly, these visual examples are provided for two consecutive years, 2004 and 2005, which helps to qualitatively assess the model's consistency and ability to handle temporal variations within the classification task.

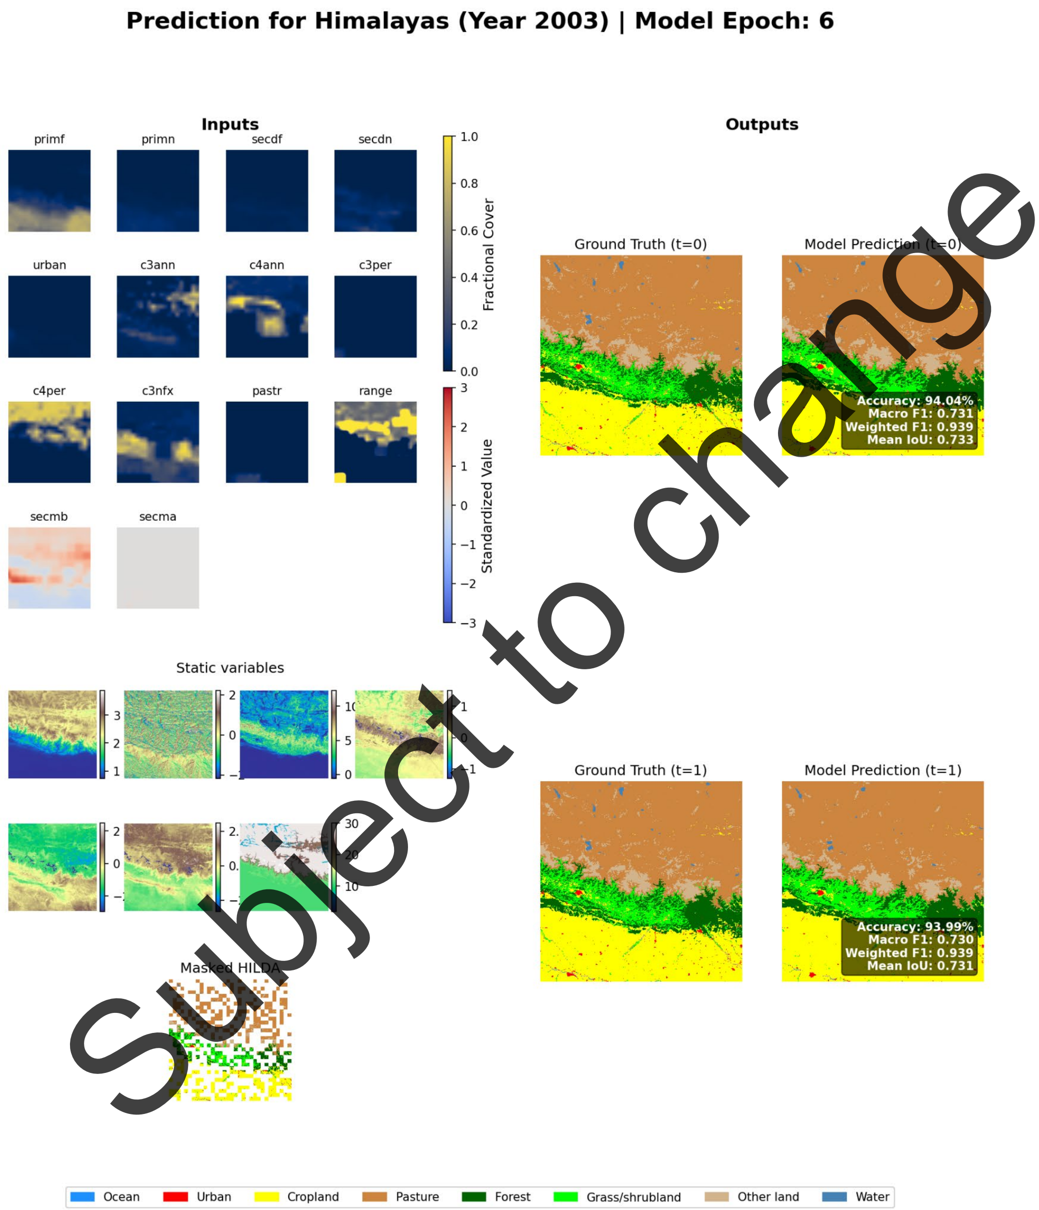

Figure 6. Qualitative results for the Himalayas region (Year 2003). The figure displays the full stack of input channels, including 12 coarse fractional LU inputs (top-left), 6 high-resolution static variables (middle-left), and the masked HILDA prior (bottom-left). The model's multi-step predictions (right) for two time steps (t=0, t=1) show strong agreement with the ground truth, achieving ~94 percent accuracy and a Mean IoU of ~0.73.

4.2. Inference

The execution of this pipeline produced a comprehensive high-resolution dataset spanning the period 1850–2100, stored in both ARCO Zarr and NetCDF formats. The primary historical output integrates model predictions for the pre-observation era (1850–1899) with consolidated HILDA+ ground truth for the observational record (1899–2020). Furthermore, the system successfully generated three distinct future projection cubes corresponding to the SSP2-4.5, SSP3-7.0, and SSP4-6.0 climate scenarios. Validation of the inference stage, specifically for the year 2014 as it shows in Figure 7, demonstrated high fidelity to the ground truth, achieving a global accuracy of 94.88% and a mIoU of 0.8569, confirming the framework's ability to accurately reconstruct complex land-use patterns on a global scale.

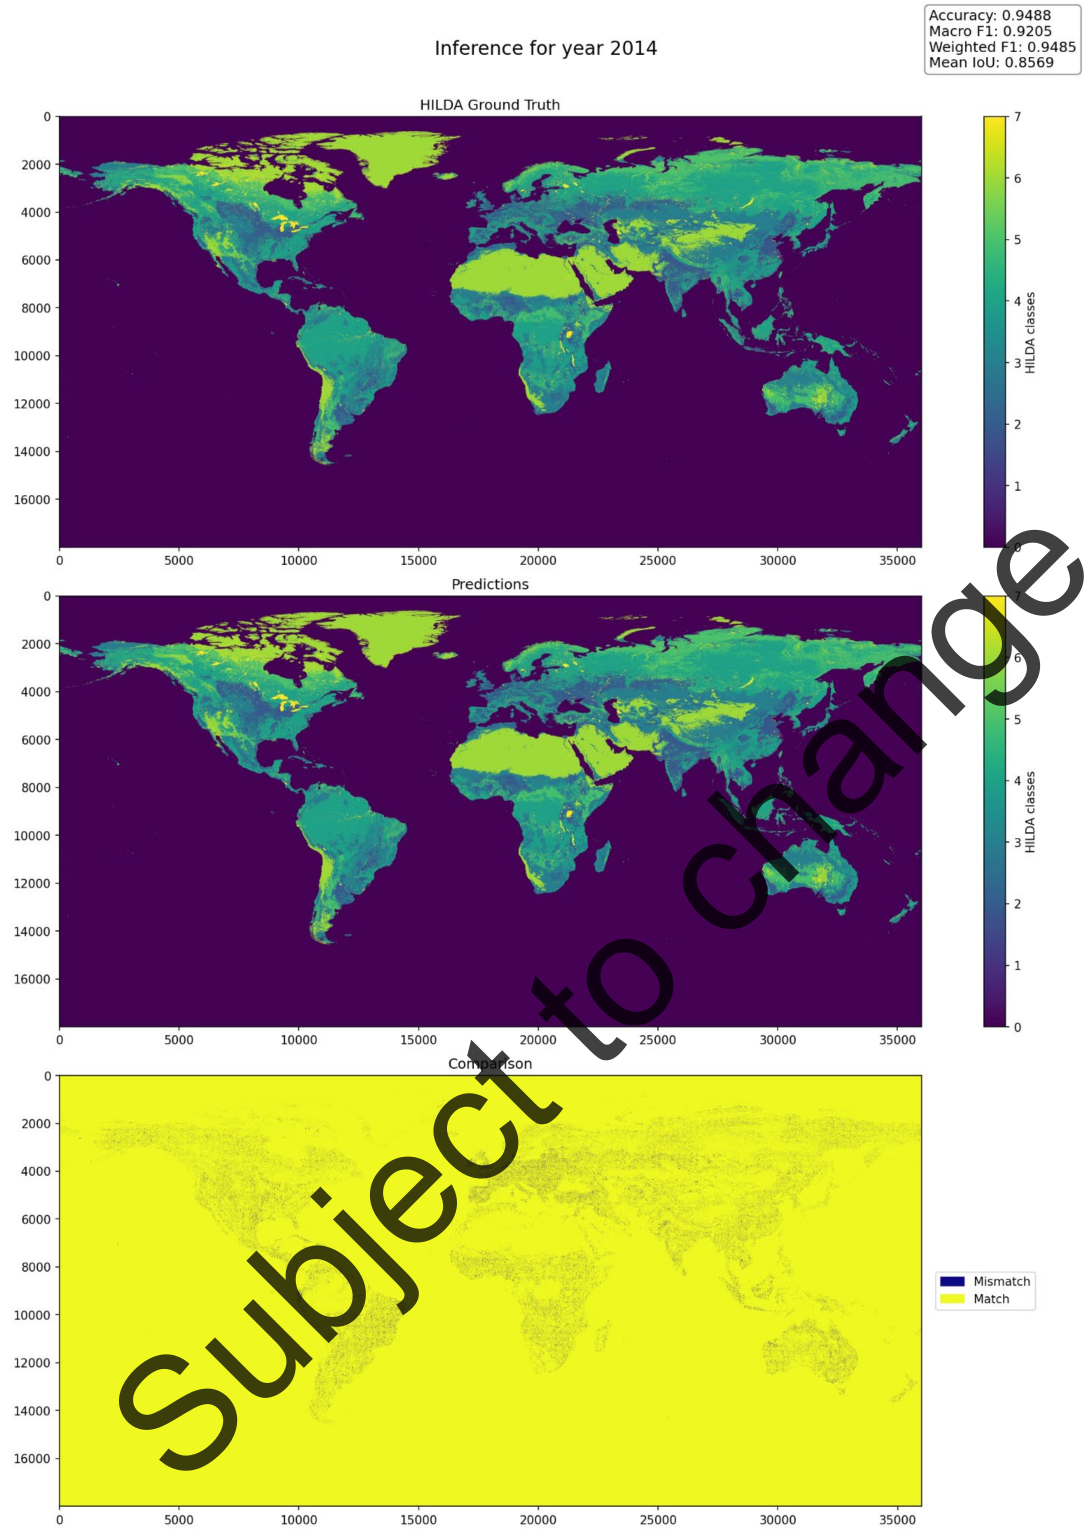

Figure 7. Worldwide inference results (Year 2014). The figure shows the ground truth LU data (top), the predictions with our UNet model (middle), and a heat map of the similarities between these two (bottom).

4.3. High-resolution global map of land use change

The inference workflow produced a single 3D Zarr store that contains LU predictions from 1850 to 1899, followed by ground truth LU with a reduced number of classes from 1899 to 2020 for the historical period, with the name 'lu_historical'. In addition, we have created three Zarr stores for SSP2 - 4.5, SSP3 - 7.0 and SSP4 - 6.0, named 'lu_future_ssp2', 'lu_future_ssp3', 'lu_future_ssp4', respectively. We also created the NetCDF4 version of these files with the same naming convention.

5. Discussion and Conclusion

In conclusion, we have successfully addressed the scaling mismatch in current land surface modelling by creating a consistent, high-resolution (1 km) LU dataset covering the historical period (1850–2020) and three future scenarios (up to 2100). Qualitative analysis confirms the model's ability to maintain temporal consistency and smooth spatial transitions without edge artifacts.

We acknowledge that the CONCERTO LU dataset is inherently impacted by the quality and biases of the source data, specifically the uncertainty in HILDA+ (ground truth) and the forcing data from LUH2. Consequently, our model reflects these inherent biases. While global accuracy is high (94.88%), performance dropped for underrepresented classes, specifically urban areas.

This current dataset (v0.1) serves as a baseline. To address the identified class imbalances, we have developed a methodology to incorporate dynamic annual Köppen-Geiger climatic zones (1 km) and ISIMIP population data (5 km) as new predictors. We expect these additions to significantly improve the urban class classification. These improved data and models will form the basis for prescribing Land Cover and Leaf Area Index for Deliverable D2.2.

6. Acknowledgements

This work received funding from the European Union's Horizon Europe Framework Programme through the project CONCERTO (Grant Agreement 101185000).

This work, as part of TerraDT - Digital Twin of Earth System for Cryosphere, Land Surface and related interactions, has received funding from the European Union's Horizon Europe Framework Programme (HORIZON) under Grant Agreement no. 101187992.

D2.1 [Creation of a high-resolution global map of land use change covering the historical and the scenario period (1850-2100)]
#101185000

The authors acknowledge the technical support of the EuroCC National Competence Centre. The authors also thank the organisers and participants of the Open Hackathon EuroCC AI Hackathon, especially Milos Maric and Simeon Harrison, for their contribution in optimising the training algorithm.

We acknowledge EuroHPC Joint Undertaking for awarding the project ID EHPC-DEV-2025D03-112 access to MareNostrum5 at BSC, Spain. We acknowledge EuroHPC Joint Undertaking for awarding the project ID EHPC-AIF-2025SC01-004 access to MareNostrum5 at BSC, Spain.

7. References

- Aboh, B. and Mutabazi, A., 2020. 'Satellite imagery analysis for Land Use, Land Use Change and Forestry: A pilot study in Kigali, Rwanda', NeurIPS 2020 Workshop on Tackling Climate Change with Machine Learning. https://www.climatechange.ai/papers/neurips2020/47
- Booth, B. B. B., et al., 2012. 'High sensitivity of future global warming to land carbon cycle processes', Environmental Research Letters, 7(2), pp. 024002.
- Boussetta, S., et al., 2013. 'Impact of a satellite-derived leaf area index monthly climatology in a global numerical weather prediction model', International journal of remote sensing, 34(9-10), pp. 3520-3542.
- Brasika, I. B. M., et al., 2024. 'Uncertainties in carbon emissions from land use and land cover change in Indonesia', EGUsphere, 2024, pp. 1-24.
- Chen, M., et al., 2020. 'Global land use for 2015-2100 at 0.05 resolution under diverse socioeconomic and climate scenarios', Scientific Data, 7(1), pp. 320.
- Chen, T. and Guestrin, C., 2016. 'Xgboost: A scalable tree boosting system', Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining, pp. 785-794.
- Chen, Z., et al., 2021. 'Integrating remote sensing and a markov-FLUS model to simulate future land use changes in Hokkaido, Japan', Remote Sensing, 13(13), pp. 2621.
- Copernicus Climate Change Service, 2019. 'Land cover classification gridded maps from 1992 to present derived from satellite observation', Copernicus Climate Change Service (C3S) Climate Data Store (CDS). https://doi.org/10.24381/cds.006f2c9a
- Copernicus Global Land Service / EEA, 2020. 'Leaf Area Index 1999–2020 (raster 1 km), global, 10-daily – version 2', Copernicus Land Monitoring Service. https://doi.org/10.2909/d5fdc595-2e03-4cbe-a39e-5f006f9cef07
- Copernicus Global Land Service / JRC, 2023. 'Leaf Area Index 2014–present (raster 300 m), global, 10-daily – version 1', Copernicus Land Monitoring Service. https://doi.org/10.2909/219fdc9f-616b-444b-a495-198f527b4722
- Fang, H., et al., 2019. 'An overview of global leaf area index (LAI): Methods, products, validation, and applications', Reviews of Geophysics, 57(3), pp. 739-799.
- Feng, L., et al., 2025. 'Convolutional neural network-based deep learning for landslide susceptibility mapping in the Bakhtegan watershed', Scientific Reports, 15(1), pp. 13250.

D2.1 [Creation of a high-resolution global map of land use change covering the historical and the scenario period (1850-2100)]
#101185000

- Friedlingstein, P., et al., 2025. 'Global Carbon Budget 2025', Earth System Science Data Discussions, 2025, pp. 1-139.
- Frieler, K., et al., 2017. 'Assessing the impacts of 1.5 C global warming--simulation protocol of the Inter-Sectoral impact model intercomparison project (ISIMIP2b)', Geoscientific Model Development, 10(12), pp. 4321-4345.
- Gao, P., et al., 2023. 'CLUMondo-BNU for simulating land system changes based on many-to-many demand–supply relationships with adaptive conversion orders', Scientific Reports, 13(1), pp. 5559.
- GEBCO Compilation Group, 2024. 'GEBCO 2024 Grid', Distributed by GEBCO, British Oceanographic Data Centre. https://doi.org/10.5285/1c44ce99-0a0d-5f4f-e063-7086abc0ea0f
- Gier, B. K., et al., 2024. 'Representation of the terrestrial carbon cycle in CMIP6', Biogeosciences, 21(22), pp. 5321-5360.
- Guo, C. and Berkhahn, F., 2016. 'Entity Embeddings of Categorical Variables', arXiv. https://arxiv.org/abs/1604.06737
- Hersbach, H., et al., 2020. 'The ERA5 global reanalysis', Quarterly journal of the royal meteorological society, 146(730), pp. 1999-2049.
- Hoffmann, P., et al., 2022. 'High-resolution land use and land cover dataset for regional climate modelling: historical and future changes in Europe', Earth System Science Data Discussions, 2022, pp. 1-50.
- Hurtt, G. C., et al., 2020. 'Harmonization of global land-use change and management for the period 850--2100 (LUH2) for CMIP6', Geoscientific Model Development Discussions, 2020, pp. 1-65.
- Jia, Y., et al., 2019. 'Super-resolution land cover mapping based on the convolutional neural network', Remote Sensing, 11(15), pp. 1815.
- Jian, J., et al., 2022. 'Historically inconsistent productivity and respiration fluxes in the global terrestrial carbon cycle', Nature Communications, 13(1), pp. 1733.
- Khikmah, F., et al., 2024. 'Modelling Vegetation Health and Its Relation to Climate Conditions Using Copernicus Data in the City of Constance', Remote Sensing, 16(4), pp. 691.
- Lee, J., Smith, C. and Wang, H., 2023. 'LandRes: High-resolution land cover and vegetation datasets improve Earth system model predictions', Geophysical Research Letters, 50(6), pp. e2022GL101234.
- Li, H., et al., 2024. 'A dataset of 0.05-degree leaf area index in China during 1983--2100 based on deep learning network', Scientific Data, 11(1), pp. 1122.

- Li, L., et al., 2024. 'Global 1 km land surface parameters for kilometer-scale Earth system modeling', Earth System Science Data, 16(4), pp. 2007-2032.
- Li, Z., et al., 2024. 'Learning without Exact Guidance: Updating Large-scale High-resolution Land Cover Maps from Low-resolution Historical Labels', arXiv. https://arxiv.org/abs/2403.02746
- Liu, T., et al., 2022. 'Estimation of vegetation leaf-area-index dynamics from multiple satellite products through deep-learning method', Remote sensing, 14(19), pp. 4733.
- Luo, M., et al., 2022. '1 km land use/land cover change of China under comprehensive socioeconomic and climate scenarios for 2020--2100', Scientific data, 9(1), pp. 110.
- Lv, J., et al., 2025. 'Land system changes of terrestrial tipping elements on Earth under global climate pledges: 2000--2100', Scientific Data, 12(1), pp. 163.
- Ma, H. and Liang, S., 2022. 'Development of the GLASS 250-m leaf area index product (version 6) from MODIS data using the bidirectional LSTM deep learning model', Remote sensing of environment, 273, pp. 112985.
- Mozaffari, A., et al., 2025a. 'Historical Reconstruction and Future Projection of Land Surface Boundary Conditions', NeurIPS 2025 Workshop on Tackling Climate Change with Machine Learning.
- Mozaffari, A., et al., 2025b. 'Reconstruction and downscaling of historical land surface boundary conditions with Machine Learning', Copernicus Meetings.
- Muñoz-Sabater, J., et al., 2021. 'ERA5-Land: A state-of-the-art global reanalysis dataset for land applications', Earth system science data, 13(9), pp. 4349-4383.
- Newman, D. J., 2024. 'Zarr Storage Specification Version 2: Cloud-Optimized Persistence Using Zarr', NASA Earth Science Data and Information System. https://doi.org/10.5067/DOC/ESCO/ESDS-RFC-048v1
- NOAA Global Monitoring Laboratory, 2025. 'Monthly Average Mauna Loa CO2: June 2025 = 429.61 ppm', NOAA Global Monitoring Laboratory. https://gml.noaa.gov/ccgg/trends/
- Prentice, I. C., et al., 2024. 'Principles for satellite monitoring of vegetation carbon uptake', Nature Reviews Earth & Environment, 5(11), pp. 818-832.
- Ronneberger, O., Fischer, P. and Brox, T., 2015. 'U-Net: Convolutional Networks for Biomedical Image Segmentation', Medical Image Computing and Computer-Assisted Intervention (MICCAI), pp. 234-241. https://arxiv.org/abs/1505.04597

- Rumelhart, D. E., Hinton, G. E. and Williams, R. J., 1986. 'Learning representations by back-propagating errors', Nature, 323(6088), pp. 533-536.
- Schulzweida, U., 2019. 'CDO User Guide', Zenodo. https://doi.org/10.5281/zenodo.3539275
- Shi, X., et al., 2015. 'Convolutional LSTM network: A machine learning approach for precipitation nowcasting', Advances in neural information processing systems, pp. 802-810.
- Sun, M., et al., 2023. 'Reconstruction of a monthly 1 Km NDVI time series product in china using random forest methodology', Remote Sensing, 15(13), pp. 3353.
- Truong, V. T., et al., 2024. 'JAXA's new high-resolution land use land cover map for Vietnam using a time-feature convolutional neural network', Scientific Reports, 14(1), pp. 3926.
- Van Jaarsveld, B., Hauswirth, S. M. and Wanders, N., 2024. 'Machine learning and global vegetation: random forests for downscaling and gap filling', Hydrology and Earth System Sciences, 28(11), pp. 2357-2374.
- Winkler, K., et al., 2020. 'HILDA+ Global land use change between 1960 and 2019 [Dataset]', PANGAEA.
- Wolfram, S., 1983. 'Statistical mechanics of cellular automata', Reviews of modern physics, 55(3), pp. 601.
- Xu, T., Gao, J. and Coco, G., 2019. 'Simulation of urban expansion via integrating artificial neural network with Markov chain–cellular automata', International Journal of Geographical Information Science, 33(10), pp. 1960-1983.
- Yang, J. and Tao, B., et al., 2020. 'Integration of remote sensing, county-level census, and machine learning for century-long regional cropland distribution data reconstruction', International Journal of Applied Earth Observation and Geoinformation, 91, pp. 102151.
- Yang, J. and Huang, X., 2021. '30 m annual land cover and its dynamics in China from 1990 to 2019', Earth System Science Data Discussions, 2021, pp. 1-29.
- Zhao, S. Q., et al., 2010. 'A spatial resolution threshold of land cover in estimating terrestrial carbon sequestration in four counties in Georgia and Alabama, USA', Biogeosciences, 7(1), pp. 71-80.
- Zhao, X., et al., 2023. 'Combining LSTM and PLUS models to predict future urban land use and land cover change: A case in Dongying City, China', Remote Sensing, 15(9), pp. 2370.

- Zhu, Y. and Newsam, S., 2015. 'Land use classification using convolutional neural networks applied to ground-level images', Proceedings of the 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems, pp. 1-4.