

# m14xcube: Machine Learning Toolkits for Earth System Data Cubes

Julia Peters<sup>1,2</sup>, Anja Neumann<sup>1,2</sup>, Marco Jaeger<sup>1,2</sup>, Lukas Gienapp<sup>1,2</sup>, Josefine Umlauf<sup>1,2</sup>

<sup>1</sup>Leipzig University, Leipzig, Germany

<sup>2</sup>Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI), Leipzig University, Germany  
julia.peters@informatik.uni-leipzig.de, josefine.umlauft@uni-leipzig.de

## Abstract

Rapidly changing climate conditions and the increase in extreme events are posing severe challenges to human life and infrastructure, requiring sophisticated analytical capabilities for hazard prediction and disaster risk management. Earth System Data Cubes (ESDCs) have become an essential tool in Earth System Sciences (ESS) by organizing large-scale, multivariate environmental datasets into a structured, scalable and analysis-ready format. However, modern machine learning techniques are not yet being utilized to their full potential on ESDCs. This is due to the lack of proper tooling, domain-specific challenges, and high barriers of entry for practitioners. We introduce `m14xcube`, an open-source Python framework designed to assist ESS domain experts in applying ML techniques on ESDCs for advanced analysis and prediction of environmental variables and impacts. Through a comprehensive suite of tools, it addresses specific challenges associated with the nature of ESS data, such as the non-uniform data distribution due to dynamic gaps, or spatio-temporal autocorrelation of environmental variables. Due to its modular architecture, it covers the complete analysis process, from data exploration, and preparation, to model development, result interpretation and evaluation. With support for distributed computing, it handles large ESDC datasets efficiently. In order to ease the adoption it includes extensive documentation and tutorial notebooks. We demonstrate `m14xcube`'s capabilities through three examples, showcasing its potential and capabilities for integrating machine learning with ESDC data.

**Code** — <https://github.com/deepesdl/ML-Toolkits>

## 1 Introduction

Extreme events, such as heatwaves, droughts, heavy precipitation and tropical cyclones have become increasingly frequent and severe across the entire globe (IPCC 2023). They pose substantial threats to human lives and infrastructure (Noy 2016; Fowler et al. 2024), as well as to natural ecosystems and biodiversity (Flach et al. 2020; Mahecha et al. 2024). Especially, people living in regions with development constraints suffer from reduced food and fresh water security, leading to a 15 times higher mortality rate due to extreme events compared to less vulnerable regions (IPCC 2023). Impacts of recent catastrophic events, such

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

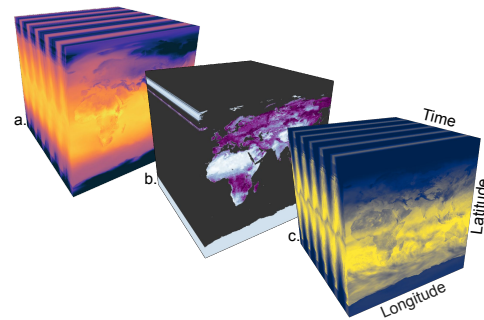


Figure 1: Lexcube visualization (Söchting et al. 2023) of an Earth System Data Cube displaying the environmental variables (a) air temperature, (b) soil moisture and (c) radiation aligned along common spatial and temporal coordinates.

as the 2018 European Compound Heatwave and Drought (Rousi et al. 2023) and the 2021 Pacific Northwest Heatwave (White et al. 2021), highlight the need for precise predictive models and proactive disaster management. To this end, advances in remote sensing technologies have revolutionized Earth System Sciences (ESS) by collecting vast amounts of Earth observational data through drones, aerial systems and satellites. These technologies provide crucial information for understanding climate extremes, with a single satellite capturing up to terabytes of data each day (Schieler et al. 2019). However, managing and analyzing this abundance of data in order to answer the relevant scientific questions poses substantial challenges for Earth System scientists.

Earth System Data Cubes (ESDCs, Mahecha et al. 2020) have emerged as a viable solution. They organize extensive multidimensional datasets into structured grids, where various environmental variables are mapped to the same spatio-temporal coordinate system. This unified structure is essential for efficiently synchronizing and correlating diverse environmental datasets. Figure 1 shows exemplary variables stored in one ESDC, here tracking air temperature, soil moisture, and solar radiation measurements. ESDCs were developed as part of the Deep Earth System Data Lab (DeepESDL) project<sup>1</sup>, utilizing Earth observation variables acquired by the European Space Agency. DeepESDL provides

<sup>1</sup><https://www.earthsystemdatalab.net/>

an online platform for accessing, managing, analyzing, and visualizing large-scale environmental data. Fully committed to the FAIR principles (Wilkinson et al. 2016), DeepESDL enables public access to the ESDCs through a user-friendly API, establishing them as one of the primary ways for interacting with large-scale Earth observational data.

Breakthroughs in machine learning (ML) have revolutionized numerous domains (Sarkar, Shiuly, and Dhal 2024; Islam 2024). Specifically in ESS ML offers significant potential (Hsieh 2022; Reichstein et al. 2019) including training neural networks to learn the response of the vegetation to climate drivers (Martinuzzi et al. 2023), assessing drought vulnerability (Saha et al. 2021), forecasting weather using deep learning (Pasche et al. 2024), or modelling urban heat distribution with Random Forests (Zumwald et al. 2021). However, as for now, the large set of multivariate data that ESDCs offer are not being utilized to their full potential. This is, on the one hand, due to the nature of environmental sensor data. The inherent data characteristics and heterogeneity amongst variables together with the data’s extensive volume hinder traditional analysis approaches. Further, the Earth’s spherical shape together with external factors such as satellite specific trajectories and cloud cover may result in autocorrelation among data points (Loaiza et al. 2023) and a non-uniform data distributions due to data gaps (Sarafanov et al. 2020). Moreover, domain knowledge is required to generate more plausible predictions (Kerrigan, Hullman, and Bertini 2021). Hence, there is a strong need to reduce the technical barrier of existing tools for ML methods on ESDCs.

In order to improve the adoption of ML techniques on ESDCs, we introduce `m14xcube`, a Python-based open-source toolkit specifically developed to pave the way for straight forward ML analysis on ESDCs. It facilitates data handling and processing for domain experts in ESS while respecting the fundamental challenges inherent to Earth observational data. `m14xcube` features a modular architecture, encompassing modules dedicated to the various stages of an ML pipeline, from data preparation and exploration to model training and result interpretation. It offers a unified API that simplifies interaction with complex data structures and ML models, making it accessible to domain experts. `m14xcube` is designed to address specific challenges in ML-based remote sensing data analysis. For instance, to manage gaps in data due to cloud cover and satellite orbit paths, it integrates tools to ensure data is complete and analytically viable before entering the modeling phase. It supports the application of filters that help researchers focus their analysis on relevant regions or timeframes. For managing autocorrelation in spatial data and prevent overfitting, it offers a tailored train-test split methodology, ensuring the training and testing subsets reflect true data characteristics. With support for multi-GPU setups, large datasets can be leveraged for insight.

`m14xcube` facilitates standard ML workflows while accommodating to the requirements of ESS and its practitioners. It supports the development of predictive models and strategies for a wide range of Earth system use cases by enabling more effective data analysis, with the primary objective of estimating and preventing the social and environmental impacts of extreme events.

## 2 Background & Related Work

We structure our review of related work into three parts. Initially, we revisit the formal definition of ESDCs and how they are commonly implemented. Next, we explore possible use cases for both ESDCs and ML techniques in ESS. Finally, we review existing solutions and frameworks suitable for analyzing multivariate environmental datasets including ESDCs.

### 2.1 Earth System Data Cubes

Formally, an ESDC  $C$  is a triplet  $(L, G, X)$  (Mahecha et al. 2020). Here,  $L$  is a set of axis labels, describing the cube’s dimensions and  $G = \text{grid}(l)_{l \in L}$  is a collection of grids, specifying the discrete points along the domain of its axis label  $l$ . These grids determine the resolution and locations of data. The entire collection  $G$  defines multidimensional indices, or grid points, at which the data is stored. Data corresponding to these indices are stored in  $X$ . Operations on ESDC variables typically result in new data cubes with potentially different labels, grids, or data (Mahecha et al. 2020).

In order to implement the concept of an ESDC, `xarray` (Hoyer and Hamman 2017), a Python package that efficiently manages extensive multidimensional datasets in structured grids, is commonly used (Mahecha et al. 2020). It allows to align variables along spatio-temporal coordinates and supports lazy loading, which enables to load data on-demand to minimize memory usage and speed up operations. `xarray` also implements chunking, i.e., dividing the dataset into smaller, manageable blocks for independent processing. It further allows for parallel computing with `dask` (Rocklin et al. 2015). `m14xcube` also adopts `xarray` as its primary data backend through the `xcube`<sup>2</sup> compatibility layer, streamlining the analysis of ESDCs and increasing interoperability with existing tools.

### 2.2 Use cases for ESDCs

The scalable structure of ESDCs facilitates the management of large and complex datasets, allowing for data exploration and long-term monitoring of continuously growing environmental data. ESDCs excel in integrating various types of Earth observation data into a single consistent view, including, but not limited to, atmospheric, terrestrial or hydro-spheric parameters. This allows scientists to analyze multiple data variables simultaneously, leading to a more comprehensive understanding of the complex interactions within the Earth system (Mahecha et al. 2020).

The utilization of a common spatio-temporal grid within ESDCs ensures that data from different sources are standardized. This is crucial for researchers comparing and combining a set of different data variables for common analysis (Montero et al. 2024). By transforming raw satellite data into analysis-ready data cubes, ESDCs facilitate data management and accessibility, making them immediately usable for analysis. This allows researchers to efficiently query and retrieve data, significantly speeding up the analysis process (Baumann et al. 2019).

<sup>2</sup><https://xcube.readthedocs.io/en/latest/>

Applying ML algorithms to ESDCs allows for answering ESS related research questions in a data-driven way, e.g. when studying how vegetation reacts to climate drivers (Martinuzzi et al. 2023) or quantifying drought legacy effects and their temporal dynamics on gross primary production (Yu et al. 2022). Furthermore, large multivariate remote sensing datasets enable the detection of patterns and trends in climate data (Liu et al. 2012; Winkler et al. 2021), which are essential for predicting future climate scenarios and understanding long-term environmental changes (Mahecha et al. 2020). It has been demonstrated, that anomalies (Flach et al. 2017) and extreme events (Mahecha et al. 2017) can be detected, which are significant environmental scenarios that require further investigation.

### 2.3 Tooling for ESDCs

Multiple ML-based tools have been developed well-suited for the analysis of ESDCs since their inception. XCast (Hall and Acharya 2022) allows users to train ML models directly on gridded datasets with minimal preprocessing, leveraging high-performance computing methods like chunk-wise parallelism and cluster computing through `dask`. Its API mirrors traditional Python data science tools, facilitating an easy transition for users. Similarly, the `nd` framework (Hansen 2022) specializes in analyzing n-dimensional data cubes, facilitating the integration of ML techniques by providing interfaces to Python’s scientific ecosystems. Advanced visualization tools like `lexcube` (Söchting et al. 2023) and `vapor` (Li et al. 2019) enhance the interpretability of large-scale ESDC datasets by providing interactive and multi-dimensional visualization capabilities, allowing researchers to gain deeper insights into environmental processes and phenomena. `scores` (Leeuwenburg et al. 2024) is a framework for quantitative evaluation, offering over 50 metrics, statistical techniques, and data processing tools designed to verify and evaluate models in the geosciences, particularly those developed using multivariate `xarray` data, similar to ESDCs.

However, existing tools face significant shortcomings. They lack adequate support for the data-driven analysis of ESDCs and tend to address only specific tasks rather than offering a comprehensive framework that encompasses the entire ESDC analysis lifecycle, including data management, exploration, modeling, and evaluation. Additionally, these tools typically do not support advanced ML techniques such as deep learning, instead focusing on traditional ML algorithms. `ml4xcube` aims to overcome these limitations by providing an integrated approach towards advanced ML on ESDCs, with extensive support for domain-specific challenges and the integration of modern ML tools.

## 3 ml4xcube

`ml4xcube` streamlines the entire ML pipeline for ESDC data. It thus encompasses tooling for data exploration, preparation, processing, model training, and result interpretation. Each of these steps is associated with modules within the `ml4xcube` framework, as illustrated in Figure 2. This section explores these components and illustrates the

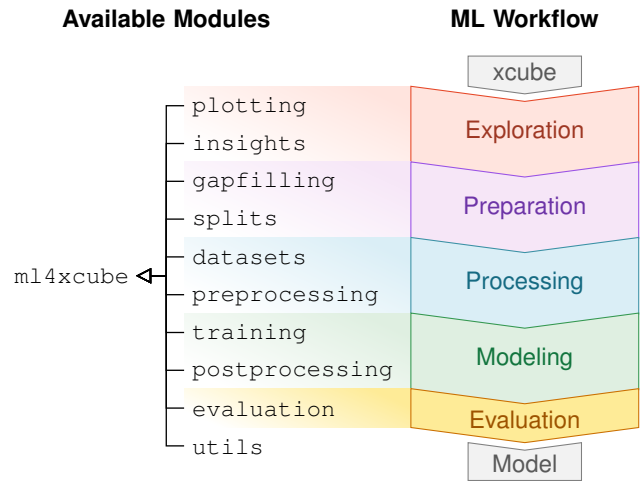


Figure 2: Module overview of `ml4xcube`, along with the corresponding ML workflow for deriving a model from an ESDC, loaded with `xcube`.

domain-specific capabilities they offer to support ESS researchers at every step of their workflow, spanning from the initial dataset loading with `xcube` to the final model.

### 3.1 Exploration

To allow researchers an initial view into the data and develop a first understanding of key characteristics, `ml4xcube` offers support for both visual and quantitative analysis through its `plotting` and `insights` modules.

**Plotting** The `plotting` module provides methods for creating and customizing ESDC visualizations. It allows to display specific variable slices from a data cube, incorporating features such as coast lines if a mask is applied. This functionality aids in the detection of patterns and anomalies in data during initial assessment.

**Insights** The `insights` module provides quantitative insight into a data cube’s characteristics, including statistical measures of data completeness, data distribution, and data quality. It offers detailed insights into the different data dimensions, value ranges, and coverage of a given cube. This module enables the computation of heatmaps representing the availability of data for identifying patterns and gaps in data coverage. Figure 3 presents such a heatmap for the land surface temperature variable across the time dimension visualized using the `plotting` module. It indicates the amount of available data points per spatial coordinate ranging from 0 to 10. Notably, the visible streaks correspond to the satellite’s paths during data collection, highlighting variations in data capture due to the satellite’s trajectory. Leveraging these insights, researchers can strategically decide on the necessary data preparation measures. For example, they might choose to implement gap filling to enhance data integrity or opt to omit incomplete data samples, focusing on fully available datasets instead.




Figure 3: Heatmap of data gaps in the land surface temperature variable over time, generated using `insights` module. The number of available data ranges from 0 to 10, corresponding to the 10 frames in the analyzed cube.

### 3.2 Preparation

In order to prepare a data cube for the modeling stage, `ml4xcube` includes two essential components: a module for gap filling, which extends beyond the conventional statistical methods found in the `preprocessing` module (Section 3.3), and a module for constructing train/test splits. Both components address key challenges specific to Earth systems model training.

**Gap Filling** The `gapfilling` module provides a method to address missing data in ESDCs, specifically designed for remote sensing datasets (Sarafanov et al. 2020). Figure 4 illustrates an example of a cube with the land surface temperature (LST) variable before and after gap filling. This method employs support vector regression (SVR) to predict missing values leveraging available latitude/longitude pixels. Each pixel in the dataset is treated individually, with a dedicated SVR model trained for this missing value. The training process uses surrounding coordinates within the same time slice as features and corresponding values from other slices as training data. For effective training, at least 50 valid neighboring points within a slice are required. In scenarios where this condition is not met, nearest neighbor interpolation is applied instead.

**Splitting** When dividing data into train and test set, remotely sensed datasets naturally challenge ML applications: they exhibit significant autocorrelation, where data points in close spatio-temporal vicinity share similar characteristics. Traditional random splitting may overlook these correlations, potentially introducing bias and affecting model generalizability (Sweet et al. 2023). To mitigate this issue, the `splits` module provides specialized splitting techniques, illustrated in Figure 5. The block split method segments data into contiguous blocks based on geographical and temporal proximity, assigning data points from these blocks to either training or test sets with a specific probability. This strategy keeps closely related data points together, reducing information leakage across the train-test divide and enhancing testing integrity. The module also offers a traditional random splitting option for scenarios where strict adherence to geospatial dependencies is not crucial.

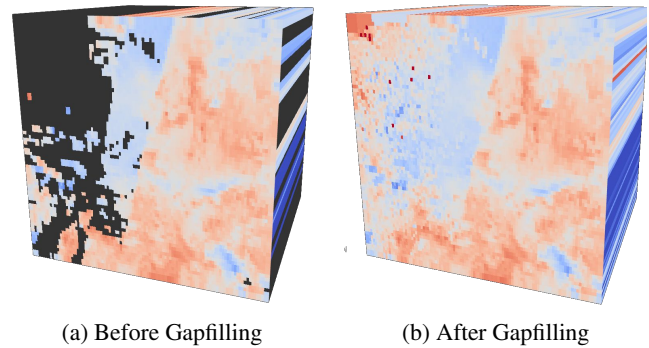


Figure 4: Lexcube visualization (Söchting et al. 2023) of a subcube from the ESDC's land surface temperature variable before (a) and after (b) gap filling.

### 3.3 Processing

While the preparation stage considers transforms of the content of the data cube, the processing stage is concerned with transforming the format, representation, and compatibility, i.e., technical aspects of the data cube. In order to easily get access to the required amounts of data for model training, `ml4xcube` provides the `datasets` and `preprocessing` submodules.

**Datasets** The `datasets` module provides tools for efficient preparation and management of geospatial data. The key purpose of this module is to expose cube-like datasets of varying sizes and complexities via a standardized API to models. It further offers a unified interface for downstream processing, supporting the application of filters and the definition of custom callback functions for additional processing steps. Thus, data cubes can be integrated seamlessly with the major ML frameworks PyTorch (Imambi, Prakash, and Kanagachidambaresan 2021) and TensorFlow (Abadi et al. 2016) to ensure optimal data formatting for various analytical tasks.

**Preprocessing** The `preprocessing` module implements common data processing tasks on ESDCs. The primary purpose of this module is to restrict data to relevant areas while ensuring its availability and compatibility for analysis. This includes tasks such as masking, which al-




Figure 5: Visualization of training (blue) and test (white) data splits. The random split assigns data points randomly to training and test sets. The block split divides data into contiguous geographic blocks, maintaining spatial coherence.

lows for focused analysis on specific geographical regions or timeframes; filtering, which can exclude data that fails to meet certain criteria, like insufficient coverage; interpolation, which fills in missing values using methods like mean replacement or a constant value; and data normalization and standardization, which ensure consistency across different scales and features in multi-faceted datasets.

### 3.4 Modeling

`ml4xcube` is designed to simplify the training process on ESDC data. During the modeling phase, it functions as a link between ESDC data and ML models implemented in any of the supported third-party ML libraries.

**Training** The `training` module provides trainer classes that streamline the training process by interfacing between model and ESDC data. To implement models with high flexibility, at present, PyTorch, TensorFlow, and scikit-learn (Pedregosa et al. 2011) are supported. The module includes essential training functionalities such as early stopping, model checkpointing, and training progress visualization. To address the challenges posed by remote sensing data, which often includes vast amounts of high-resolution data, the module enables efficient handling of such datasets through distributed training across multiple GPUs or nodes, relying on a distributed data parallel approach. This can significantly reduce the time required for training and inference (Kim et al. 2019).

**Postprocessing** Following training and inference methods, the `postprocessing` module can be used to reverse preprocessing operations applied previously, such as standardization and normalization, to transform model predictions back to the original scale. This allows for clear and meaningful interpretation of ML outcomes and compatibility with downstream use of model predictions.

### 3.5 Evaluation & Utilities

Finally, `ml4xcube` provides functionality to evaluate the effectiveness of models, and includes an assortment of utilities to support common ESDC tasks.

**Evaluation** To assess the effectiveness of the developed models, the `evaluation` package provides an interface to standard metrics from sklearn, PyTorch, and TensorFlow, and also allows for custom evaluation functions to be supplied with this standard interface. This ensures a consistent and simple, yet flexible evaluation process.

**Utilities** The `utils` module provides helper functions to leverage the chunked characteristic of ESDCs backed by an `xarray` data structure. This includes functions for chunking cubes into different chunking schemes, fetching specific data chunks by their index, iteration over chunks, splitting of chunks into smaller segments with adjustable overlap. These are commonly occurring data operations which enable efficient processing and preparation of data samples for model input.

## 4 Use Cases and Applications

This section highlights the capabilities of `ml4xcube` through three use cases. The first use case serves as a simple introductory example, predicting LST using linear regression. This approach focuses on single-coordinate predictions, treating each data point independently without accounting for spatial or temporal dependencies. The second use case, also predicting LST, involves a convolutional neural network (CNN) for time series analysis, a common approach in ESS, utilizing temporal and spatial dimensions (Han et al. 2023; Zhao and Ji 2022; Shirvani, Abdi, and Goodman 2023). Although both tasks aim to predict LST, these use cases illustrate the progression from a basic model to a more advanced, temporally and spatially aware approach. The third use case showcases the parallelization of this time series analysis, emphasizing the parallel processing capabilities of `ml4xcube`.

For all the use cases presented, a prerequisite is that an ESDC `ds` is loaded. Here, we utilize the `xcube` Python package as demonstrated below, restricting the available data to the relevant variables and time frame:

```
from xcube.core.store import new_data_store

ds = (
    new_data_store(...).open_data("<file>.zarr")
    [
        "land_surface_temperature",
        "air_temperature_2m",
        ...
    ]
    .sel(time=slice("2002-05-21", "2002-08-01"))
)
```

### 4.1 Predictive Task

To prepare the loaded data cube for modeling, it is first pre-processed using a mask that is applied to isolate terrestrial regions. This step ensures that subsequent analyses are relevant only to areas where LST values are meaningful. Then, the data is segmented into blocks along the three axes *time*, *latitude*, and *longitude*, to consider the spatio-temporal characteristics of the data when conducting the train-test split.

```
from ml4xcube.preprocessing import assign_mask
from ml4xcube.splits import assign_block_split

ds = assign_mask(ds, land_mask)
ds = assign_block_split(
    ds = ds,
    block_size = [("time", 10), ("lat", 100), ("lon", 100)],
    split = 0.8
)
```

Subsequently, the dataset is divided into test and a training sets using a sampler. By default, the data is standardized to a uniform scale, with configurable options for using the available normalization function or custom scaling functions. The standardization parameters are stored in the sampler's `scale_params` attribute for future predictions.

NaN values are discarded based on prior analyses with the `insights` module, and a land mask, if assigned, is applied as a standard configuration to filter for relevant values.

```
from ml4xcube.datasets.xr_dataset import XrDataset

sampler = XrDataset(
    ds = ds,
    num_chunks = 3,
    to_pred = "land_surface_temperature"
)
train_ds, test_ds = sampler.get_datasets()
```

Following data preparation, a linear regression from the scikit-learn package is trained. The model is evaluated using any metrics supplied in the metrics dictionary, which can be obtained from the `evaluation` module. Once trained, the model is saved to the given path for future use.

```
from ml4xcube.training.sklearn import Trainer
from sklearn.linear_model import SGDRegressor

trainer = Trainer(
    model = SGDRegressor(),
    train_data = train_ds,
    test_data = test_ds,
    model_path = ...,
    metrics = {...}
)
sgd_reg = trainer.train()
```

## 4.2 Time Series Analysis

This analysis utilizes historical temporal samples from selected variables to predict the most recent one. In this example, four time frames are captured for each set of spatial coordinates in the ESDC, with the first three predicting the fourth. Each sample consists of a  $20 \times 20$  grid for latitude and longitude, across the four time steps. The `MultiProcSampler` leverages multiprocessing, making it particularly suited for handling large datasets. It is configured to replace gaps with the sample mean, while samples with all missing values are dropped. The processed train and test subsets are subsequently stored in `zarr` format (Nguyen et al. 2023), which is compatible with the PyTorch-specific dataset implementation within `ml4xcube`.

```
from ml4xcube.datasets.multiproc_sampler import
MultiProcSampler

sampler = MultiProcSampler(
    ds = ds,
    train_ds = "train.zarr",
    test_ds = "test.zarr",
    sample_size = [("time", 4), ("lat", 20), ("lon", 20)],
    nproc = 5,
    chunk_size = (32, 4, 20, 20),
    array_dims = ("samples", "time", "lat", "lon"),
    drop_nan = "if_all_nan",
    fill_method = "sample_mean"
)
train_ds, test_ds = sampler.get_datasets()
```

The dataset is further prepared for training by converting it into a PyTorch-specific dataset implementation that allows GPU integration. During this process, users can provide a custom function that is applied to each sample (`map_fn`). In this example, the custom function is designed for time-series sampling, selecting the first three time frames as predictors and the last time frame as the dependent variable.

```
from ml4xcube.datasets.pytorch import PTXrDataset,
prep_data_loader

train_set = PTXrDataset(train_ds)
test_set = PTXrDataset(test_ds)

train_loader, test_loader = prep_data_loader(train_set,
test_set, callback=map_fn)
```

A CNN model implemented in PyTorch is supplied to the trainer, which handles the entire training process. Standard training options can be specified, such as the choice of optimizer, loss function, and number of training epochs, allowing for flexibility and control during the modeling phase.

```
from ml4xcube.training.pytorch import Trainer

trainer = Trainer(
    model = cnn,
    train_data = train_loader,
    test_data = test_loader,
    optimizer = optimizer,
    loss = mse_loss,
    model_path = model_path,
    epochs = 50
)
trained_cnn = trainer.train()
```

## 4.3 Distributed Training

To conduct the example time series analysis in a distributed setup, only minor adjustments are required. The distributed training process needs to be initialized, and the dataset must be prepared with the `parallel` argument set to `True`. Apart from these adjustments, all other configurations remain consistent with those used for training on a single GPU, rendering distributed training accessible without introducing significant overhead.

```
from ml4xcube.training.pytorch_distributed import
ddp_init, Trainer

ddp_init()

train_loader, test_loader = prep_data_loader(...,
parallel=True)

trainer = Trainer(
    model = cnn,
    train_data = train_loader,
    test_data = test_loader,
    ...
)
trained_cnn = trainer.train()
```

## 5 Limitations and Future Plans

With `m14xcube`, we focus on the technical aspects of supporting ML research on ESDC data. We identify three major limitations with this approach. First, while `m14xcube` is designed to facilitate the development and deployment of ML models, it does not guarantee the performance, fairness, correctness, or safety of the models created using it. Researchers must carefully consider issues such as bias, transparency, and accountability when utilizing this framework to build ML systems, but also be mindful of the limitations of `m14xcube` itself. Second, `m14xcube` inherits limitations from the ESDC approach itself (Mahecha et al. 2020). One significant challenge is its equal treatment of data cube dimensions across spatial, temporal, and nominal variable dimensions. This approach, although flexible, can lead to inappropriate model applications that do not consider the unique nature of each dimension. Moreover, the need to reformat or remap data to fit within a common grid, especially when working with data produced at varying resolutions, can compromise data integrity and potentially lead to suboptimal ML outcomes. Third, despite efforts to enhance the accessibility of `m14xcube`, effectively utilizing `m14xcube` still requires a certain level of technical expertise in ML and data science, which can vary greatly depending on the specific task at hand. For example, users must carefully select appropriate strategies for data preparation and feature selection. Additionally, when training a neural network, users are responsible for designing a suitable architecture and configuring various aspects of the pipeline for optimal results.

While the first two are underlying limitations inherited from ML and ESS as research fields as a whole, we see several areas of future improvements for the usability and applicability `m14xcube`. Incorporating AutoML, or Automated ML, (Karmaker et al. 2021) offers a promising solution for supporting domain experts who may not have deep ML expertise. AutoML can automate different steps within the ML pipeline, including data preparation (Shah and Kumar 2019), feature engineering (Ravishankar and Battineni 2022), hyperparameter optimization (Vincent and Jidesh 2023), and neural architecture search (Salehin et al. 2024), which seeks to build well-performing neural architectures. These techniques have already proven instrumental in several ML applications, particularly in analyzing high-dimensional remote sensing data (Kheir et al. 2024; Wasala et al. 2024; Babaeian et al. 2021). Integrating AutoML capabilities into `m14xcube` could enhance its usability and accessibility for domain experts, allowing them to focus more on extracting scientific insights rather than on the technical complexities.

Apart from ESS, `m14xcube` could be expanded to support developers and scientists from various domains. In medicine, MRI data can be arranged as data cubes to improve disease detection and diagnosis (Schmale, Seidel, and Paul 2017). In urban planning, data cubes integrate various datasets such as traffic patterns, land use, and environmental monitoring to optimize city development and management (Dhu et al. 2019; Park et al. 2013). In energy management, data cubes consolidate information on energy consumption, and production to improve grid management and the efficiency of renewable energy sources (Noh et al. 2017; Grasso

et al. 2019). Additionally, in the field of astronomy, data cubes can be employed to organize and analyze data from telescopes, enhancing the study of celestial objects and phenomena (Perkins et al. 2014). For these diverse domain specific applications, primarily statistical and visual analyses were conducted to derive scientific insights. By expanding the capabilities of `m14xcube`, the tool can support specialists across diverse disciplines through the integration of ML with data cubes.

## 6 Conclusion

In this paper, we introduced `m14xcube`, a versatile Python-based toolkit designed to enhance the analysis of ESDC data. ESDCs enable the management of large, complex ESS datasets by consolidating diverse types of Earth observation data into a coherent format. They are instrumental in detecting and analyzing the dynamics of climate extremes and developing predictive models to aid disaster management and mitigation. `m14xcube` offers a comprehensive and flexible framework for the integration of ML in ESS, allowing its application to ESDC data. This was previously associated with difficulties, as specialized data processing techniques, such as managing non-uniform data distributions, handling spatio-temporal autocorrelation, and addressing data gaps are required for successful modeling of ESS data. `m14xcube` addresses these challenges by providing a framework that spans the entire ML pipeline, from data preparation to result interpretation. Its modular design features specialized tools for efficient ML workflows, including capabilities for data exploration, preparation, and processing, a unified training approach compatible with modern ML frameworks, and evaluation functionalities. Additionally, `m14xcube` supports deep learning and distributed computing, equipping it to handle the complexities of vast ESDC datasets effectively. Through three use cases, we demonstrated the practical applications and benefits of `m14xcube`. These examples illustrate how `m14xcube` simplifies the integration of ML techniques with ESDC data for Earth system researchers and practitioners. In addition to its technical capabilities, `m14xcube` prioritizes accessibility and user support. The toolkit is complemented by comprehensive documentation and tutorial notebooks.<sup>3</sup>

As an open-source toolkit, `m14xcube` also facilitates collaborative research, allowing for community-driven improvements and innovation. While requiring a certain level of ML and data science expertise to be fully utilized, possible future enhancements aim to further simplify the use of `m14xcube`, and increase its accessibility for domain experts. In summary, `m14xcube` contributes to the integration of ML with environmental data analysis, offering both technical solutions and enhanced accessibility, ultimately aiming to mitigate the impacts of climate extremes through enhanced data-driven decision-making.

## Acknowledgments

This work was conducted with the support of the European Space Agency as part of the DeepESDL project.

<sup>3</sup><https://deepesdl.readthedocs.io/en/latest/guide/ml-toolkits/>

## References

- Abadi, M.; Barham, P.; Chen, J.; Chen, Z.; Davis, A.; Dean, J.; Devin, M.; Ghemawat, S.; Irving, G.; Isard, M.; et al. 2016. {TensorFlow}: a system for {Large-Scale} machine learning. In *12th USENIX symposium on operating systems design and implementation (OSDI 16)*, 265–283.
- Babaeian, E.; Paheding, S.; Siddique, N.; Devabhaktuni, V. K.; and Tuller, M. 2021. Estimation of root zone soil moisture from ground and remotely sensed soil information with multisensor data fusion and automated machine learning. *Remote sensing of environment*, 260: 112434.
- Baumann, P.; Misev, D.; Merticariu, V.; and Huu, B. P. 2019. Datacubes: Towards space/time analysis-ready data. *Service-Oriented Mapping: Changing Paradigm in Map Production and Geoinformation Management*, 269–299.
- Dhu, T.; Giuliani, G.; Juárez, J.; Kavvada, A.; Killough, B.; Merodio, P.; Minchin, S.; and Ramage, S. 2019. National open data cubes and their contribution to country-level development policies and practices. *Data*, 4(4): 144.
- Flach, M.; Brenning, A.; Gans, F.; Reichstein, M.; Sippel, S.; and Mahecha, M. D. 2020. Vegetation modulates the impact of climate extremes on gross primary production. *Biogeosciences Discussions*, 2020: 1–20.
- Flach, M.; Gans, F.; Brenning, A.; Denzler, J.; Reichstein, M.; Rodner, E.; Bathiany, S.; Bodesheim, P.; Guaniche, Y.; Sippel, S.; et al. 2017. Multivariate anomaly detection for Earth observations: a comparison of algorithms and feature extraction techniques. *Earth System Dynamics*, 8(3): 677–696.
- Fowler, H.; Blenkinsop, S.; Green, A.; and Davis, P. 2024. Precipitation extremes in 2023. *nature reviews: earth & environment*, (5): 250–252.
- Grasso, F.; Talluri, G.; Giorgi, A.; Luchetta, A.; Paolucci, L.; et al. 2019. Peer-to-peer energy exchanges model to optimize the integration of renewable energy sources: The e-cube project. *L'Energia Elettrica*, 96: 0–0.
- Hall, K. J. C.; and Acharya, N. 2022. XCast: A python climate forecasting toolkit. *Frontiers in Climate*, 4: 953262.
- Han, Z.; Zhang, C.; Gao, L.; Zeng, Z.; Zhang, B.; and Atkinson, P. M. 2023. Spatio-temporal multi-level attention crop mapping method using time-series SAR imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 206: 293–310.
- Hansen, J. N. 2022. nd—A Framework for the Analysis of n-dimensional Earth Observation Data. *Journal of Open Research Software*, 10(1).
- Hoyer, S.; and Hamman, J. 2017. xarray: ND labeled arrays and datasets in Python. *Journal of Open Research Software*, 5(1): 10–10.
- Hsieh, W. W. 2022. Evolution of machine learning in environmental science—A perspective. *Environmental Data Science*, 1: e3.
- Imambi, S.; Prakash, K. B.; and Kanagachidambaresan, G. 2021. PyTorch. *Programming with TensorFlow: solution for edge computing applications*, 87–104.
- IPCC, . 2023. IPCC, 2023: Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team and Lee, H. and J. Romero (eds.)]. *IPCC, Geneva, Switzerland.*, 35–115.
- Islam, M. M. 2024. Unveiling the Power of Deep Learning: Insights into Advanced Neural n Network Architectures. *Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023*, 3(1): 1–14.
- Karmaker, S. K.; Hassan, M. M.; Smith, M. J.; Xu, L.; Zhai, C.; and Veeramachaneni, K. 2021. Automl to date and beyond: Challenges and opportunities. *ACM Computing Surveys (CSUR)*, 54(8): 1–36.
- Kerrigan, D.; Hullman, J.; and Bertini, E. 2021. A survey of domain knowledge elicitation in applied machine learning. *Multimodal Technologies and Interaction*, 5(12): 73.
- Kheir, A. M.; Govind, A.; Nangia, V.; Devkota, M.; El-nashar, A.; Omar, M. E. D.; and Feike, T. 2024. Developing automated machine learning approach for fast and robust crop yield prediction using a fusion of remote sensing, soil, and weather dataset. *Environmental Research Communications*, 6(4): 041005.
- Kim, S.; Yu, G.-I.; Park, H.; Cho, S.; Jeong, E.; Ha, H.; Lee, S.; Jeong, J. S.; and Chun, B.-G. 2019. Parallax: Sparsity-aware data parallel training of deep neural networks. In *Proceedings of the Fourteenth EuroSys Conference 2019*, 1–15.
- Leeuwenburg, T.; Loveday, N.; Ebert, E. E.; Cook, H.; Khanarmuei, M.; Taggart, R. J.; Ramanathan, N.; Carroll, M.; Chong, S.; Griffiths, A.; et al. 2024. scores: A Python package for verifying and evaluating models and predictions with xarray. *Journal of Open Source Software*, 9(99): 6889.
- Li, S.; Jaroszynski, S.; Pearse, S.; Orf, L.; and Clyne, J. 2019. Vapor: A visualization package tailored to analyze simulation data in earth system science. *Atmosphere*, 10(9): 488.
- Liu, Y. Y.; Dorigo, W. A.; Parinussa, R.; de Jeu, R. A.; Wagner, W.; McCabe, M. F.; Evans, J.; and Van Dijk, A. 2012. Trend-preserving blending of passive and active microwave soil moisture retrievals. *Remote sensing of environment*, 123: 280–297.
- Loaiza, D. M.; Kraemer, G.; Anghela, A.; Camacho, C. L. A.; Brandt, G.; Camps-Valls, G.; Cremer, F.; Flik, I.; Gans, F.; Habershon, S.; et al. 2023. Data Cubes for Earth System Research: Challenges Ahead.
- Mahecha, M.; Bastos, A.; Bohn, F.; Eisenhauer, N.; Feilhauer, H.; Hickler, T.; Kalesse-Los, H.; Migliavacca, M.; Otto, F.; Peng, J.; Sippel, S.; Tegen, I.; Weigelt, A.; et al. 2024. Biodiversity and Climate Extremes: Known Interactions and Research Gaps. *Earth's Future*, 12(6): 1–18.
- Mahecha, M. D.; Gans, F.; Brandt, G.; Christiansen, R.; Cornell, S. E.; Fomferra, N.; Kraemer, G.; Peters, J.; Bodesheim, P.; Camps-Valls, G.; et al. 2020. Earth system data cubes unravel global multivariate dynamics. *Earth System Dynamics*, 11(1): 201–234.
- Mahecha, M. D.; Gans, F.; Sippel, S.; Donges, J. F.; Kaminiski, T.; Metzger, S.; Migliavacca, M.; Papale, D.; Rammig,

- A.; and Zscheischler, J. 2017. Detecting impacts of extreme events with ecological in situ monitoring networks. *Biogeosciences*, 14(18): 4255–4277.
- Martinuzzi, F.; Mahecha, M. D.; Camps-Valls, G.; Montero, D.; Williams, T.; and Mora, K. 2023. Learning extreme vegetation response to climate forcing: A comparison of recurrent neural network architectures. *EGUsphere*, 2023: 1–32.
- Montero, D.; Aybar, C.; Ji, C.; Kraemer, G.; Söchting, M.; Teber, K.; and Mahecha, M. D. 2024. On-Demand Earth System Data Cubes. *arXiv preprint arXiv:2404.13105*.
- Nguyen, D. M. T.; Cortes, J. C.; Dunn, M. M.; and Shiklomanov, A. N. 2023. Impact of Chunk Size on Read Performance of Zarr Data in Cloud-based Object Stores. *Authorea Preprints*.
- Noh, B.; Son, J.; Park, H.; and Chang, S. 2017. In-depth analysis of energy efficiency related factors in commercial buildings using data cube and association rule mining. *Sustainability*, 9(11): 2119.
- Noy, I. 2016. A global comprehensive measure of the impact of natural hazards and disasters. *Global Policy*, 7(1): 56–65.
- Park, D.; Yu, J.; Park, J.-S.; and Kim, M.-S. 2013. NetCube: a comprehensive network traffic analysis model based on multidimensional OLAP data cube. *International Journal of Network Management*, 23(2): 101–118.
- Pasche, O. C.; Wider, J.; Zhang, Z.; Zscheischler, J.; and Engelke, S. 2024. Validating deep-learning weather forecast models on recent high-impact extreme events. *Artificial Intelligence for the Earth Systems*.
- Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. 2011. Scikit-learn: Machine learning in Python. *The Journal of Machine Learning Research*, 12: 2825–2830.
- Perkins, S.; Questiaux, J.; Finniss, S.; Tyler, R.; Blyth, S.; and Kuttel, M. M. 2014. Scalable desktop visualisation of very large radio astronomy data cubes. *New Astronomy*, 30: 1–7.
- Ravishankar, S.; and Battineni, G. 2022. A Survey on Recent Advancements in Auto-Machine Learning with a Focus on Feature Engineering. *Journal of Computational and Cognitive Engineering*.
- Reichstein, M.; Camps-Valls, G.; Stevens, B.; Jung, M.; Denzler, J.; Carvalhais, N.; and Prabhat, F. 2019. Deep learning and process understanding for data-driven Earth system science. *Nature*, 566(7743): 195–204.
- Rocklin, M.; et al. 2015. Dask: Parallel computation with blocked algorithms and task scheduling. In *SciPy*, 126–132.
- Rousi, E.; Fink, A. H.; Andersen, L. S.; Becker, F. N.; Beobide-Arsuaga, G.; Breil, M.; Cozzi, G.; Heinke, J.; Jach, L.; Niermann, D.; et al. 2023. The extremely hot and dry 2018 summer in central and northern Europe from a multifaceted weather and climate perspective. *Natural Hazards and Earth System Sciences*, 23(5): 1699–1718.
- Saha, S.; Gogoi, P.; Gayen, A.; and Paul, G. C. 2021. Constructing the machine learning techniques based spatial drought vulnerability index in Karnataka state of India. *Journal of Cleaner Production*, 314: 128073.
- Salehin, I.; Islam, M. S.; Saha, P.; Noman, S.; Tunj, A.; Hasan, M. M.; and Baten, M. A. 2024. AutoML: A systematic review on automated machine learning with neural architecture search. *Journal of Information and Intelligence*, 2(1): 52–81.
- Sarafanov, M.; Kazakov, E.; Nikitin, N. O.; and Kalyuzhnaya, A. V. 2020. A machine learning approach for remote sensing data gap-filling with open-source implementation: An example regarding land surface temperature, surface albedo and NDVI. *Remote Sensing*, 12(23): 3865.
- Sarkar, K.; Shiuly, A.; and Dhal, K. G. 2024. Revolutionizing concrete analysis: An in-depth survey of AI-powered insights with image-centric approaches on comprehensive quality control, advanced crack detection and concrete property exploration. *Construction and Building Materials*, 411: 134212.
- Schieler, C.; Robinson, B.; Guldner, O.; Bilyeu, B.; Garg, A.; Riesing, K.; Chang, J.; Hakimi, F.; Brown, J.; Khatri, F.; et al. 2019. NASA's terabyte infrared delivery (TBIRD) program: Large-volume data transfer from LEO.
- Schmale, S.; Seidel, P.; and Paul, S. 2017. Permuted cubes wavelet thresholding for mask-sensed MRI. In *2017 22nd International Conference on Digital Signal Processing (DSP)*, 1–5. IEEE.
- Shah, V.; and Kumar, A. 2019. The ML data prep zoo: Towards semi-automatic data preparation for ML. In *Proceedings of the 3rd International Workshop on Data Management for End-to-End Machine Learning*, 1–4.
- Shirvani, Z.; Abdi, O.; and Goodman, R. C. 2023. High-resolution semantic segmentation of woodland fires using residual attention UNet and time series of Sentinel-2. *Remote Sensing*, 15(5): 1342.
- Söchting, M.; Mahecha, M. D.; Montero, D.; and Scheuermann, G. 2023. Lexcube: Interactive visualization of large earth system data cubes. *IEEE Computer Graphics and Applications*.
- Sweet, L.-b.; Müller, C.; Anand, M.; and Zscheischler, J. 2023. Cross-validation strategy impacts the performance and interpretation of machine learning models. *Artificial Intelligence for the Earth Systems*, 2(4): e230026.
- Vincent, A. M.; and Jidesh, P. 2023. An improved hyperparameter optimization framework for AutoML systems using evolutionary algorithms. *Scientific Reports*, 13(1): 4737.
- Wasala, J.; Marselis, S.; Arp, L.; Hoos, H.; Longépé, N.; and Baratchi, M. 2024. AutoSR4EO: An AutoML Approach to Super-Resolution for Earth Observation Images. *Remote Sens.* 2024, 16, 443.
- White, R.; Anderson, S.; Booth, J.; Braich, G.; Draeger, C.; Fei, C.; Harley, C.; Henderson, S.; Jakob, M.; Lau, C.-A.; Admasu, L.; Narinesingh, V.; Rodell, C.; Roocroft, E.; Weinberger, K.; and West, G. 2021. The unprecedented Pacific Northwest heatwave of June 2021. *Nature Communications*, 12(14): 1–20.
- Wilkinson, M. D.; Dumontier, M.; Aalbersberg, I. J.; Appleton, G.; Axton, M.; Baak, A.; Blomberg, N.; Boiten, J.-W.; da Silva Santos, L. B.; Bourne, P. E.; et al. 2016. The FAIR

Guiding Principles for scientific data management and stewardship. *Scientific data*, 3(1): 1–9.

Winkler, A. J.; Myneni, R. B.; Hannart, A.; Sitch, S.; Haverd, V.; Lombardozzi, D.; Arora, V. K.; Pongratz, J.; Nabel, J. E.; Goll, D. S.; et al. 2021. Slowdown of the greening trend in natural vegetation with further rise in atmospheric CO<sub>2</sub>. *Biogeosciences*, 18(17): 4985–5010.

Yu, X.; Orth, R.; Reichstein, M.; Bahn, M.; Klosterhalfen, A.; Knohl, A.; Koebsch, F.; Migliavacca, M.; Mund, M.; Nelson, J. A.; Stocker, B. D.; Walther, S.; and Bastos, A. 2022. Contrasting drought legacy effects on gross primary productivity in a mixed versus pure beech forest. *Biogeosciences*, 19(17): 4315–4329.

Zhao, L.; and Ji, S. 2022. CNN, RNN, or ViT? An evaluation of different deep learning architectures for spatio-temporal representation of sentinel time series. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 16: 44–56.

Zumwald, M.; Baumberger, C.; Bresch, D. N.; and Knutti, R. 2021. Assessing the representational accuracy of data-driven models: The case of the effect of urban green infrastructure on temperature. *Environmental Modelling & Software*, 141: 105048.