



Deutscher Wetterdienst  
Wetter und Klima aus einer Hand



# E-AI Spring Workshop 2026: Products and Services 20 – 24 April 2026

## Agenda

Location [Hotel Zonar, Trg Krešimira Čosića 9, Zagreb, Croatia](#)

Times in CEST, Zagreb.

Stream and Material: <https://www.eumetnet-ai.eu/2026/workshop6/links/> (Password was sent on 15 April via email)

### Monday, 20 April:

#### Training Day

9:00 –	<i>Registration</i>
9:30 – 10:30	Getting Started with Pytorch on your device (Marek Jacob, DWD)
10:30 – 11:00	Coffee break
11:00 – 12:30	<a href="#">Low-cost sensors &amp; data fusion for PM2.5 - a visual transformer approach</a> (Marko Kvakić, DHMZ)
12:30 – 13:30	Lunch break
13:30 – 15:00	DAWID for Weather Services: Hands-on Introduction and Local Setup (Nikolas Porz, DWD)
15:00 – 15:30	Coffee break
15:30 – 17:00	Sustainable concepts for ML-software systems (Leif Denby, Danish Meteorological Institute)

### Tuesday, 21 April:

#### Opening

Chair: Marko Kvakić (DHMZ)

9:00 – 9:15	Welcome (Ivan Güttler, Director-General of DHMZ; Organizers)
9:15 – 9:35	E-AI Programme survey (Roland Potthast, E-AI)
9:35 – 10:20	<a href="#">Ethics and Regulatory Considerations for AI Developers and Operators</a> (Presentation and Discussion, Roland Potthast, Marek Jacob)
10:20 – 10:50	Coffee break

<b>Keynote and MLOps</b>	
Chair: Marko Kvakić (DHMZ)	
10:50 – 11:50	Keynote: <a href="#">Deep-learning prediction of high-frequency sea-level oscillations in the Adriatic Sea</a> (Iva Međugorac, University of Nova Gorica & ARSO)
11:50 – 12:15	<a href="#">From research to operations: deploying AI model for real-time regional weather forecasting</a> (Hugues de Laroussilhe, Meteo Swiss)
12:15 – 13:15	Lunch break
<b>LLMs and their Application for Meteorological Services</b>	
Chair: Steven Ramsdale (MetOffice)	
13:15 – 13:45	DAWID – LLM-based AI-Assistant (Nikolas Porz, DWD)
13:00 – 13:15	DAWID – Discussion
14:00 – 14:45	Building Services with LLMs for Meteorological Agencies – Discussion (Roland Potthast, DWD)
14:45 – 15:15	Coffee break
<b>AI in the Aviation Support Programme</b>	
Chair: Marek Jacob (DWD)	
15:15 – 16:05	Guidance for the Cross-Border Convection Forecast (CBCF) Product: A Collaborative Initiative between the Aviation Support Programme (ASP) and E-AI <ul style="list-style-type: none"> <li>❖ Introduction of the ASP and the CBCF Use Case (Mihaly Szucs, EUMETNET)</li> <li>❖ AI-Guidance for CBCF (Marek Jacob, DWD)</li> <li>❖ Lightning-based Guidance and Operational Concerns (Guido Schröder, Manuel Baumgartner, DWD)</li> <li>❖ Discussion</li> </ul>
16:05 – 16:25	<a href="#">AI from Aviation Forecaster's Perspective</a> (Stjepko Jančijev, Vinko Šoljan, Croatia Control)
16:25 – 17:15	<b>Poster Session</b> <ul style="list-style-type: none"> <li>❖ <a href="#">Streamlining Training Dataset Generation for ML Model Development at the Met Office</a> (Simon Boardman, Met Office)</li> <li>❖ <a href="#">Third Eye: A Platform for Environmental Data Integration and Machine Learning</a> (Marcel Prelogovic, Marshalling Ltd)</li> <li>❖ <a href="#">Machine Learning Applications in Post-Processing at DHMZ</a> (Jakov Lozuk, DHMZ)</li> <li>❖ <a href="#">Improving forecast quality through AI-based multi-model ensembles</a> (Grzegorz Zakrzewski, IMWM-NRI)</li> <li>❖ <a href="#">Using Crowdsourced Data to Improve User Communication</a> (Falk Anger, DWD)</li> <li>❖ <a href="#">obsweatherscale: a Python library for ML-based probabilistic interpolation and downscaling of surface weather fields</a> (Icía Lloréns Jover, Meteo Swiss)</li> <li>❖ GEO-Ring Level 1 (Carlos Horn, EUMETSAT)</li> </ul>
17:15 –	Icebreaker

Wednesday, 22 April

### Machine-learning based Nowcasting, Probabilistic Downscaling & Extreme Event Detection

Chair: Leif Denby (DMI)

- 9:00 – 9:30 [An update from WG6: the Open-Source MLCast Initiative for Machine Learning-Based Weather Nowcasting](#) (Lesley De Cruz, RMI)
- 9:30 – 9:35 [Machine learning based solar nowcasting system combining satellite observations and NWP data, poster introduction](#) (Soma Oláh, Pascal Gfaeller)
- 9:35 – 10:00 [Extreme precipitation detection in global climate models using ML](#) (Patrik Jureša, GFZ)
- 10:00 – 10:05 Group picture
- 10:05 – 10:30 Coffee break

### Data Curation

Chair: Roope Tervo (EUMETSAT)

- 10:30 – 10:55 [ECMWF's ARCO Data Access Strategy: Complementary Paths for Analysis-Ready Workflows](#) (James Hawkes, ECMWF)
- 10:55 – 11:20 [Advancing Environmental Prediction with the Canadian Surface Reanalysis \(CaSR\): New Developments, AI/ML Applications, and Future Directions](#) (Milena Dimitrijevic, ECCO)
- 11:20 – 11:35 Updates on UKMO Data Curation activities (Simon Boardman, MetOffice)
- 11:35 – 11:40 [Update on EUMETSAT ML-datasets including GeoRing and OPERA-SEVIRI-NORDLIS data](#) (Anna-Lena Erdmann, Lauren Biermann, Carlos Horn, Roope Tervo, EUMETSAT)
- 11:40 – 11:50 Early results from E-AI Data Curation Gap Analysis and updates (Lauren Biermann, EUMETSAT)
- 11:50 – 12:10 Interactive Data Gap Interviews (Lauren Biermann, EUMETSAT)
- 12:10 – 12:30 Discussion: Data Issues and Data Curation (Roope Tervo, EUMETSAT)
- 12:30 – 13:30 Lunch break

### Reusable Neural Network Architectures

Chair: Richard Williams (DWD)

- 13:30 – 14:15 [RACCOONN: A component-level framework for DOP exploration](#) (Joel Bedard, ECCO)
- 14:15 – 15:00 [MFAI: Toolbox library for meteorological machine learning](#) (Oscar Dewasmes, Meteo France)
- 15:00 – 15:30 Coffee break

### Machine Learning for Weather Prediction

Chair: Marko Kvakić (DHMZ)

- 15:30 – 16:00 Update from the ECMWF Machine-Learning Pilot Project (Katrin Ehlert, Meteo Swiss)
- 16:00 – 16:45 Relation of Physics and AI; Future of NWP and MLWP (Impulse discussion)
- afterwards Walking Dialogue (DHMZ)

Thursday, 23 April

**Flexible Framework for Applications in Meteorology**

Chair: Joel Bedard (ECCC)

9:00 – 9:20 *FRAIM* Architecture (Alexander Pasternack, DWD)

9:20 – 9:40 Downscaling with the DWD AI Toolbox (FRAIM) (Alexander Pasternack, DWD)

9:40 – 10:00 *SynCast* – Observation-oriented multi-model post-processing (Boran Frank, DWD)

10:00 – 10:20 [Prep2Floods – Indicators for Early Flood Warning](#) (Simon Zech, DWD)

10:20 – 10:50 Coffee break

**Frameworks**

Chair: Boran Frank (DWD)

10:50 – 11:50 [Earthkit. Open-Source Python Ecosystem Toward Operational-Ready Workflows](#) (Tiago Quintino, ECMWF)

11:50 – 12:30 Discussion: Towards Compatible and Synergetic Frameworks and Libraries (Boran Frank, DWD)

12:30 – 13:30 Lunch break

**Hands-on session: Zarr Hackathon**

Chair: Anna-Lena Erdmann, Carlos Horn (EUMETSAT)

13:30 – 15:00 Introduction to Zarr (Anna-Lena Erdmann, Carlos Horn, EUMETSAT; James Hawkes, Tiago Quintino, ECMWF)

15:00 – 15:30 Coffee break

15:30 – 17:30 Zarr Hackathon – Mastering Zarr together (Anna-Lena Erdmann, Carlos Horn, Roope Tervo, EUMETSAT; James Hawkes, Tiago Quintino, ECMWF)

Friday, 24 April

**Data Recovery**

Chair: Felicitas Hansen (DWD)

09:00 – 10:30 [From Paper to Proof: Revealing Congo Basin Warming Through Rescued Climate Archives](#) (Derrick Muheki, VUB)

10:30 – 11:00 Coffee break

**Planning session**

Chair: Roland Potthast (DWD)

11:00 – 11:30 Workshop Wrap-up

11:30 – 12:30 Planning the E-AI Service Portfolio

# Abstracts

## Monday, 20th of April

- Training day
  - [Low-cost sensors & data fusion for PM2.5 - a visual transformer approach](#)

## [Low-cost sensors & data fusion for PM<sub>2.5</sub> - a visual transformer approach](#)

*Marko Kvakić, Croatian Meteorological and Hydrological Service*

The recent tightening of air quality (AQ) standards in the EU has created a demand for regional AQ estimates at a much finer spatial resolution, as complementary modelling tools do not yet necessarily conform with the performance requirements set out by the new directive. To bridge this operational gap, low-cost sensors networks are increasingly looked upon, as more and more cities Europe-wide deploy this type of technology, and in turn provide much denser estimates of local air pollution levels; albeit at a higher uncertainty.

To address the topic of integrating low-cost sensor networks into existing regional AQ products, a dedicated working group was made within the FAIRMODE community (WG6), where best practices and experiences are exchanged, with the goal of providing guidelines and recommendations for operational use. Within the group, existing methodologies are evaluated and compared through a common benchmark experiment, in which hourly PM<sub>2.5</sub> low-cost sensor and reference data are provided for the year 2024 over the whole of Netherlands.

The workshop will go through a simplified example which focuses on the use of already existing frameworks (like PyTorch Lightning and mfai) in order to drastically speed up the process of development, in form of a JupyterLab notebook with a subset of the benchmark data which, highlights the key steps in a typical ML development workflow.

## Tuesday, 21st of April

- Opening
  - [Ethics and Regulatory Considerations for AI Developers and Operators](#)
- Keynote an MLOps
  - [Deep-learning prediction of high-frequency sea-level oscillations in the Adriatic Sea](#)
  - [From research to operations: deploying AI model for real-time regional weather forecasting](#)
- AI in the Aviation Support Programme
  - [AI from Aviation Forecaster's Perspective](#)
- Poster session
  - [Streamlining Training Dataset Generation for ML Model Development at the Met Office](#)
  - [Third Eye: A Platform for Environmental Data Integration and Machine Learning](#)
  - [Machine Learning Applications in Post-Processing at DHMZ](#)
  - [Improving forecast quality through AI-based multi-model ensembles](#)
  - [Using Crowdsourced Data to Improve User Communication](#)
  - [obsweatherscale: a Python library for ML-based probabilistic interpolation and downscaling of surface weather fields](#)

## Ethics and Regulatory Considerations for AI Developers and Operators

*Roland Potthast, E-AI*

*Marek Jacob, E-AI*

To support an interactive discussion, we would like to collect community input on current practices and open challenges in relation to ethical and regulatory concerns of AI-based systems. The exercise aims to identify which institutions or projects have already developed ethical documentation and guidelines for operational procedures.

For example, we would like to gather experience on approaches to handling hallucinations, unphysical artefacts, and other model failure modes, as well as on mechanisms for transparency, accountability, and responsible decision-making. Further topics include shared governance structures, intellectual property rights (IPR) of code and models, documentation standards, and the allocation of responsibility when AI-supported outputs inform scientific or operational decisions.

The collected input is intended to provide an overview of existing approaches, reveal gaps, and stimulate a focused discussion on good practices for trustworthy and responsible AI in our community.

## Deep-learning prediction of high-frequency sea-level oscillations in the Adriatic Sea

*Iva Međugorac, University of Nova Gorica & ARSO*

The eastern Adriatic coast is a known hotspot of strong meteorologically induced high-frequency sea-level oscillations (HFOs), occurring at periods shorter than 1 hour. The most extreme of these events, known as meteorological tsunamis (meteotsunamis), can reach heights of several meters and are generated by specific meteorological conditions that are spatially limited (from several tens to a few hundred kilometers). Although the physical mechanisms driving extreme HFOs are well understood, existing forecasting systems based on hydrodynamic models remain unreliable and computationally demanding.

To address these limitations, we develop deep-learning models to predict HFOs at two Adriatic locations, Bakar and Ploče, using convolutional neural networks driven by past sea-level observations and atmospheric predictors from ERA5 and CERRA reanalyses. These stations, located in the northern and central Adriatic, were selected because they provide the longest available 1-minute sea-level records (2003–2025) along the eastern (Croatian) Adriatic coast. Although they are not among the most meteotsunami-prone locations, their long records enable model training that can be transferred to more exposed sites with shorter datasets (e.g. Stari Grad, Vela Luka, Mali Lošinj, Sobra).

Our analysis shows that HFO extremes are larger in Bakar ( $A > 60$  cm) than in Ploče ( $A < 35$  cm), while occurring with similar frequency (~6 events per year) and predominantly during the warm season. Three types of extremes were identified based on the meteorological background (Bora, Sirocco, and resonance type), and were considered when splitting the time series into training, validation, and testing datasets. Two deep-learning architectures (SE-CNN and JE-CNN) successfully reproduce the observed variability, with higher skill for typical conditions than for extreme events. The JE-CNN performs best overall and under typical conditions, whereas SE-CNN more effectively captures extremes, although these remain systematically underestimated. Model performance is higher at Ploče, likely due to its smaller sea-level range and simpler response to forcing. ERA5 outperforms the higher-resolution CERRA, suggesting limited added value of increased spatial resolution. Ablation experiments further indicate that predictors improving average performance do not necessarily enhance the representation of extremes.

Overall, the results demonstrate the potential of deep learning for predicting high-frequency sea-level oscillations in the Adriatic, while highlighting persistent challenges in forecasting rare, high-amplitude events.

## [From research to operations: deploying AI model for real-time regional weather forecasting](#)

*Hugues de Laroussilhe, MeteoSwiss*

This presentation outlines the architecture of the real-time inference pipeline designed for the next generation of AI-based weather forecasting system at MeteoSwiss. As the meteorological community transitions toward data-driven weather prediction, a primary challenge lies in bridging the gap between flexible, experimental AI research and the rigid reliability required for safety-critical operations. We demonstrate a robust end-to-end workflow that successfully transitions AI research into a hardened High-Performance Computing operational environment.

The inference pipeline is built around a modular, multi-stage design where data preprocessing, model inference, and post-processing are decoupled into independent containerized components. This separation of concerns allows for optimized resource allocation, isolating GPU-intensive inference from data-heavy CPU tasks.

By leveraging MLflow as the central source of truth for model versioning and artifact management, we ensure strict lineage tracking from specific training hyper-parameters to real-time production forecasts. An automated CI/CD pipeline facilitates this transition, managing the building, testing, and publishing of these containerized stages to a registry, ensuring that every deployed version is verified, reproducible, and compliant with operational standards.

Finally, we describe the orchestration layer, where these containerized stages are deployed onto a dedicated HPC cluster at the Swiss National Supercomputing Center using the Slurm workload manager, controlled by an ecFlow scheduling suite. This setup mirrors the workflow of traditional Numerical Weather Prediction suites while introducing the agility of modern AI frameworks. Crucially, the system facilitates the delivery of inference products directly to forecasters' workstations. This establishes a vital human-in-the-loop feedback mechanism, where qualitative meteorological insights directly guide the next iteration of model training and refinement.

## AI from Aviation Forecaster's Perspective

*Stjepko Jančijev, Croatian Control*  
*Vinko Šoljan, Croatian Control*

Aviation meteorologists operate at the critical intersection of two complex disciplines: meteorology and aviation operations. Given that the majority of aviation accidents involve adverse weather as a contributing factor, forecast accuracy is paramount. This presentation examines how artificial intelligence can enhance two primary categories of aviation forecasting: point forecasts (Terminal Aerodrome Forecasts - TAF) and area forecasts (eGAFOR route forecasts).

We discuss current methodologies including Model Output Statistics (MOS) and manual polygon generation by experienced forecasters. While AI demonstrates significant potential in automating aspects of forecast generation and improving pattern recognition across multiple data sources, a fundamental challenge exists: aviation meteorology specifically requires the detection and forecasting of extreme phenomena, yet AI models characteristically smooth out extreme values.

This presentation argues that the optimal approach is "human-AI collaboration", where artificial intelligence provides intelligent automation and data synthesis support, while experienced forecasters maintain oversight, validation, and decision-making authority. We demonstrate how AI can generate initial "first guess" forecasts that accelerate the forecasting workflow while preserving the critical human expertise necessary for aviation safety.

## Streamlining Training Dataset Generation for ML Model Development at the Met Office

*Simon Boardman, Met Office*

At the Met Office we generate a huge amount of data from a multitude of sources on a daily basis. The ML Datasets team at the Met Office is responsible for building efficient workflows to ensure as much of this data as practical can be used in ML model development.

There are several common problems to face when converting legacy data to be ML Ready. Solutions will need to be pragmatic for today's compute ecosystem, yet forward thinking enough to avoid obsolescence in a rapidly evolving space.

- Data are difficult to search for and access, often requiring users to know what they are looking for rather than being able to browse what is available. Utilising Spatio Temporal Asset Catalogs (STAC) improves data discovery.
- Legacy datasets use a variety of naming conventions with different levels of completeness. A large amount of the data transformation work that needs to take place will be to apply a consistent metadata standard where possible and provide translation between other conventions.
- Data tend to be in file formats that are inefficient for ML Training due to data and metadata access. The use of appropriate Analysis Ready, Cloud Optimised formats will allow lazy loading of data in subsets alongside independent metadata querying.

To solve this non-exhaustive list of problems the ML Datasets team is undertaking the following work using international standards to ensure data interoperability:

1. Developing efficient extraction methods to retrieve legacy datasets from archival storage, databases, and across a diverse compute estate
2. Re-using operational code systems for standardisation where possible (primarily for NWP and Climate data)
3. Designing and developing conversion workflows for producing ARCO data
4. Producing STAC catalogues for ML Ready datasets

At present the focus is on making data previously produced, but ultimately the tooling and processes are being designed with application to future research and operational workflows to ensure consistency in our growing pool of ML Ready data for internal and external use.

## Third Eye: A Platform for Environmental Data Integration and Machine Learning

*Marcel Prelogovic, Marshalling Ltd*

With advances in data storage and processing technology, as well as the development of advanced machine learning algorithms, the use of data science to produce models in environmental science and monitoring has become the norm. However, data accessibility remains a major roadblock in this endeavor. Even when data exists out there, finding and processing it for use in highly interdisciplinary projects requires a significant amount of man-hours and a wide spectrum of expertise.

The goal of the Third Eye data collection and analysis platform is to help the users obtain and integrate high quality environmental data from various sources such as satellite images, seismological and geomagnetic data, drone footage, air and water data (temperature, visibility, pollution level...), long term meteorological data and short term forecasts, etc. We also provide a machine learning framework to facilitate testing of users' models on various datasets as well as building models from scratch. In the spirit of facilitating accessibility, the platform is fully online, with fast database access and an intuitive user interface, so even users that do not have a lot of experience in IT can still easily accomplish their goals. We also ensure the highest level of cybersecurity to make sure any proprietary data or models can be used with our platform without fear of it being accessed by other parties.

Though still in development, Third Eye shows potential as a data collection and machine learning tool for users that work with environmental data, whether they are a part of a government agency, a private company or a research institute.

## [Machine Learning Applications in Post-Processing at DHMZ](#)

*Jakov Lozuk, Croatian Meteorological and Hydrological Service*

In recent years, machine learning (ML) techniques have become widely used across a broad range of disciplines, helping to solve various complex problems. Meteorology is no exception, as ML methods can be applied to various tasks that involve a large amount of data. In this study, we present ongoing work on how ML techniques are implemented at DHMZ for post-processing of numerical weather prediction (NWP). We put the focus on key meteorological variables: temperature, wind speed, precipitation, and visibility. These variables are crucial not only for the general public, but also for specialized end users who rely on accurate forecasts to respond timely to hazardous extreme events that may cause significant material damage.

Although currently operational analog method approach provides notable progress in reducing errors for the most frequent wind speed or temperature categories, we evaluated the performance of several ML approaches in comparison to the analog method. Preliminary results indicate how ML-based post-processing can further enhance forecast performance, particularly for the most common categories, similarly to the analog method, and in some cases even outperform raw NWP for extreme categories. Additionally, we have also focused on developing new products, such as probabilistic forecasts of precipitation occurrence and visibility, where we exhibit promising preliminary results.

With continued refinement and operational integration, these advances can further improve forecast performance and help end users make more informed decisions when facing weather-related challenges.

## [Improving forecast quality through AI-based multi-model ensembles](#)

*Grzegorz Zakrzewski, Institute of Meteorology and Water Management*

In traditional Numerical Weather Prediction, combining different models into a multi-model ensemble is a proven strategy to achieve results superior to single-model ensembles. Despite the rapid rise of deep learning models – which offer remarkably low computational costs and ease of setup on modern GPUs – combining them into a multi-model ensemble has received little recognition.

To address this gap, we evaluated several ensemble strategies using popular neural network architectures and benchmarked them against ECMWF's operational AIFS-ENS. First, we tested an ensemble of deterministic models (GraphCast, FourCastNetV2, and PanguWeather) initialized with perturbed conditions using random or correlated noise. We also ran this ensemble within an assimilation cycle using a Localised Ensemble Transform Kalman Filter. While this setup yielded interesting spread-to-error properties, its overall forecast skill was significantly lower than AIFS-ENS.

Next, we explored ensemble models equipped with internal perturbation mechanisms, specifically FourCastNetV3 and GenCast. Individually, these achieve forecast skills only slightly below AIFS-ENS. However, when we combined AIFS-ENS, FourCastNetV3, and GenCast into a multi-model ensemble, the resulting system outperformed any of its constituent models running alone in terms of both forecast spread and skill.

Given the low computational cost and straightforward setup of ready-to-use deep learning models, multi-model ensembles represent a highly effective, accessible method for improving forecast quality.

## Using Crowdsourced Data to Improve User Communication

*Falk Anger, Deutscher Wetterdienst*

Crowdsourced data provides valuable information that complements numerical weather prediction (NWP) model output. It not only reflects how users perceive specific weather events but also captures impact-oriented insights. This makes such data particularly relevant for modern weather warning systems, as e.g. the RainBoW programme developed by the DWD (German Meteorological Service), which aim to ensure that both professionals and the general public clearly understand warnings and can take appropriate action.

In this study, we analyse crowdsourced observations collected via the WarnwetterApp, the weather app of the DWD. We combine these user reports with ensemble NWP data (ICON-D2-EPS) and explore the potential of this integrated dataset using XGBoost. Our focus is on its ability to enhance user communication by generating virtual crowd-sourced reports that reflect perceived weather impacts.

## [obsweatherscale: a Python library for ML-based probabilistic interpolation and downscaling of surface weather fields](#)

*Icía Lloréns Jover, Meteo Swiss*

*obsweatherscale* is an open-source Python library for machine learning-based probabilistic interpolation and regression of surface weather variables using Gaussian Processes (GPs). Built on GPyTorch, the library provides a modular and extensible framework for constructing neural-augmented GP models that incorporate trainable mean and kernel functions that accept arbitrary input features. Key features include plug-and-play data transformations, support for uncertainty quantification, GPU acceleration via PyTorch, and training and inference routines.

The library was developed to generate high-resolution wind maps over Switzerland, where complex alpine terrain and sparse observations challenge traditional methods. Using *obsweatherscale*, we downscale hourly ICON reanalysis wind fields from 1 km to sub-kilometer resolution integrating a wide variety of predictors, showcasing the model's ability to improve spatial detail and observational consistency, as well as provide calibrated uncertainty estimates. Beyond wind downscaling, *obsweatherscale* generalizes to a range of meteorological applications, including bias correction of model outputs and probabilistic spatial interpolation of observational datasets.

## Wednesday, 22nd of April

- Machine-learning based Nowcasting, Probabilistic Downscaling & Extreme Event Detection
  - [An update from WG6: the Open-Source MLCast Initiative for Machine Learning-Based Weather Nowcasting](#)
  - [Machine learning based solar nowcasting system combining satellite observations and NWP data](#)
  - [Extreme precipitation detection in global climate models using ML](#)
- Data curation
  - [ECMWF's ARCO Data Access Strategy: Complementary Paths for Analysis-Ready Workflows](#)
  - [Advancing Environmental Prediction with the Canadian Surface Reanalysis \(CaSR\): New Developments, AI/ML Applications, and Future Directions](#)
  - [Update on EUMETSAT ML-datasets including GeoRing and OPERA-SEVIRI-NORDLIS data](#)
- Reusable Neural Network Architectures
  - [RACCOONN: A Neural Field Framework for Satellite Data Assimilation and Beyond](#)
  - [MFAI: Toolbox library for meteorological machine learning](#)

## [An update from WG6: the Open-Source MLCast Initiative for Machine Learning-Based Weather Nowcasting](#)

*Lesley De Cruz, Royal Meteorological Institute of Belgium*

Within E-AI Working Group 6 (Nowcasting), strong collaboration between several European weather services and research institutes has led to the creation of the MLCast community and software framework. This initiative aims to make machine-learning nowcasting reproducible, validated, and easier to deploy operationally. The first release focuses on radar-based precipitation nowcasting and includes three main components: (1) ML-ready radar datasets from several European networks (Denmark, Italy, Germany, UK, Netherlands, Belgium) at 1 km resolution and 5-10 minute time steps, provided in CF-compliant Zarr format on the European Weather Cloud; (2) reference implementations of cutting-edge ML nowcasting models, trained on the mlcast-datasets; ; (3) a modular API designed for easy operational integration, including compatibility with pySTEPS.

In the open spirit of E-AI, all datasets, trained models, and source code are released under permissive open-source licenses, enabling NMHSs to experiment with state-of-the-art ML nowcasting even without large GPU resources.

Looking ahead, MLCast is evolving toward a multi-modal nowcasting framework combining radar with other observations such as Meteosat satellite imagery and lightning data. The goal is to support hybrid nowcasting systems that can predict multiple variables, including precipitation, lightning, and solar irradiance, which are increasingly relevant for applications such as renewable energy forecasting. MLCast demonstrates how collaborative development within E-AI can accelerate the operational adoption of machine learning while maintaining transparency and reproducibility.

[Machine learning based solar nowcasting system combining satellite observations and NWP data](#)  
[- poster preview](#)

*Soma Oláh, HungaroMet*  
*Pascal Gfaeller, HungaroMet*

Solar energy has become a key component of Hungary's electricity system, with total installed photovoltaic capacity now exceeding 8 GW and expected to continue growing in the coming years. To maintain grid stability and optimize energy trading accurate estimation of near-future energy production is required. This presents a significant challenge for weather-dependent renewable energy sources, such as photovoltaic (PV) systems, where power output is strongly influenced by changing meteorological conditions.

To support these efforts a machine learning based nowcasting system providing solar irradiance forecasts with a 3-hour lead time and 15-minute temporal resolution on a 0.05° grid is being developed for Hungary. The system integrates a near-real-time processing chain with a deep learning model centered on a ConvLSTM encoder-decoder architecture adapted from IrradianceNet (Nielsen et al., 2021). The primary dynamic input is the EUMETSAT LSA-SAF MDSSFTD satellite data product, complemented by static metadata such as latitude, longitude, elevation, day-of-year, and hour-of-day.

Ongoing research focuses on the integration of AROME/HU regional NWP fields as auxiliary predictors. This approach is driven by an assumed benefit of further supporting features, enabling a more accurate estimation of future atmospheric conditions that influence solar irradiance at the ground level. Preliminary results suggest improvements in forecast performance when combining satellite and NWP features to nowcast solar irradiance.

## Extreme precipitation detection in global climate models using ML

*Patrik Jureša<sup>1</sup>, Filip Kolarić<sup>1</sup>, Danijel Belušić<sup>1,2</sup>, Ramón Fuentes-Franco<sup>2</sup>, Mikhail Ivanov<sup>2</sup>*

*<sup>1</sup>University of Zagreb, Faculty of Science, Department of Geophysics, Zagreb, Croatia*

*<sup>2</sup>Rosby Centre, Swedish Meteorological and Hydrological Institute*

Extreme precipitation is increasing in intensity and frequency due to climate change and the impacts are considerable across multiple domains – societal, economic, and agricultural. To adapt and mitigate, it is important to identify, investigate and ultimately predict these events.

To capture the short-term and localised nature of extreme precipitation events, one possible approach is dynamical downscaling, which implies running a numerical climate model of finer resolution. Using a given high-resolution climate model, it explicitly simulates physical processes, rather than assuming stationarity. The drawback of such an approach is the computational cost of continuous, long-lasting simulations. This work addresses that limitation and solves it by running high-resolution model just for periods identified as extreme events in a global climate model (GCM).

The methodology proposed here is for extreme precipitation events identification from coarse-resolution GCM EC-Earth2's output and subsequently downscaling it using a regional climate model (RCM) HCLIM ALADIN. Event detection is performed by a Convolutional Neural Network (CNN), trained via supervised learning on daily GCM fields as predictors and daily extreme precipitation masks from the RCM as labels. Model skill is assessed using the Fractions Skill Score across seasons and predictor configurations.

Preliminary results over northern Europe demonstrate promising capabilities for ML-based detection, particularly during the cold season when extreme events are more under the influence of synoptic-scale processes. The proposed ML detection model increases the ability to detect and potentially predict such events, which could lead to better climate adaptation and mitigation strategies.

## [ECMWF's ARCO Data Access Strategy: Complementary Paths for Analysis-Ready Workflows](#)

*James Hawkes, ECMWF*

ECMWF delivers Analysis-Ready Cloud-Optimised (ARCO) data access through complementary paths: full-field access, curated Zarr datasets, Polytope-based feature extraction on native data, and Virtual Zarr. Rather than enforcing a single access pattern, this multi-path strategy balances flexibility, performance, and cost — supporting use cases from low-latency operational workflows to large-scale analytics and AI training pipelines across heterogeneous infrastructures.

A core design principle is preserving the operational reliability of primary storage while enabling analysis-ready workflows where they bring clear value. Curated Zarr products serve high-impact use cases selectively, while Polytope and Virtual Zarr reduce unnecessary data movement by exposing efficient, chunk-aware access over native archives — allowing users to retrieve targeted subsets without duplicating entire datasets. A common standards layer, including STAC and open interfaces, unifies discovery, navigation, and access across services, making the ecosystem more consistent and interoperable for both internal and external users.

This talk will present ECMWF's ARCO strategy, illustrate how these complementary approaches serve different segments of the EUMETNET community, and discuss opportunities for alignment within the E-AI initiative."

## Advancing Environmental Prediction with the Canadian Surface Reanalysis (CaSR): New Developments, AI/ML Applications, and Future Directions

*Milena Dimitrijević, Environment and Climate Change Canada*

This presentation will showcase the latest developments and future plans for the Canadian Surface Reanalysis (CaSR), a state-of-the-art, 10-km reanalysis system that delivers comprehensive meteorological and surface variables for North America. CaSR is distinguished by its enhanced spatial and temporal resolution, extended variable coverage, making it an indispensable resource for scientific research in environmental prediction and climate studies. The increasing availability of training datasets like CaSR is driving advancements in AI-driven meteorological, environmental, and hydrological prediction applications. We will discuss how updated versions of CaSR, including recent improvements in assimilation and data quality, extend their value for a wide range of AI/ML applications. These developments support not only weather and climate research but are also highly beneficial to hydrology, agriculture, and environmental prediction.

A key highlight of this presentation will be a case study demonstrating the successful application of CaSR within an AI/ML applied to hydrology framework. More precisely, a LSTM was trained over the Great-Lakes region using CaSR forcings, leading to strong improvements with regard to streamflow prediction, accuracy over traditional methods. By leveraging CaSR's grided data, significant gains in predictive skill for streamflow are achieved in Great Lakes watershed. This emphasizes the potential of such datasets to advance hydrological forecasting and water resource management.

Looking ahead, we will outline the vision and plans for the next generation of CaSR, with targeted enhancements in assimilation of new data sources, improved representation and details of terrestrial environments (lakes, towns, soils) and that will further increase the value of CaSR as a foundational dataset for the growing AI/ML community. Collaborations and user feedback will play a central role in shaping these future developments, ensuring that CaSR continues to meet evolving scientific and operational requirements in weather, climate, and hydrological prediction.

## [Update on EUMETSAT ML-datasets including GeoRing and OPERA-SEVIRI-NORDLIS data](#)

*Anna-Lena Erdmann<sup>1</sup>, Lauren Biermann<sup>1</sup>, Carlos Horn<sup>1</sup>, Roope Tervo<sup>1</sup>*  
*<sup>1</sup>EUMETSAT*

Machine learning (ML) applications bring a specific set of technical requirements towards Earth observation data. There is a growing need for data that is not only accessible but also optimised for large-scale, cloud-based processing. This presentation shares an overview of newly curated datasets in Zarr, as well as learnings from the conversion process.

EUMETSAT created a fused dataset combining surface radar (OPERA) composites with geostationary satellite radiances from SEVIRI rapid-scan data, and lightning data from the NORDLIS ground lightning network in Northern Europe to an aligned OPERA radar data grid. The dataset covers the period 2008–2013. Recent efforts have focused on converting the Opera-SEVIRI-Nordli dataset from NetCDF to Zarr. This required addressing challenges related to large-scale conversion, including chunking strategies and ensuring dataset consistency. Additional Zarr-datasets are currently being considered to support scalable analysis and machine learning workflows.

This presentation introduces the updates on the OPERA-SEVIRI-NORDLIS dataset, access mechanisms, and supporting resources, and shares lessons learned during the conversion process.

## [RACCOONN: A component-level framework for DOP exploration](#)

*Joel Bedard, Benoit Tremblay, Mark Buehner, Kuo-Hsien Chan, and Graigory Sutherland,  
Environment and Climate Change Canada*

Satellite radiances contain extensive information that can improve atmospheric analyses and numerical weather prediction (NWP). However, only part of this information is currently exploited due to assumptions of spatially uncorrelated errors, thinning strategies, and the exclusion of many channels—especially in cloudy conditions or over land and sea ice. To address these limitations, we are developing RACCOONN (Retrievals of Atmospheric Conditions Computed using Observations and Optimized Neural Networks), a deep learning framework designed to investigate more flexible, data driven approaches to atmospheric retrievals.

RACCOONN leverages coordinates based neural networks (neural fields) to represent thermodynamic profiles as continuous functions of space and time. Instead of using radiances as direct inputs, the neural field predicts atmospheric states (e.g., temperature, humidity, particle properties), and radiances appear in a variational style loss function. An observation term compares measured brightness temperatures with those simulated via a differentiable radiative transfer model, while an additional term regularizes the neural field predictions toward background profiles. Although the objectives are similar to those of classical 1D Var, the neural field formulation may relax on previously fixed assumptions. Initial experiments using synthetic datasets are underway, and upcoming evaluations will compare neural field retrievals with 1D Var to assess potential benefits and limitations.

Beyond retrievals, ECCC aim to examine whether neural fields could help blend outputs from multiple numerical models, potentially reducing discontinuities at model boundaries. Likewise, ECCC is exploring — at the conceptual level — whether the RACCOONN framework could support additional AI enabled applications. The continuous nature of neural fields suggests the possibility of extending the approach to rapid update or nowcasting contexts, where smooth spatio temporal evolution is advantageous. These directions remain exploratory, and part of our ongoing research is to rigorously test the suitability and feasibility of neural field architectures for such tasks.

## [MFAI: Toolbox library for meteorological machine learning](#)

*Oscar Dewasmes, Météo-France*

Over the years, Météo-France's AILab has gathered reusable code for meteorological machine learning applications in its MFAI library. It contains a variety of PyTorch neural network architectures (CNNs, Vision Transformers, small LLMs, small multimodal LMs, etc.), a NamedTensor class, losses, metrics, and PyTorch Lightning training strategies.

All the elements of the mfai library are implemented from research papers and have been improved, tested, and proven to work in real-world operational meteorological applications at Météo-France.

MFAI is not a framework, but it provides elements to use in your preferred one. In this presentation, we will go over:

- MFAI's content.
- How easily it can be integrated in your pure PyTorch projects.
- How a real-life operational PyTorch Lightning project uses MFAI.
- How to contribute.

## Thursday, 23rd of April

- Flexible Framework for Applications in Meteorology
  - [Prep2Floods - Indicators for Early Flood Warning](#)
- Frameworks
  - [Earthkit: Open-Source Python Ecosystem Toward Operational-Ready Workflows](#)

## Prep2Floods - Indicators for Early Flood Warning

*Simon Zech, Deutscher Wetterdienst*

Timely, spatially detailed flood signals are essential for anticipating impacts from heavy rainfall events. We present an AI based indicator for early flood warning implemented as an application in FRAIM, DWD's interface framework for AI applications in meteorology. It is demonstrated in the DestinE Flood Signals use case to showcase the use of the DestinE Extremes Digital Twin for weather extremes.

The method combines time resolved precipitation with static terrain information (orography) and applies a U-Net that has been trained using physically motivated prior model runoff/indicator fields (instantaneous values, short term maxima, and assimilated variants). The resulting near real time, kilometre scale flood signal maps provide a fast screening product that complements established hydrological forecasts and gauge based information.

In the use case, we demonstrate how the indicator can act as a lightweight trigger to prioritise subsequent high resolution simulations and downstream decision support.

## [Earthkit: Open-Source Python Ecosystem Toward Operational-Ready Workflows](#)

*Tiago Quintino, ECMWF*

Earthkit is an open-source Python ecosystem that simplifies how meteorological and climate data are accessed, processed, analysed, and visualised. It provides interoperable, high-level components — spanning data access, geospatial and meteorological computation, regridding, visualisation, and domain applications such as hydrology — while embracing GPU acceleration, cloud-native data access, and in-memory processing. Users work consistently across formats (GRIB, NetCDF, GeoTIFF, CoverageJSON), storage backends, and scientific tasks from operational pipelines to AI workflows.

For EUMETNET, Earthkit offers interoperable, scalable data workflows adoptable across ECMWF's Member and Co-operating States, aligning with E-AI goals for joint package development and shared operational frameworks.

This talk will introduce Earthkit's core capabilities, demonstrate its relevance for national meteorological services, and outline integration with initiatives including Forecast-in-a-Box, Copernicus Data Store modernisation, and Destination Earth. We will highlight Earthkit's evolution toward a 1.0 operational-ready release and discuss opportunities for collaboration within EUMETNET.

## Friday, 24th of April

- Data recovery
  - [From Paper to Proof: Revealing Congo Basin Warming Through Rescued Climate Archives](#)

## [From Paper to Proof: Revealing Congo Basin Warming Through Rescued Climate Archives](#)

*Derrick Muheki, Vrije Universiteit Brussel*

The Congo Basin in Central Africa remains one of the few regions globally where the intergovernmental Panel on Climate Change (IPCC) has not reported observed trends in hot extremes or attributed such changes to anthropogenic influences, primarily due to the scarcity of in situ observational data. Similarly, observed changes in extreme daily precipitation since the 1950s have not been assessed for this region. Although extensive daily weather records exist, spanning from the 1900s to the early 2000s and covering numerous stations across the basin, the majority of these remain archived on paper, limiting their accessibility for climate analysis.

Here, we present our historical weather data rescue project entailing archived data from 37 weather stations in the Democratic Republic of the Congo (DRC). We outline the digitization process of these archival records, comprising over 1 million individual observations, and describe the subsequent transcription using MeteoSaver version 1.1. From these records, we construct daily time series of daily maximum, minimum, and average temperatures, precipitation, as well as dry bulb and wet bulb temperatures measured at three times per day (06:00, 15:00, and 18:00) across these stations. This newly transcribed dataset provides a critical foundation for undertaking hydroclimatic trend analysis in the Congo Basin, one of the world's most data scarce yet climatically significant regions.

Using this data, we conduct an analysis of multi-decadal temperature trends across the basin. Our findings reveal a consistent and accelerating warming signal since the 1960s, characterized by a rightward shift in the distribution of daily maximum, minimum, and average temperatures with each successive decade. Median trends across the stations are 0.24°C, 0.09°C, and 0.18°C per decade for daily maximum, minimum, and average temperatures, respectively, corresponding to approximately 0.7°C, 0.3°C, and 0.6°C of warming over 30 years. We further find an increasing frequency of hot extremes and a decreasing frequency of cold events with each successive decade during the period 1960-1990, across the aggregated station data. Specifically, the most recent decade exhibits approximately twice as many hot days per year and about half as many cold days compared to the earliest decade.

Overall, this analysis of newly digitized historical weather data for the DRC highlights the urgent need to close the knowledge gap on climate trends in the Congo Basin.