

# Newsletter

No. 183 | Spring 2025

EarthCARE data begin to make  
an impact

---

Operational release of AIFS Single 1.0

---

Analysis of humidity in the  
stratosphere


---

ECMWF contributes to Swiss  
supercomputing project

---

Extreme precipitation in Valencia region

---



© Copyright 2025

European Centre for Medium-Range Weather Forecasts, Shinfield Park, Reading, RG2 9AX, UK

The content of this document, excluding images representing individuals, is available for use under a Creative Commons Attribution 4.0 International Public License. See the terms at <https://creativecommons.org/licenses/by/4.0/>. To request permission to use images representing individuals, please contact [pressoffice@ecmwf.int](mailto:pressoffice@ecmwf.int).

The information within this publication is given in good faith and considered to be true, but ECMWF accepts no liability for error or omission or for loss or damage arising from its use.

---

### **Publication policy**

The ECMWF Newsletter is published quarterly. Its purpose is to make users of ECMWF products, collaborators with ECMWF and the wider meteorological community aware of new developments at ECMWF and the use that can be made of ECMWF products. Most articles are prepared by staff at ECMWF, but articles are also welcome from people working elsewhere, especially those from Member States and Co-operating States.

The ECMWF Newsletter is not peer-reviewed.

Any queries about the content or distribution of the ECMWF Newsletter should be sent to [Georg.Lentze@ecmwf.int](mailto:Georg.Lentze@ecmwf.int)

Guidance about submitting an article and the option to subscribe to email alerts for new Newsletters are available at [www.ecmwf.int/en/about/media-centre/media-resources](http://www.ecmwf.int/en/about/media-centre/media-resources)

# EarthCARE and the AIFS Single



This has been a long time coming. The EarthCARE satellite, a joint venture between the European Space Agency (ESA) and the Japan Aerospace Exploration Agency (JAXA), was launched in May 2024. Its four instruments provide new insights into the interactions between clouds, aerosols, precipitation and radiation, as well as helping to establish the initial conditions of weather forecasts. EarthCARE is part of the Earth Explorer missions, which address key scientific challenges identified by the Earth observation community. For us, there are two areas which are important: the first is new data from which to infer the properties of clouds and how they interact with solar and thermal-infrared radiation. The results can be used to improve how cloud processes are represented in weather models. The second is the assimilation of the satellite's data into our Integrated Forecasting System (IFS) to improve the initial conditions of weather forecasts. This is exciting as it will be the first time that radar and lidar data are assimilated operationally in a global data assimilation system. An article in this Newsletter provides an overview of our use of EarthCARE data, and it reports on first positive results in data assimilation. Our forecasts are expected to benefit from the data as early as this year.

Another development highlighted in this Newsletter is the operational release of data-driven forecasting. Our Artificial Intelligence Forecasting System (AIFS) became operational on 25 February with the AIFS Single 1.0 version. It produces a single forecast, and it will be complemented later this year with an AIFS ensemble system. The AIFS still depends on traditional weather forecasting because it uses ECMWF's ERA5 reanalysis dataset as well as the IFS operational

analysis for training purposes and for initialisation. However, on that basis, starting from the operational analysis, it produces a forecast solely by means of machine-learning (ML) methods. And it does so very successfully because it achieves better results than traditional forecasts in a number of headline scores. The AIFS Single 1.0 is the first operational ML weather forecasting model. It has a greater number of forecast variables than other data-driven systems.

The growing influence of AI and ML methods has also resulted in ECMWF collaborating with the Swiss National Supercomputing Centre (CSCS) and MeteoSwiss to establish efficient access to our forecast products and archive to train AI and ML models. An article in this Newsletter describes how the project creates new opportunities for research in meteorology and scientific computing. There is also a look at the next upgrade of the IFS to Cycle 50r1 later this year, with an article on reintroducing an analysis of humidity in the stratosphere. The article demonstrates that this step will improve forecasts of humidity and temperature across all lead times, in particular in the Upper Troposphere Lower Stratosphere (UTLS) layer, but improvements can also be seen further down in the troposphere.

**Florence Rabier**  
Director-General

## Contents

### Editorial

EarthCARE and the AIFS Single ..... 1

### News

Extreme precipitation in Spain's Valencia region ..... 2  
Forecasts for Storm Éowyn ..... 4  
Operational release of AIFS Single 1.0 ..... 6  
Towards a FAIRer Data Stores service ..... 7  
The dawn of a new ERA Explorer ..... 9  
The current state of the climate in Europe ..... 10  
Support for regional re-forecasts over Montenegro ..... 12  
Operationalisation of global climate reanalysis led by ECMWF ..... 14  
Impact assessment of Chinese hyperspectral infrared sounder in preparation for MTG-S IRS ..... 15  
Evolution of the Regional Meteorological Data Communication Network ..... 16

AI Weather Quest: Advancing sub-seasonal forecasting with AI/ML ..... 18

Kilometre-scale modelling with the nonhydrostatic dynamical core ..... 19

### Earth system science

EarthCARE data begin to make an impact ..... 21

Reintroducing the analysis of humidity in the stratosphere ..... 28

### Computing

ECMWF contributes to Swiss supercomputing project in weather and climate science ..... 32

### General

ECMWF publications ..... 36

ECMWF Calendar 2025 ..... 36

Contact information ..... 36



# Extreme precipitation in Spain's Valencia region

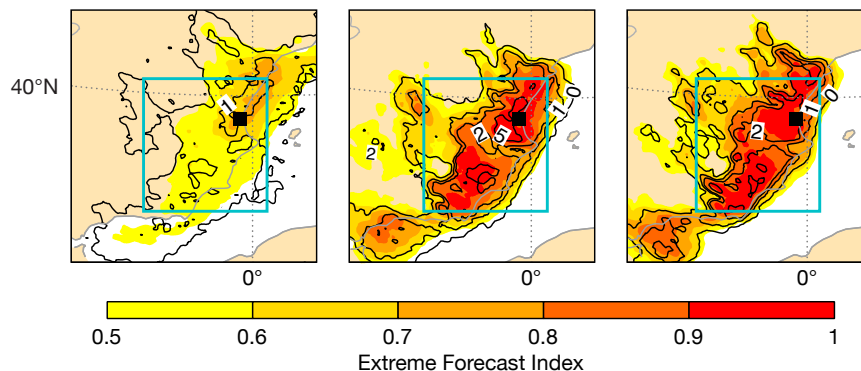
Estibaliz Gascón, Linus Magnusson, Tim Hewson (all ECMWF), Jaime Rey, Javier Rodríguez (both AEMET, Spain)

Between 28 October and 4 November 2024, a prolonged period of intense rainfall affected Spain's Mediterranean and adjacent provinces. On 29 October, the rainfall led to flash floods and river flooding, resulting in devastating impacts in the province of Valencia and, to a lesser extent, in other surrounding regions (Albacete, Cuenca and Málaga). At least 232 people lost their lives. Infrastructure damage and economic losses were severe, totalling around €16.5 billion.

## Extreme rainfall and flooding

The weather system associated with this event was a cut-off low-pressure system, sometimes referred to in Spanish as a 'dana' ('depresión aislada en niveles altos'). Warm, moist air from the Mediterranean Sea was pushed toward the Spanish coast, creating powerful convective structures. A detailed description of the event, including the storm structure, chronology, observed rainfall and historical context, is available in Spanish from the Spanish national meteorological service, AEMET (<https://tinyurl.com/5y7wtfty>). Further information is provided in ECMWF's severe event catalogue (<https://tinyurl.com/yfwxn488>). The storms persisted for several hours over the same area. As a result, during this event, the national records for total rainfall in one (184.6 mm), six (620.6 mm) and twelve hours (720.4 mm) were all broken. The 1-hour rainfall total of 184.6 mm is three times AEMET's threshold defining 'torrential rain', which is 60 mm. On 29 October, 771.8 mm of rain in 24 hours was seen at Turís Mas de Calabarra. This is the second highest 24-hour rainfall total on record for Spain, after 817.0 mm observed in Oliva (Valencia) during a similar event in 1987.

The most impacted areas were located immediately south of the city of Valencia, specifically the Horta Sud region (including Paiporta). Although they received relatively little direct precipitation, their location downstream and downhill of the aforementioned extreme totals was critical, with the



**EFI for precipitation.** Extreme Forecast Index (EFI) (colours) and Shift of Tails (black lines) in eastern Spain for 24-hour total precipitation valid on 29 October 2024 and based on forecasts starting at 00 UTC on 24 (left), 26 (middle) and 28 (right) October 2024. The box indicates the area shown in the next figure, and the black square shows the position of Valencia.

seasonal river Rambla del Poyo and its small catchment playing a key role in flooding. Many sections of the rivers Turia, Júcar, Cabriel and Magro, which have much larger catchments, also flooded as waves of water moved quickly downhill towards coastal towns.

## Extreme precipitation forecast

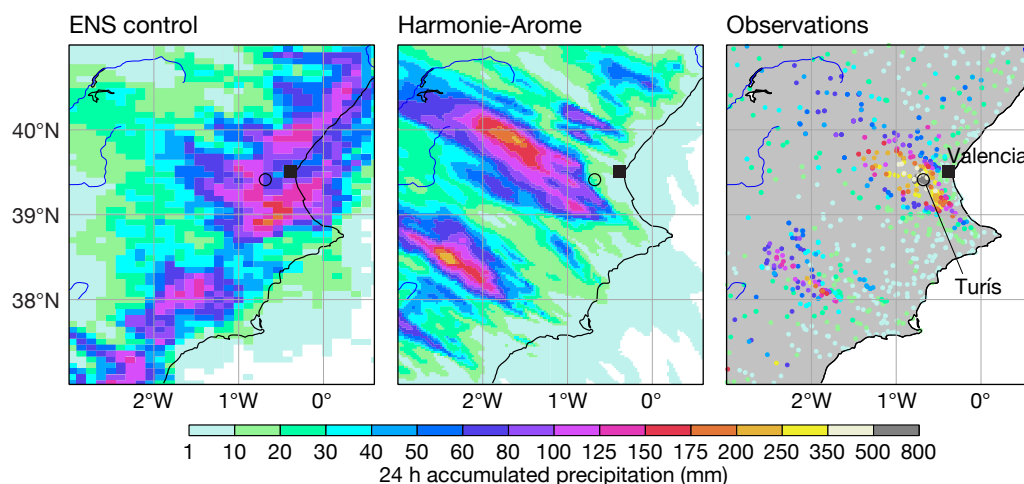
ECMWF's Extreme Forecast Index (EFI) for 24-hour precipitation indicated a strong signal for a rare precipitation event over the Mediterranean regions of Spain, appearing as early as six days in advance, with values ranging between 0.7 and 0.8. By four days prior to the event, EFI values neared 1, suggesting an increasing confidence in an extreme precipitation scenario. This high predictability was due to the accurate representation of the cut-off low location in the forecast several days in advance. While some uncertainty remained in the exact spatial distribution of the anomalous weather at lead times of 5–6 days, forecast consistency improved significantly from day 4 onward. From that time, the Integrated Forecasting System ensemble prediction system (IFS ENS) maintained a stable spatial representation of the expected precipitation patterns. Shift of Tails (SOT) values, which provide information about how extreme an event could potentially be, exceeded 5, particularly for Valencia province in the forecast

issued on 26 October, with a three-day lead time. By the forecast issued on 28 October (one day before the event), the high SOT values were widely spread across neighbouring regions.

A comparison between the 48 h IFS ensemble control forecast and the Harmonie-Arome regional model from AEMET valid for 24 h precipitation accumulations on 29 October reveals significant differences in both intensity and spatial distribution of the precipitation. The Harmonie-Arome model, with its higher horizontal grid spacing (2.5 km), predicted stronger precipitation intensities near the affected areas than the IFS ensemble control with a grid spacing of 9 km. The regional model simulated maximum 24-hour accumulations of 175–250 mm, while the global IFS forecast produced maximum values in the range of 150–200 mm. However, high-density observational datasets from AEMET reveal a notable underestimation of the predicted precipitation maxima in both models, with actual observations indicating clearly a sizeable region (about 15 km by 15 km) with accumulations between 500 and 800 mm, which neither forecast system successfully captured.

In terms of spatial distribution, the IFS ensemble control exhibited limited inland propagation of precipitation from the coastal areas along the prevailing southeast-to-northwest





**Predicted and observed precipitation.** 24 h accumulated precipitation valid from 29 October 00 UTC to 30 October 00 UTC 2024, in 24–48 h forecasts from 28 October 00 UTC, from ENS control (left) and the Harmonie-Arome operational forecast from AEMET (middle), with observations shown on the right. Sources of observations: AEMET, SAIH, SIAR, SUREMET, AVAMET, Sisritel and IVIA.

wind direction. In contrast, Harmonie-Arome reproduced this progression more realistically but positioned the convection initiation and peak precipitation too far inland.

## Understanding model errors

Whilst underprediction of extreme rainfall totals at grid scale has been seen before, high observation density makes this case unusually clearcut. Nonetheless, underprediction is atypical; on average the IFS overpredicts slightly. Therefore, different weather situations may have different forecast biases. ECMWF's ecPoint rainfall post-processing framework can provide further insights because, as well as anticipating sub-grid variability (blue lines on the cumulative distribution function figure panels), situation-dependent (grid scale) bias correction is also incorporated (green lines).

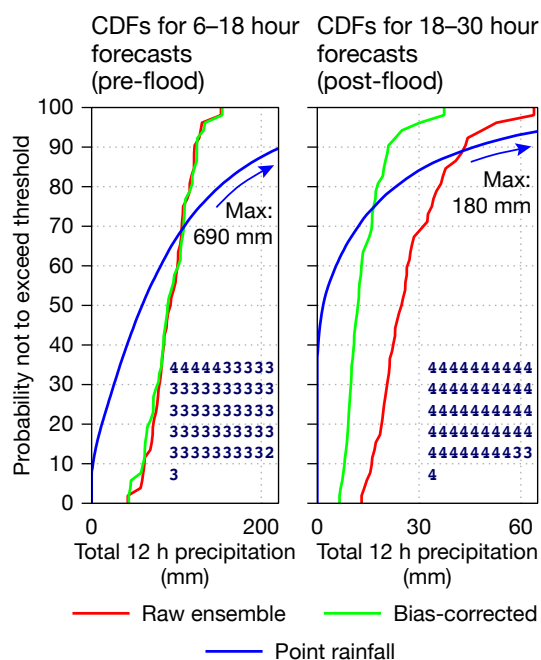
The relationship between raw IFS (red) and bias-corrected (green) differs between the two forecasting periods. On the left, raw output hardly changed, whilst on the right raw totals were reduced by about 50%. A physical interpretation can be made using 'weather types'. These are indices used for post-processing, that denote the ongoing meteorological situation in each member. On the left these mostly showed 50–75% of rainfall to be convective (3s in the table of numbers), but high Convective Available Potential Energy (not shown), which together suggest active convection that the IFS is trying to upscale. On the right, convective rainfall prevails (4s for >75% fraction) and mid-level winds are strong,

suggesting local over-prediction, in turn because parametrized convection does not move in the IFS. So, in the later period the chances of under-prediction in the raw IFS output were much smaller.

Meanwhile, using ecPoint calibration software offline, with millions of cases, we built on the above and identified another issue that could allow for further ecPoint and/or model physics refinements, to elevate rainfall total forecasts for this type of case. Most notably, biases also relate strongly to IFS near-surface relative humidity: high values, seen in this case, can link to underprediction, and vice-versa. Thus, from a model physics perspective droplet evaporation and warm-rain processes could be useful targets. The characteristics of Arome are expected to be rather different.

## Conclusion

The synoptic-scale predictability of the event was remarkably high, with signs of unusual and extreme weather well captured by the IFS ensemble forecast. However, current operational numerical weather prediction systems still have limitations and did not accurately predict either the maximum precipitation accumulation or its location, likely due to resolution constraints, simplifications in model physics, and challenges in representing convective-scale processes. Future improvements in high-resolution modelling, together with advances in model physics and suitable post-processing techniques will be essential to enhance the precision of extreme precipitation predictions in such high-impact events.



**Cumulative distribution function (CDF) plots of forecast rainfall.** Plots are based on IFS forecasts from 00 UTC on 29 October, for an ~18 km x 18 km grid box southwest of Valencia, for 12 h periods ending at 18 UTC on 29 October (left, mostly pre-flood) and at 06 UTC on 30 October (right, mostly post-flood): raw ENS (red), bias-corrected grid scale ecPoint equivalent (green), and point-scale ecPoint (blue). Digits (lower right) are part of the five-variable weather type index, for the 51 members, for the valid periods, showing convective rainfall fraction (2, 3, 4 for 25–50, 50–75 and >75% respectively). Blue curves have been chopped to save space; labels show the nominal 100% intersection point.

# Forecasts for Storm Éowyn

Linus Magnusson (ECMWF), Sinéad Duffy, Brandon Creagh (both Met Éireann)

Storm Éowyn hit the west coast of Ireland in the early hours of 24 January 2025. According to Met Éireann, the storm (based on preliminary evaluation) broke the land records for Ireland of 10-minute surface mean wind speed (142 km/h, 39.4 m/s) and wind gusts (184 km/h, 51.1 m/s) at the coastal station at Mace Head, County Galway. The storm caused widespread disruption in the Republic of Ireland and Northern Ireland and also affected Scotland later on. Ahead of the storm, Met Éireann issued red warnings for the whole country, and the UK Met Office did the same for Northern Ireland and parts of Scotland. 768,000 customers were without power in the Republic of Ireland, with nearly 326,000 customers affected in Northern Ireland. Estimates suggest that twice as many trees fell during the storm as there would normally be felled in a year. Here we present forecasts of the event provided by ECMWF's Integrated Forecasting System (IFS) and Artificial Intelligence Forecasting System (AIFS).

## Meteorological situation

The cyclone formed over the western Atlantic on 22 January. It quickly moved across the Atlantic on 23 January and rapidly intensified

under an extremely strong jet stream ( $>100$  m/s at 250 hPa in ECMWF's analysis). The lowest measured pressure (940 hPa) was on the station Belmullet at 03–04 UTC on 24 January, and the 06 UTC analysis from ECMWF had a minimum pressure off the coast of 941 hPa. Satellite images suggest that the most extreme wind was likely due to a sting jet feature southwest of the storm centre.

## ECMWF's forecasts

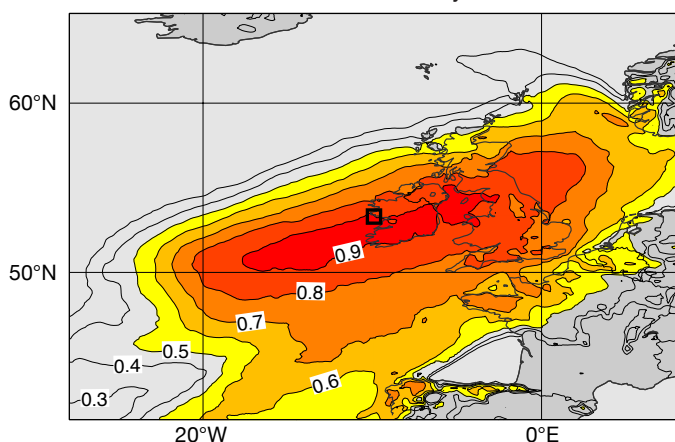
The signal for the event started to develop in ECMWF ensemble forecasts (IFS-ENS) from 18 January, about six days before the event. The examples of the Extreme Forecast Index (EFI) for 24-hour maximum wind gusts on 24 January included here are from 00 UTC on 19 and 21 January (see the figure). In the forecast from 19 January, the centre of the highest EFI was to the south of where the highest wind speeds were recorded (black square symbol), but the forecast still provided an early warning for most of the island. In the forecast from 00 UTC on 21 January, large parts of Ireland had an EFI of about 0.9. For Mace Head (black square), 50% of the ensemble members predicted maximum wind gusts higher than the maximum in the

model climate for this time of the year (not shown).

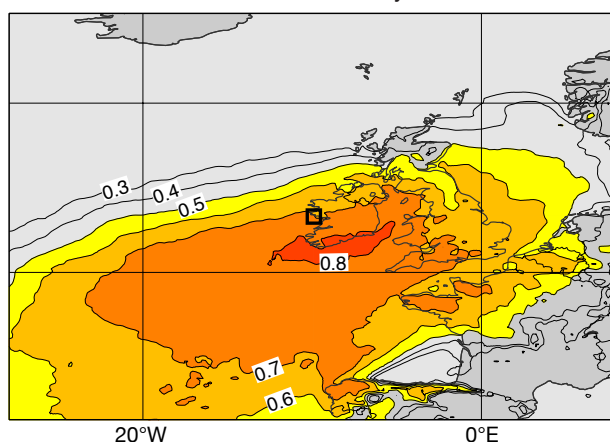
The plots in the second figure show the mean sea level pressure and the surface mean wind speed valid at 06 UTC on 24 January in the ECMWF analysis and in 30-hour forecasts (initialised at 00 UTC on 23 January). The forecasts are the IFS control forecast (IFS-CF) and the AIFS Single v1.0, which was pre-operational at the time. Both forecasts captured the position of the storm very well at this stage. For the minimum pressure, the AIFS predicted an even deeper cyclone (935 hPa) than the IFS control forecast (940 hPa). However, at the same time the maximum mean wind speed was considerably lower in the AIFS than in the IFS control forecast (30 m/s vs. 44 m/s). The analysis provided a wind speed of 38 m/s. As the maximum wind speed is believed to be in the narrow region of the sting jet, it is very difficult to verify the maximum wind speed over the sea as we only have a few buoy measurements, and none of the available buoys were located near the extreme.

To look further into the prediction of maximum mean wind at Mace Head, the third figure shows the ensemble

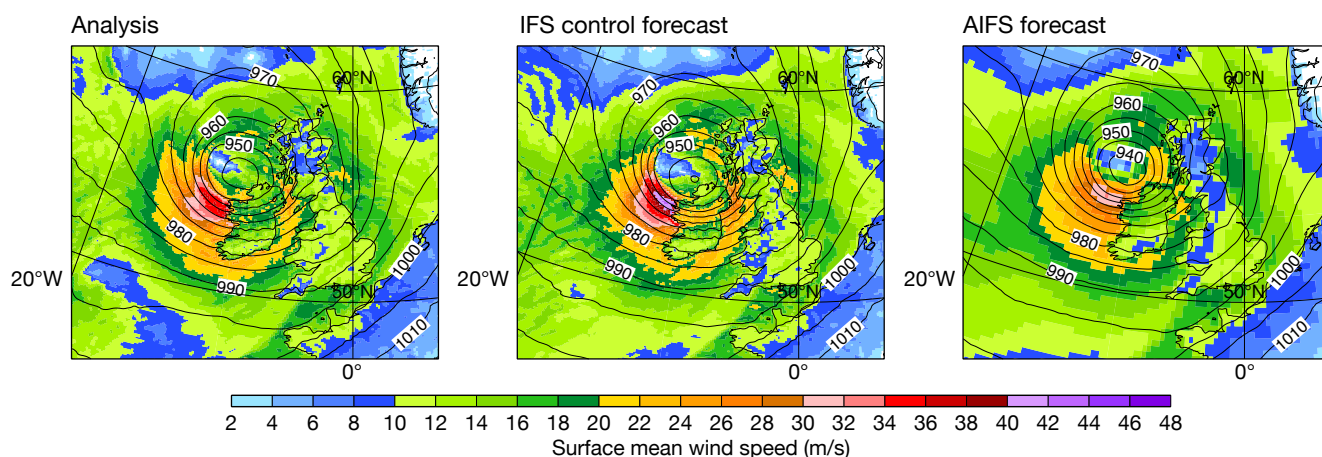
Forecast from 00 UTC on 21 January



Forecast from 00 UTC on 19 January



**IFS Extreme Forecast Index (EFI).** Forecasts of the EFI for 24-hour maximum wind gusts on 24 January are shown from 00 UTC on 21 January 2024 (left) and 00 UTC on 19 January 2024 (right), with Mace Head marked with a black square.



**Surface mean wind speed.** The panels show the mean sea level pressure (lines, in hPa) and the surface mean wind speed at 06 UTC on 24 January 2025 in ECMWF's analysis (left), a 30-hour forecast from the IFS control (middle), and a 30-hour forecast from the AIFS v1.0 (right).

forecast from 00 UTC on 23 January, using hourly output. The plot also includes hourly (not maximum) 10-minute mean wind from SYNOP weather station observations. The IFS-ENS control and many members provided a peak in the wind speed with similar magnitude and timing as the observations. However, a few ensemble members predicted very extreme mean wind speed, with the most extreme one reaching above 50 m/s. While it is not possible to verify if this was a potential scenario for a single case, the predicted value was more than 10 m/s above the observation, which is now a preliminary record for the Republic of Ireland.


The feedback from Met Éireann was

that the IFS performed in general well with the structure of the system, as well as performing well on a number of parameters, including mean winds for the storm, predicting hurricane force 12 mean wind speeds in the vicinity of Mace Head. Predicted gusts were also well forecast in terms of spatial distribution and peak amplitude gusts, while the rainfall was too low. Some ensemble members with extreme wind values would generally have been discounted as outliers. A good consistency in the runs with regard to track and depth of the predicted low led to good confidence in the forecast room, leading to high-level warnings issued relatively early on.

## Outlook

In summary, the ECMWF forecasts gave an early warning for the storm around 5–6 days in advance. The AIFS captured the pressure of the storm very well but underestimated the maximum wind speed. Future evaluation will also include the new AIFS ensemble that is under development. This result is in line with previously evaluated windstorms. For the IFS ensemble, we identified a few members with very extreme wind speed. After the implementation of IFS Cycle 49r1 in November 2024, we became aware of occasional ensemble members giving very high mean wind speeds during windy days. This triggered a deeper analysis of the ensemble statistics and further experimentation. It was identified that perturbations related to the surface momentum flux in the new Stochastically Perturbed Parametrizations (SPP) model uncertainty scheme played a major role in the members with extreme wind speed. While it is difficult to verify the reliability of forecasts for the most extreme winds, we are considering ways to address this issue in IFS Cycle 50r1 planned for implementation later this year (see ECMWF's 'Known IFS forecasting issues' page for more details:

<https://confluence.ecmwf.int/display/FCST/Known+IFS+forecasting+issues>).



**Surface mean wind speed at Mace Head.** Forecasts from 00 UTC on 23 January and observations for surface mean wind speed at Mace Head in Co. Galway, Republic of Ireland.



# Operational release of AIFS Single 1.0

Ewan Pinnington, Ana Prieto Nemesio, Gabriel Moldovan

In a significant milestone, ECMWF has implemented the first operational data-driven forecasting system, the AIFS (Artificial Intelligence Forecasting System), demonstrating the Centre's commitment to pursue the best possible weather forecasts with both physics-based and machine learning (ML) methods. There has been an explosion in scientific publications featuring ML-based forecasting systems over recent years. The vast majority of these systems depend on ECMWF's ERA5 reanalysis dataset and Integrated Forecasting System (IFS) operational analysis for training and initialisation of their forecasts. They have showcased how traditional physics-based numerical weather prediction (NWP) forecast models are outperformed for a number of headline scores at a much-reduced computational cost. Now this technology has been leveraged for operational meteorology with the implementation of the AIFS Single. This first version, AIFS Single 1.0, is named to highlight that it produces a single forecast, to be complemented later this year with a first AIFS ensemble system. Having implemented the AIFS Single, we can now begin to understand its performance at an event-based level with feedback from forecasters, as it runs operationally alongside the IFS.

## Basic characteristics




The AIFS Single has an encoder–processor–decoder structure, where the encoder first projects the model state onto a lower-resolution grid, using a graph neural network (GNN). The processor then uses a Transformer to update the model state 6-hours in time before the decoder reprojects the model state back to the target resolution. Further details on the model structure can be found in Lang et al., 2024a (<https://arxiv.org/abs/2406.01465>). The AIFS Single 1.0 was trained on ECMWF's ERA5 reanalysis dataset (1979–2022) and then fine-tuned to the IFS operational analysis (2016–2022).

The AIFS Single is now fully built under the open-source Anemoi framework developed by ECMWF and its Member States (see ECMWF Newsletter No. 181: <https://tinyurl.com/bdh5bc5a>). This ensures the development, training and implementation of such data-driven models are reproducible, traceable and scalable. More information on the Anemoi framework can be found within the set of Anemoi webinars (<https://events.ecmwf.int/event/446/>). The AIFS Single 1.0 has been released as a model on the ML model-sharing platform Hugging Face, including a guide on how to use this from ECMWF's open data stream (<https://huggingface.co/ecmwf/aifs-single-1.0>).

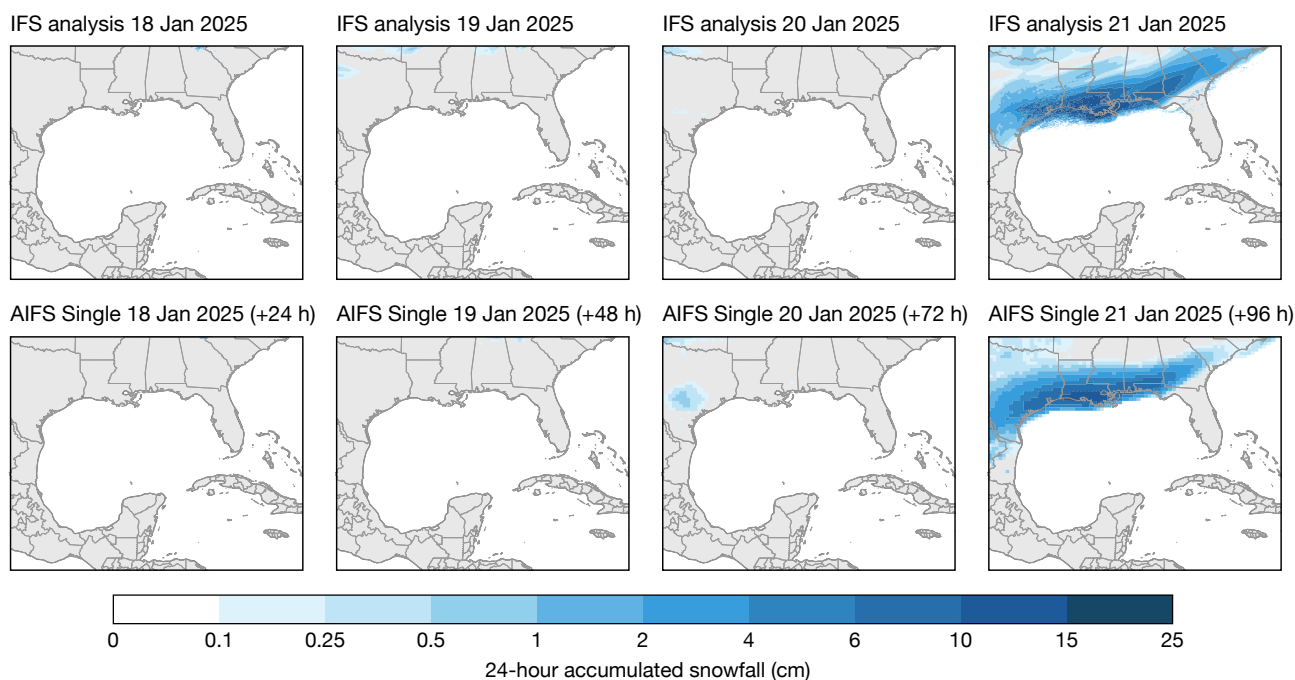
## New features

The AIFS Single has expanded the number of forecast variables compared with other data-driven systems. These include solar radiation and 100 m wind speeds, useful for energy forecasting and pricing. We now include cloud cover, snowfall and additional land surface variables of soil temperature, soil moisture and runoff. The physical consistency of these new parameters can be seen in the figure, where the model has an internally consistent representation of clouds for total cloud cover and solar radiation at the surface. Some of these new variables can already be explored through ECMWF charts (<https://charts.ecmwf.int/?query=aifs-single>).

We have increased the physical consistency of the model by including 'boundings' as an additional model layer after the decoder. We explicitly bound total precipitation, specific humidity, runoff and solar radiation by zero. Soil moisture and total cloud cover are bounded between zero and one. Snowfall, convective precipitation and low/medium/high cloud cover are bound to be fractions of total precipitation and total cloud cover respectively. This has led to increased forecast scores for precipitation and reduced previous issues, where it was noted the model produced an excess of light rainfall.



**New variables.** A selection of new variables is available from AIFS Single forecasts: cloud cover (left), surface solar radiation (centre), and 100 m wind speed/direction (right).



**Snowfall forecast over the Gulf Coast of America.** The figure shows how the snowfall event was forecast accurately four days ahead.

## Early performance

The AIFS Single is already showing strong performance while running operationally, including forecasting rare events that have no clear analogies in the model's training data. This can be seen clearly in the extreme snow event over the Gulf Coast of America on 21 January 2025 (see the second figure). This had a return period of roughly 100 years, i.e. such events are very rare or do not appear at all in the ERA5 reanalysis or the operational archive. These types of events provide confidence that the model is not only reproducing its training data but has learnt a representation of atmospheric physics. The ability of an earlier AIFS to capture extremes is also noted by Ben Bouallègue et al., 2024 (<https://doi.org/10.1175/BAMS-D-23-0162.1>).

[org/10.1175/BAMS-D-23-0162.1](https://doi.org/10.1175/BAMS-D-23-0162.1).

In general, the AIFS Single, due to training to minimise the mean squared error, still struggles to represent small-scale extreme events. First steps to address this will be made with the upcoming AIFS ensemble system.

## Outlook

The operational release of the AIFS Single has been the culmination of a lot of work across ECMWF. That includes work from the start of the process – by those who contribute to the creation of the ERA5 reanalysis and IFS operational analysis, vital to train and initialise the AIFS – to the end – by those who integrate AIFS models into real-time production and dissemination pipelines for forecasts.

New developments are already in progress for the AIFS. An ensemble version of the AIFS, trained using a Continuous Ranked Probability Score (Lang et al., 2024b: <https://arxiv.org/abs/2412.15832>), is targeted for operational release in the summer of 2025. Work is also under way to 'fine-tune' the ~28 km AIFS Single towards the IFS operational resolution of ~9 km, making full use of operational data, which is currently only used by the AIFS at ~28 km resolution. Within the EU's Destination Earth initiative and the Copernicus programme, to which ECMWF contributes, more Earth system components are being actively trained using ML methods, including waves, the ocean, sea ice, atmospheric composition, and hydrology.


# Towards a FAIRer Data Stores service

Marcus Zanicchi, Eduardo Damasio da Costa, Angel López Alós

The Copernicus Climate Data Store (CDS) and Atmosphere Data Store (ADS) provide a vast array of information from the EU's Copernicus Climate Change Service (C3S) and Atmosphere Monitoring Service (CAMS) run by ECMWF. The CDS can be found at [https://](https://cds.climate.copernicus.eu/)

[cds.climate.copernicus.eu/](https://cds.climate.copernicus.eu/), and the ADS at <https://ads.atmosphere.copernicus.eu/>. The two Data Stores have recently been developed with a strong emphasis on ensuring that our digital assets fully adhere to the FAIR principles – making them Findable, Accessible,

Interoperable, and Reusable (<https://www.go-fair.org/>). This transformation is not only a technical upgrade; it represents a strategic step toward maximising the impact of our data catalogues and serving a broader and more diverse user community. ECMWF is



	Score earned:	Fair level:
<b>Findable:</b>	7 of 7	<b>advanced</b>
<b>Accessible:</b>	3 of 3	<b>advanced</b>
<b>Interoperable:</b>	2 of 4	<b>initial</b>
<b>Reusable:</b>	5 of 10	<b>initial</b>

**FAIR assessment.** The modernised ECMWF data catalogues (CDS & ADS) have achieved an impressive FAIR assessment score of 70%, as measured by the FAIRsFAIR F-UJI assessment tool, as here shown for ERA5 hourly data on pressure levels from 1940 to the present. This signifies an important step towards maximising the impact of our Data Stores.

not only enhancing climate and atmospheric data accessibility for researchers and policymakers but also laying the groundwork to support the next generation of machine learning weather and climate applications and European-wide federated data platforms.

## Achieving a new benchmark in FAIR compliance

According to our recent FAIR assessment, our modernised data catalogues have achieved an impressive 70% FAIR score as measured by the FAIRsFAIR F-UJI assessment tool (see an example in the figure). The tool is a European Open Science Cloud initiative to assess the FAIRness of digital assets (<https://fairsfair.eu/fairsfair-eosc>). This score is a marked improvement in comparison to previous assessments, underlining the success of our efforts to implement modern data management strategies. The assessment highlights four findings:

- **Findability:** Each catalogue entry is now assigned a persistent Digital Object Identifier (DOI) registered with Datacite, ensuring that every dataset and its accompanying metadata remain uniquely locatable over time.
- **Accessibility:** The introduction of a STAC-compliant metadata API means that both human users and machines can now query catalogue metadata programmatically, significantly enhancing data discoverability and accessibility (<https://stacspeg.org/en>).
- **Interoperability:** By also adopting the Open Geospatial Consortium (OGC) API - Records, we are better

aligned with other European and global data-sharing infrastructures. This facilitates seamless integration and cross-referencing (<https://ogcapi.ogc.org/records/>).

- **Reusability:** Our plan to move towards using standard open licences (e.g. Creative Commons licences, CC BY) and the provision of richer, well-documented metadata now enable users to confidently reuse our data, generating new knowledge and increasing the impact value of the open data provided by the Data Stores.

## Standardisation: a catalyst for integration and innovation

Standardisation plays a pivotal role in the modernisation drive around the Data Stores. By aligning our data catalogues with internationally recognised standards, we ensure that our datasets are not only high quality but also readily consumable across a variety of platforms and applications. This is particularly important in an era where machine learning and artificial intelligence increasingly harvest large, standardised, datasets to drive innovation. As part of ECMWF's data services strategy, we have created a centralised, automated method for exposing metadata, enabling us to quickly update records with minimal intervention and onboard new metadata standards or adjust to evolving ones.

These improvements are instrumental in supporting an integrated data-sharing ecosystem that aligns with broader European Union initiatives and the European Open Science Cloud, highlighting ECMWF's role in supporting Europe's digital and green transitions.

## What next?

Whilst we are pleased with our progress so far, we recognise that this is a step on a longer journey toward optimal FAIR compliance. Future work will focus on:

- **Enriching metadata:** Further refining our metadata to provide detailed data provenance and using consistent controlled vocabulary.
- **Strengthening interoperability:** Enhancing links between our catalogues and external resources to boost overall data integration.
- **Improving reusability:** Transitioning more datasets to standard open licences like CC BY to maximise their reuse and increase their social and economic impact for the public.

## Explore our catalogues' FAIRness assessment

With fairer weather accompanying this spring edition of the ECMWF Newsletter, we invite you to assess our FAIRer catalogue. Simply visit the FAIRsFAIR F-UJI assessment page (<https://www.f-uji.net/>) and enter one of our catalogue DOIs (for example, <https://doi.org/10.24381/cds.bd0915c6> for ERA5 hourly data on pressure levels) to see our progress for yourself.

Are you interested in learning more? We hosted a webinar featuring external expert Dr Ge Peng from the University of Alabama. She presented her insights gained from assessing the FAIR compliance of NASA Earth science data products and also provided an overview of existing tools used for evaluating data FAIRness. You can re-listen to the excellent talk via our website (<https://events.ecmwf.int/event/471/>).



# The dawn of a new ERA Explorer

Matthew Menary, James Varndell, Edward Comyn-Platt, Chiara Cagnazzo

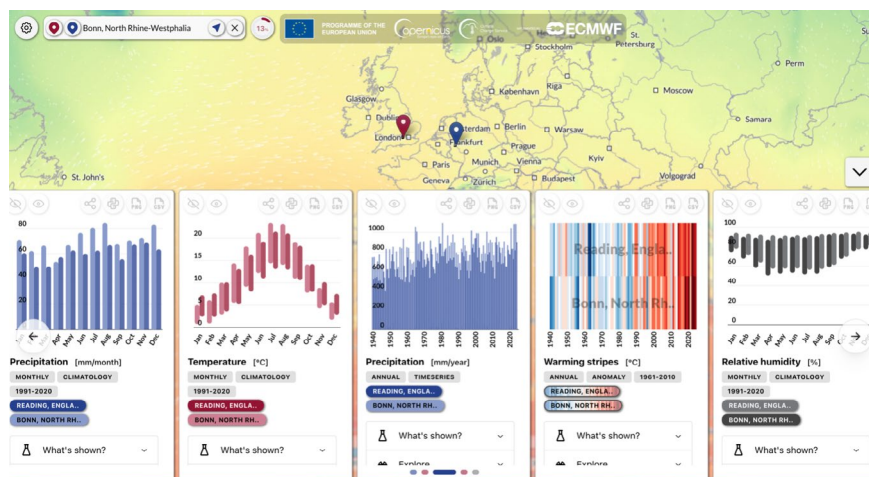
We are pleased to share the release of the new ERA Explorer web app powered by the Climate Data Store (CDS) of the Copernicus Climate Change Service (C3S). C3S is an EU service provided by ECMWF. ERA Explorer provides an interactive interface for exploring 85 years of global climate data from the ERA5 reanalysis: <https://era-explorer.climate.copernicus.eu/>. It replaces the previous ERA5 Explorer, which was introduced in 2019.

ERA Explorer is designed to demonstrate the power and richness of ERA5. It is deployed via the updated infrastructure in the C3S applications cluster. The app enables users to easily compare climate data across the globe, such as temperature, precipitation, wind and other climatological averages. For example, users can investigate local temperature thresholds or understand how the climate has changed in the last 85 years by generating warming stripes.

ERA Explorer has been designed from the ground up to be mobile-first and responsive, offering an optimal experience across devices, from smartphones to desktops. It offers a modern search and map navigation system, allowing users to quickly retrieve climate data by selecting locations or typing place names. It also facilitates deeper data exploration, allowing users to download CSV files to work with locally or access Python notebooks for more advanced analysis. In keeping with the mobile-first approach, the app also supports the creation and sharing of images directly, allowing easy sharing via social media.

## Zarr: revolutionising data processing

Behind the scenes, the newly implemented Data Stores ARCO (Analysis Ready, Cloud Optimised) Data Lake based on Zarr technology optimises data processing for ERA



**ERA Explorer in action.** Here, ERA Explorer is used to compare climatologies and trends for Reading, UK, and Bonn, Germany.

Explorer. Zarr is a modern storage format designed for large, multidimensional datasets like ERA5. It stores data in 'chunks', allowing efficient, on-the-fly access to specific subsets of data. The app uses two Zarr storage formats – one chunked by time and another by space – enabling rapid retrieval of time-series data for a single location or spatial data for broader regions. This enables ERA Explorer to process and display grid-point daily, monthly, and annual statistics, derived from ERA5's hourly data, in a matter of seconds. The use of Zarr is integral to modernising the C3S Climate Data Store (CDS), significantly enhancing the scalability and speed of climate data analysis. As an extension to the ERA Explorer and powered by the same underlying ARCO data, a new entry has been published on the CDS Catalogue, providing improved access to ERA5 point time-series: <https://cds.climate.copernicus.eu/datasets/reanalysis-era5-single-levels-timeseries>.

## Powerful features

Although using modern, minimalist design standards, ERA Explorer offers a wealth of powerful features. The base map can display

temperature or precipitation based on 30-year averages, with features like month-by-month visualisations and animated wind layers.

An interactive onboarding feature guides users into the app, inviting them to explore deeper and more advanced options, such as comparing locations side by side, choosing which climatologies to display, and toggling locations on and off. ERA Explorer provides explanations about the data and data provider, including detailed discussion of the caveats and considerations when interpreting ERA5 data. The accompanying Python notebooks, where users can recreate the same plots themselves, provide a jumping off point for a longer scientific journey.

All told, the design, implementation, technical infrastructure, and dissemination of ERA Explorer represents a significant collaborative effort by many colleagues across ECMWF, supported by the Copernicus Climate Change Service (C3S), external contractor Lobelia for the ARCO implementation, and the Horizon Europe ASPECT project. We welcome any feedback.

# The current state of the climate in Europe

Rebecca Emerton, Julien Nicolas, Freja Vamborg, Anna Lombardi, Francesca Guglielmo, Shaun Harrigan, David Lavers, Annabel Cook, Samantha Burgess, Matthieu Chevallier

On 15 April 2025, the latest European State of the Climate (ESOTC) report was published (<https://copernicus.climate.eu/ESOTC/2024>). ESOTC 2024 provides a detailed overview of climate conditions in Europe and the Arctic in 2024, covering a wide range of climate variables and topics including temperature, heat and cold stress, precipitation, wildfires, glaciers, sea ice, renewable energy resources and many more.

'Spotlight' topics this year include flooding across Europe, and extreme heat and drought in southeastern Europe during the summer. There is also an overview of key findings for trends related to these topics based on information from the Intergovernmental Panel on Climate Change (IPCC). In addition, the report discusses

climate policy and action, focusing on the resilience of the built environment to climate extremes, and provides updates on long-term trends in key climate indicators for the globe and for Europe. ESOTC aims to reach a broad and non-specialist audience, including policymakers, scientists, students, journalists and the wider public.

## What goes into producing the report?

ESOTC 2024 is published jointly by the EU's Copernicus Climate Change Service (C3S), implemented by ECMWF, and the World Meteorological Organization (WMO). It is the result of a collaborative effort involving around 100 scientists from the Copernicus and WMO networks and beyond.

Coordinated by an editorial team at ECMWF, in collaboration with the WMO, this large team of scientists analyse the latest data from around 40 different datasets and report on approximately 40 climate variables, indicators and topics across the report's 19 sections. The report is reviewed by experts from across both organisations' networks, and by representatives from national meteorological and hydrological services.

## How is the report presented?

ESOTC 2024 features a wide range of resources to highlight and communicate the current state of the climate, making this broad and complex topic more accessible:



PROGRAMME OF  
THE EUROPEAN UNION



IMPLEMENTED BY **ECMWF**



**Graphics gallery.** This illustrates the ESOTC 2024 'graphics gallery', showcasing 130 figures exploring the state of the climate in Europe and the Arctic in 2024.



- The ESOTC website enables navigation of the full report, including a range of interactive charts.
- The ESOTC graphics gallery showcases 130 images, with one-click data downloads to easily explore the data behind them (see the first figure).
- An interactive map highlights key events of 2024, alongside examples of climate resilience and adaptation initiatives in European cities.
- A summary distils and highlights key messages from the report, telling the story of the data behind them through a set of infographics.
- A PDF version of the full report is easily browsable.
- An update of key climate indicators and the 'Climate Indicator dashboard' (<https://climate.copernicus.eu/climate-indicators>) are provided.

The graphics gallery and full report PDF are new features for ESOTC 2024.

Data visualisation plays a central role, from the bespoke cover art to mobile-friendly responsive charts, while careful copy-editing ensures the language is clear and accessible to a wide audience. The WMO also provides translations of the summary into multiple languages.

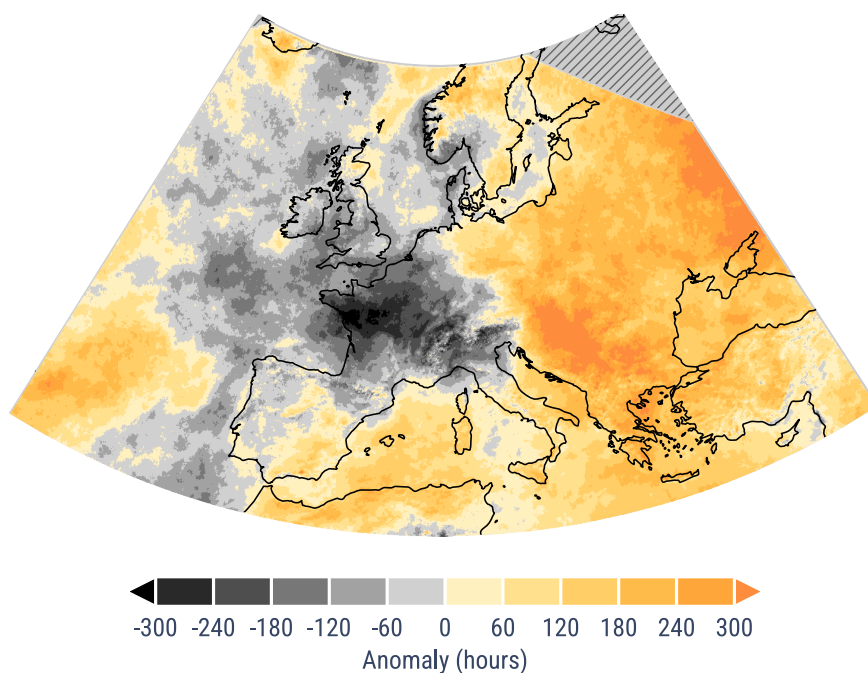
## How are the findings shared with the public?

The findings of ESOTC 2024 reach a wide audience through press coverage and public outreach coordinated by ECMWF's Copernicus communications team. An embargoed media briefing took place on 10 April, followed by an online launch event on 15 April, where key findings were presented.

## What are the key messages?

Since the 1980s, Europe has warmed twice as fast as the global average, making it the fastest-warming continent. In 2024, Europe experienced its warmest year on record.

Europe saw its most widespread flooding since 2013. An estimated 413,000 people were affected by storms and flooding, with at least 335 lives lost. The persistent rainfall



PROGRAMME OF  
THE EUROPEAN UNION



IMPLEMENTED BY **ECMWF**

**EUMETSAT**  
**CM SAF**

**Sunshine duration.** This figure shows sunshine duration anomalies over Europe for 2024, showing positive (shades of orange) and negative (shades of grey) anomalies. Grey hatching in the top right corner of the map indicates missing data. Reference period: 1991–2020. Data: CM SAF SARAH-3 CDR/ICDR. Credit: C3S/ECMWF/DWD/EUMETSAT.

from Storm Boris in September caused river flows to reach at least twice the average annual maximum along 8,500 km of rivers. In October, extreme rainfall led to devastating flooding in Valencia, Spain (see the separate article in this Newsletter). According to the IPCC, Europe is one of the regions with the largest projected increase in flood risk.

2024 was also a year of contrasting climate conditions across the continent. Eastern Europe saw warmer-than-average or record-high temperatures throughout much of the year. Southeastern Europe experienced its longest heatwave on record, lasting 13 days, as the summer brought record numbers of 'strong heat stress' days and tropical nights to the region.

Western Europe, meanwhile, had one of its ten wettest years in the analysed period since 1950, and cloud cover was above average – in some areas, there were as many as 350 fewer hours of sunshine than average (see the second figure).

The Arctic saw its third-warmest year

on record, with contrasting temperature anomalies across the European Arctic during summer. In the east, record-high temperatures contributed to glaciers in Scandinavia and Svalbard experiencing their highest annual rates of mass loss on record and the largest loss of any glacier region globally.

Svalbard, one of the fastest-warming places on Earth, recorded its highest summer temperature for the third consecutive year. Further west, temperatures were mostly near or below average, and the Greenland Ice Sheet recorded its third-smallest mass loss since 2001.

## Where can I find out more?

Visit <https://copernicus.climate.eu/ESOTC/2024> to explore the full report, summary, graphics gallery, interactive key events map, policy sections and more.

*The Copernicus Climate Change Service (C3S) is implemented by the European Centre for Medium-Range Weather Forecasts (ECMWF) on behalf of the European Commission.*



# Support for regional re-forecasts over Montenegro

Angel Marčev (Institute of Hydrometeorology and Seismology of Montenegro), Bojan Kašić

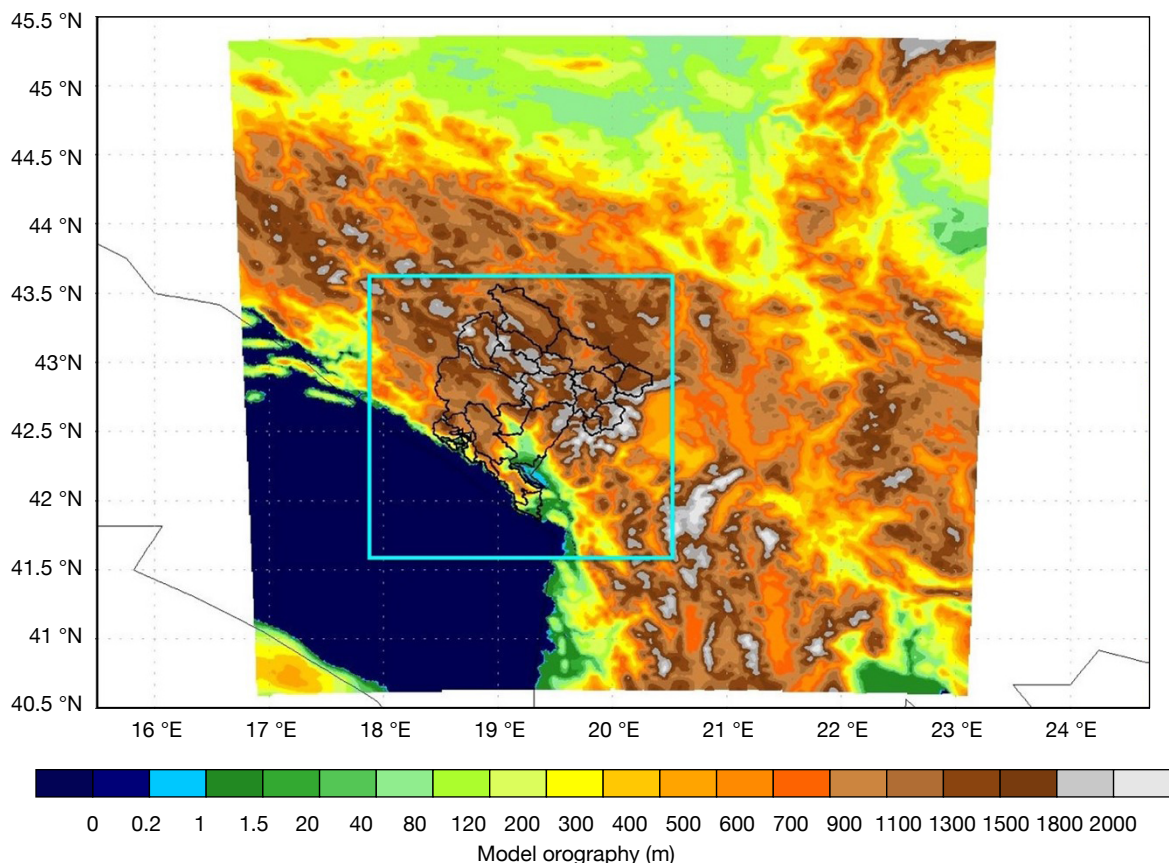
In November and December 2024, under ECMWF's Member and Co-operating State short-term secondment programme (<https://www.ecmwf.int/en/about/jobs/application-member-state-short-term-secondment-ecmwf>), ECMWF hosted Angel Marčev, a visitor from the Institute of Hydrometeorology and Seismology of Montenegro (IHMS), to run limited-area, high-resolution re-forecasts for the past year. This was to be done using a new configuration of the IHMS operational numerical weather prediction (NWP) system on ECMWF's high-performance computing facility (HPCF). Daily re-forecasts over a full year, including all seasons, are crucial for evaluating model performance and developing post-processing techniques preceding operational implementation of a new system. During the visit, with support from ECMWF User Services, the model

was installed on ECMWF's HPCF, re-forecasts for the year December 2023 to November 2024 were completed, and preliminary verification scores were obtained. Montenegro is an ECMWF Co-operating State, which has started the application process to become a Member State (MS). The visit was an opportunity to explore the benefits of having full HPCF access and preparation for becoming a Member State.

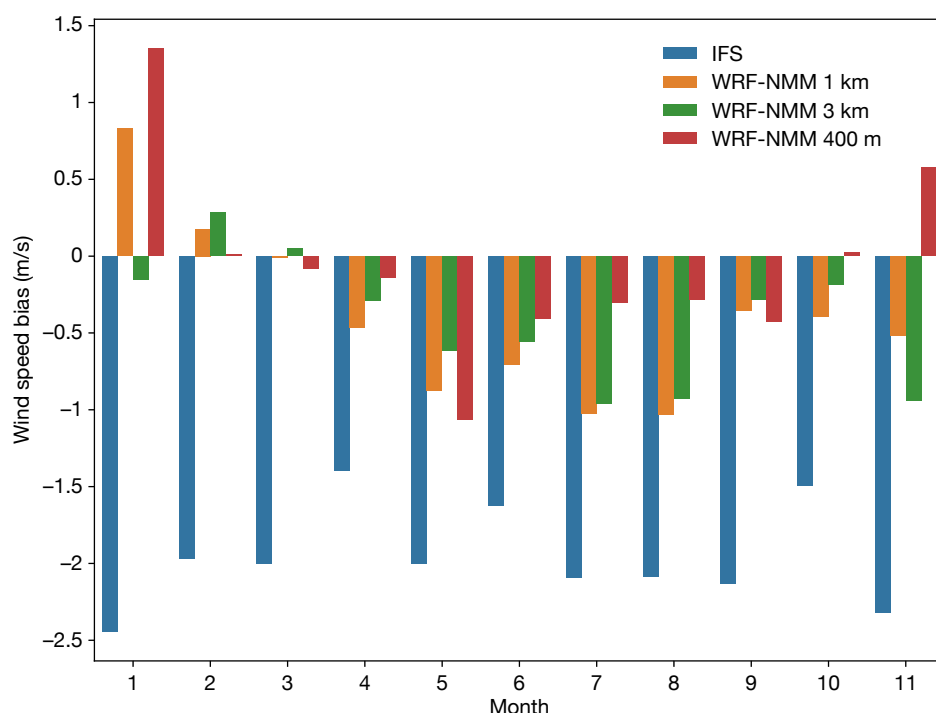
## Background and experiment description

IHMS has a long tradition of running limited-area NWP models to provide operational forecasts. Due to Montenegro's very complex orography, the horizontal resolution of current global models, e.g., 9 km for ECMWF's Integrated Forecasting System (IFS), is insufficient to resolve small-scale

weather patterns, especially those related to strong convection over mountainous terrain. Therefore, IHMS significantly relies on high-resolution non-hydrostatic regional NWP to supplement the global IFS in providing reliable local forecasts and issuing weather warnings. The operational NWP system at IHMS is based on the Non-hydrostatic Mesoscale Model (NMM) on E-grid. The dynamical core was developed by the US National Oceanic and Atmospheric Administration/National Centers for Environmental Prediction inside the Weather Research and Forecast (WRF) system known as WRF-NMM. The model scientific documentation by Janjic et al. is available at: [https://dtcenter.org/sites/default/files/community-code/hwrf/docs/scientific\\_documents/WRF-NMM\\_2018.pdf](https://dtcenter.org/sites/default/files/community-code/hwrf/docs/scientific_documents/WRF-NMM_2018.pdf). One of the main limitations for NWP developments at



**Two re-forecast domains.** This map comprises the entire Domain 2 showing model orography in metres (1 km resolution), with a nested sub-domain, which is Domain 3 (400 m resolution).



**Wind speed bias at Dubrovnik Airport.** Average monthly wind speed bias in 2024 against observations at Dubrovnik Airport for the IFS as well as for WRF-NMM at three resolutions.

IHMS is a lack of high-performance computational resources. As an ECMWF Co-operating State, Montenegro has access to all products from ECMWF's global operational forecast and to some computing facilities. This includes ECGATE Class Services (ECS) and the European Weather Cloud (EWC), but not ECMWF's full HPCF, which is made available only to Member States. Current on-premises infrastructure in Montenegro is sufficient for running high-resolution regional models, for example at 1 km over small domains, but it does not enable extensive testing of sensitivity to model configurations and evaluating performance before a specific configuration is deployed to operations. Furthermore, the current infrastructure does not enable the production of long-term re-forecasts for calculating the model climatology and archiving NWP results for a long time. For this reason, IHMS approached ECMWF asking for support in doing re-forecasts with their new NWP configuration based on WRF-NMM. The re-forecasts of 2024 on ECMWF's Atos HPCF during the visit covered three regional domains, two of which are shown in the first figure:

- **Domain 1:** with 3 km horizontal resolution, driven by initial and boundary conditions (IBC) from the IFS
- **Domain 2:** with 1 km horizontal resolution, nested in Domain 1

- **Domain 3:** with 400 m horizontal resolution, driven by Domain 2

Running all three WRF-NMM domains requires 10 nodes on the Atos HPCF to meet the operational requirement of completing one day of forecasts in seven minutes and completing one year of daily forecasts 72 hours ahead in less than one month.

All three domains were run for the full year from December 2023 to November 2024, daily, starting at 00 UTC to 72 hours ahead. Initial and boundary conditions from the IFS were extracted from ECMWF's Meteorological Archival and Retrieval System (MARS). Hourly model results from all three domains were stored on the HPCF. A subset of fields required for further processing was archived in a Simple Storage Service (S3) bucket of the EWC, using the storage allocation of Montenegro granted to Member and Co-operating States. These results have been used to evaluate forecast quality and calculate model biases with bias correction for key observation points in Montenegro. The same archive will be used later to apply existing and develop new post-processing techniques based on machine learning to improve high-resolution forecast products further.

### Initial verification

Some verification scores have already been produced for all World Meteorological Organization (WMO)

stations in Montenegro and observational stations covered by Domain 3 in the surrounding countries. In the second figure, wind speed monthly bias for the airport in Dubrovnik is shown for each WRF-NMM domain as well as for IFS direct model output. Despite the overall good match with observations, the global IFS tends to underestimate the wind speed due to its coarser resolution and a simplified representation of the complex orography around the observation station, compared to the WRF-NMM models. Negative bias for the high-resolution regional model is clearly reduced, while the Domain 3 configuration, with the highest horizontal resolution of 400 m, provides the best results overall.

As an added value of this visit, Angel improved his knowledge about ECMWF's HPCF, the EWC, and other computing facilities. He was also able to familiarise himself with software packages such as ECMWF's workflow manager ecFlow, which can be implemented locally at IHMS to support operational and research workflows.

*This work was supported by the German Early Career Fellowship programme and the Senior Research Visitors programme, which funded the secondment, and by the Republic Hydrometeorological Service of Serbia (RHMS), which generously provided the required HPCF resources from its ECMWF Member State allocation.*

# Operationalisation of global climate reanalysis led by ECMWF

Alison Cobb

In 2024, the World Meteorological Organization (WMO) approved the inclusion of global climate reanalysis in its WMO Integrated Processing and Prediction System (WIPPS), which means it is now formally an operational activity. This will ensure that global reanalyses are continually produced and made available to users, alongside clear documentation and visualisation tools. ECMWF is one of three current designated centres, contributing the ERA5 reanalysis, and it is also the lead centre of this activity.

## WMO WIPPS

The WMO has an intergovernmental mandate for coordinating the generation and exchange of weather, climate, and water information across its members. It has successfully coordinated the production and provision of weather forecasts and climate predictions from international operational centres, from short-range to medium-range, sub-seasonal, seasonal and decadal timescales.


In June 2024, the WMO Executive Council (EC-78) approved the inclusion of two WIPPS-Designated Centres (DCs) related to climate reanalysis:

- WIPPS-DCs for Global Climate Reanalysis
- WIPPS-DCs for Coordination of assessment of multiple climate reanalysis

After the approval by EC-78, climate reanalysis is now formally an operational activity within WIPPS and is on a par with forecasting activities.

## Designated centres

At present, there are three WIPPS-DCs for Global Climate Reanalysis: ECMWF, the China Meteorological Administration (CMA), and the US National Aeronautics and Space Administration (NASA).



PROGRAMME OF  
THE EUROPEAN UNION



**Surface air temperature anomalies.** This image shows surface air temperature anomalies in 2024 according to the ERA5 reanalysis. The reference period is 1991–2020.  
*Credit: C3S/ECMWF*

The Japanese Meteorological Agency (JMA) is not formally part of the project yet but hopes to be made an official WIPPS-DC during the summer of 2025.

CMA, ECMWF, and NASA will provide CRA-40 (1979–present), ERA5 (1940–present), and MERRA2 (1980–present), respectively. Finally, JMA will contribute JRA-3Q (September 1947–present) once it is officially a designated centre. The approximate horizontal resolutions of CRA-40, ERA5, MERRA2, and JRA-3Q are 34, 31,

50, and 40 km, respectively.

ECMWF is the WIPPS-DC for Coordination of assessment of multiple climate reanalysis. In this capacity, ECMWF will facilitate intercomparison of global reanalysis products from different centres, with the provision of comparable data on identical grids, graphical products, and visualisation tools.

## Development phases

During the initial phase of this activity, we are gathering historical monthly mean data from 1991–2020 (the



WMO climatological period). We are starting with three principal variables of total precipitation, 2 m temperature, and mean sea level pressure.

In the future, additional mandatory products at the surface will be added, including pressure, land-sea mask, topography, sea-surface temperature, sea-ice cover, snow water equivalent, incoming short-wave radiation, outgoing long-wave radiation, 2 m specific humidity, 10 m winds, alongside geopotential height,

temperature, winds, and specific humidity data on pressure levels.

Along with the inclusion of more variables, in future phases we will gather data at higher temporal resolution, and within no more than 60 days behind real time.

Global climate reanalyses are widely used in a range of diverse applications, and this significant step forward in making them operational will benefit many sectors. We are working closely with

CMA, JMA, and NASA on the provision of initial data for phase 1 and will ultimately host a platform with several mandatory products that are regularly updated, with the ability to visualise data and conduct intercomparisons.

ECMWF's web page on being the lead centre for global climate reanalyses is here: <https://confluence.ecmwf.int/display/GCR/WMO+Lead+Centre+for+Global+Climate+Reanalyses+LC-GCR>

## Impact assessment of Chinese hyperspectral infrared sounder in preparation for MTG-S IRS

Naoto Kusano (JMA), Chris Burrows (ECMWF)

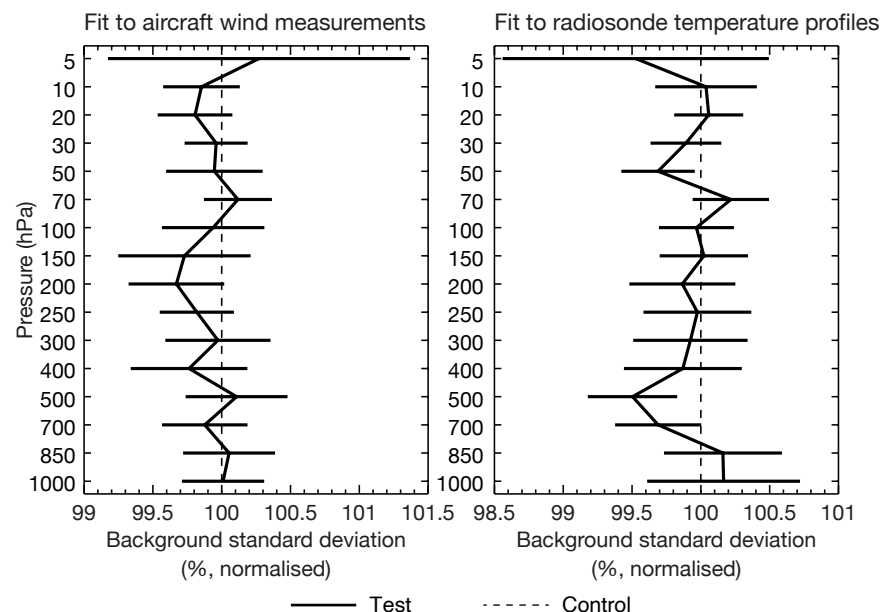
EUMETSAT's first Meteosat Third Generation Sounder satellite (MTG-S1) is planned to be launched in July 2025, carrying the next-generation hyperspectral infrared instrument (InfraRed Sounder, IRS). This instrument will measure infrared radiances from the Earth with high spectral resolution and, from its geostationary orbit, will give us detailed information on atmospheric structures both spatially and temporally over Europe and Africa. In preparation for MTG-S IRS, the Geostationary Interferometric Infrared Sounder (GIIRS) onboard China's second Fengyun-4 Series satellite (FY-4B) has been assessed, for monitoring and assimilation in ECMWF's numerical weather prediction (NWP) system. FY-4B GIIRS was launched in 2021 and observes east Asia. It is the first operational hyperspectral infrared instrument on a geostationary platform. Its nominal characteristics are very close to future hyperspectral infrared instruments on geostationary satellites, including MTG-S IRS. Therefore, experiences gained by investigating FY-4B GIIRS can help to prepare us for MTG-S IRS. Exploitation of data from GIIRS presents some challenges, some of which are specific to the instrument. However, we show here that these observations can have a positive impact on forecast scores, improving

temperature and humidity in the lower troposphere over the GIIRS domain.

### FY-4B GIIRS data assessment

FY-4B GIIRS is generally performing better than its predecessor onboard the first Fengyun-4 Series satellite (FY-4A). The mid-wave infrared band of FY-4A GIIRS had several issues

which limited its potential use in NWP. These have been greatly improved for FY-4B thanks to work by the China Meteorological Administration (CMA). However, there are still some challenging artefacts in the FY-4B GIIRS data, including systematic differences in the pixels across the detector, and horizontally-correlated observation



**Fit to independent observations.** These two panels show the fit of 12-hour IFS forecasts to aircraft wind measurements (left) and radiosonde temperature profiles (right) over the GIIRS observational domain (east Asia) when assimilating GIIRS observations for the period between 8 June and 31 August 2023. The 100% line refers to the control experiment. Error bars show the 95% confidence range.

errors in high-peaking temperature channels. These can be addressed by using only a subset of the 128 pixels and by selecting a subset of channels with relatively small horizontal correlations.

In addition, CMA performs two yaw flip manoeuvres to FY-4B in March and September to prevent solar intrusion into the sensor. In the yaw flip manoeuvre, the satellite rotates about the axis pointed toward the centre of Earth. When the manoeuvre takes place, the biases of channels with wavenumbers less than  $742.5\text{ cm}^{-1}$  suddenly jump with a change of more than 0.5 K. This is possibly due to thermal emission from the scanning mirrors. MTG-S will also have yaw flip manoeuvres, and EUMETSAT will monitor the internal temperature of the satellite after the launch of MTG-S to try to mitigate this.

### Impact assessment in ECMWF's 4D-Var system

The impact of FY-4B GIIRS data in ECMWF's 4D-Var data assimilation system was tested under clear-sky

conditions in Cycle 49r1 of the Integrated Forecasting System (IFS) for a summer season period, when the yaw flip manoeuvre was not performed. A carefully chosen subset of channels sensitive to atmospheric temperature, surface temperature, ozone and water vapour was used, and only a subset of pixels which have consistent statistics. Otherwise, the assimilation methodology is similar to that used for the other hyperspectral infrared sounders in the IFS system, including the cloud detection method, the use of an observation error covariance matrix with appropriate inter-channel correlations, and variational bias correction.

Globally, the impact on temperature, humidity and ozone forecasts is neutral, but it is positive over the GIIRS observational domain. Some improvements in the short-range forecast of temperature, humidity, wind and ozone are found with respect to independent observations, as shown in the figure. Medium-range forecast scores in the lower troposphere are also promising; the assimilation of FY-4B GIIRS results in

statistically significant forecast improvements for temperature and humidity as verified against own analyses in the lower troposphere over the GIIRS observational domain until forecast day 4. The root-mean-square error of forecasts decreases by up to about 1%.

### Summary and outlook

The assimilation of FY-4B GIIRS has a neutral-to-positive impact on forecast skill, and passive monitoring of these data will begin soon as part of our NWP SAF (Satellite Application Facility) activities. However, before considering operational assimilation of FY-4B GIIRS, additional assimilation experiments will be run for a period including a yaw flip manoeuvre, and also a winter season period will be considered to confirm the assimilation of FY-4B GIIRS has robust impact. This investigation of FY-4B GIIRS has given us some interesting insights about hyperspectral infrared sounders on geostationary platforms and helps prepare the way to assimilate MTG-S IRS when the data become available.

## Evolution of the Regional Meteorological Data Communication Network

Oliver Gorwits

The Regional Meteorological Data Communication Network (RMDCN) is a specialised communication network supporting meteorological data exchange between 53 national meteorological and hydrological services (NMHSs) and other meteorological organisations around the world. The journey of the RMDCN is one of continual renewal. Run by ECMWF and now in its 26th year, we recently completed another Technical and Commercial Refresh exercise, designed to maintain performance and value for money.

### History of the RMDCN

Before the RMDCN, meteorological data was exchanged using many bilateral links established between

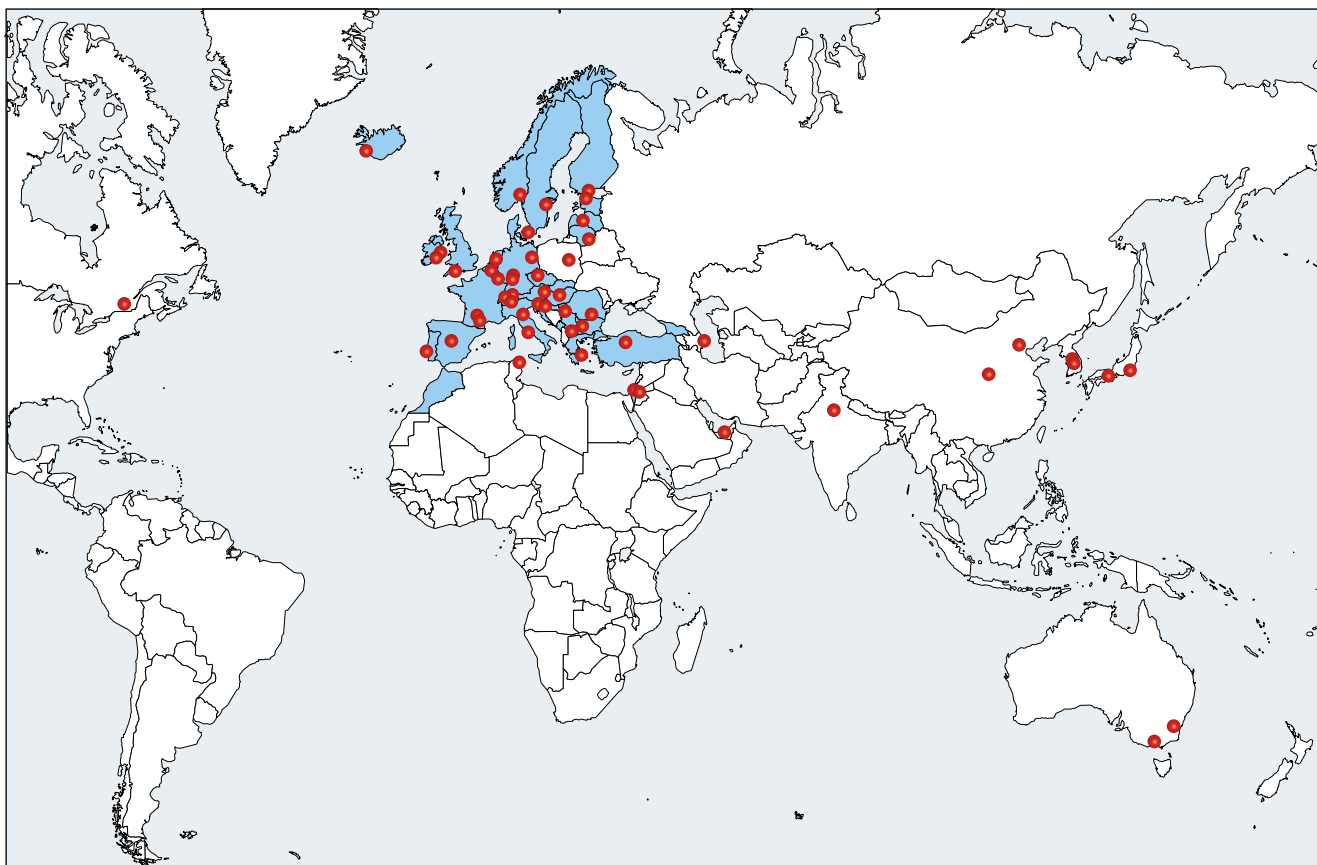
pairs of organisations wishing to connect their computer systems. In the late 1990s ECMWF proposed to replace this costly and complex mesh with a single managed service in the European region, to serve the needs of ECMWF and those of World Meteorological Organization (WMO) Region VI. In agreement with the WMO, the service is also offered to Global Information System Centres (GISCs) around the world to enable reliable and timely delivery of weather data and bulletins.

During its first few years of operation, the RMDCN was built using leased lines and Frame Relay technology. In 2006 the network was upgraded to use Multiprotocol Label Switching (MPLS) Internet Protocol Virtual Private Network (IP-VPN) –

the backbone technology of the Internet as we know it today. This enabled the service to become truly 'any-to-any', meaning any site can choose to communicate directly with another site without needing to involve ECMWF.

### Technical and Commercial Refresh

In 2012, the service evolved again with the award of a new commercial contract transitioning to a new service provider. To support long-term stability for the service, which was now an established and critical part of the global meteorological infrastructure, the new contract allowed for a periodic Technical and Commercial Refresh (TCR) to be conducted on a three-yearly cycle.



**Meteorological communication network.** RMDCN connections at meteorological organisations around the world. The blue colour indicates ECMWF Member and Co-operating States.

This exercise involves the provider of the service proposing either a bandwidth upgrade or reduced costs on a per-site basis, with sites being free to choose their preferred option. ECMWF reviews the offer together with consultants familiar with the IP-VPN market, to ensure pricing and technology provide value for the investment. Some negotiation may be required to refine the proposals to reflect current market rates.

As we have seen on previous occasions, this latest TCR resulted in an approximate doubling of bandwidth for the same price, reflecting progress in internetworking technologies in the intervening three years. The base connection speed for ECMWF's Member States was upgraded from 20 Mbps (megabits per second) to 50 Mbps, with many sites opting to purchase more, up to around 200 Mbps.

By modern standards these are not high capacity links, particularly considering the size of today's numerical weather prediction outputs. However, the very high

level of reliability and security assurance provided under the RMDCN contract has kept the service relevant, providing resilience to deliver the most essential observations and products.

Many lessons have been learnt along the way by ECMWF's Computing Department, procurement, and legal teams. The RMDCN's contract envisaged a TCR exercise every three years. In practice, however, this stretches out somewhat, as the implementation of a TCR turns out not to be straightforward. For example, it can require significant lead times for user sites to secure budget approval to make decisions. The expected time to commission a service upgrade is approximately 90 days, but it can take longer. Occasionally, unexpected challenges are encountered, and with many subcontractors involved and over 100 connections to deliver and test, project management is key to a successful implementation of the TCR.

## Future of the RMDCN

Times are changing, though, and to improve efficiency the WMO has developed WIS 2.0, a new architecture for information sharing within the weather community. Notably WIS 2.0 proposes migrating to the technologies of the Internet and the use of open-standard web technologies for data publication and exchange. WMO Members consider that the Internet is sufficiently robust and mature to support safety-critical applications – so long as those applications are built in a fault-tolerant way. Many NMHSs around the world are already using the Internet to support all or most of their time-critical and safety-related data exchange.

For this reason, the current RMDCN contract has been extended until 2028, but it is likely this recently completed Technical and Commercial Refresh will be the last. This year, ECMWF will carry out an internal task to survey readiness for WIS 2.0 among its community, and to consider any steps or service developments necessary to address gaps.



# AI Weather Quest: Advancing sub-seasonal forecasting with AI/ML

Olga Loegel, Joshua Talib, Frédéric Vitart, Jörn Hoffmann

Predicting weather beyond two weeks remains one of the greatest challenges in meteorology. While traditional models are very good at short- and medium-range forecasting, their accuracy declines significantly at sub-seasonal timescales (15 days to 2 months ahead). This gap is crucial, as improved sub-seasonal forecasts can support decision-making in sectors such as energy, agriculture, and disaster risk management.

Recent advances in artificial intelligence (AI) and machine learning (ML) offer exciting new possibilities for improving weather forecasting, particularly at the sub-seasonal scale. AI/ML models can extract patterns from vast datasets and generate forecasts that may complement traditional physics-based approaches. To address this, ECMWF is launching the AI Weather Quest, a global competition aimed at pushing the boundaries of AI-driven sub-seasonal forecasting.

This initiative, endorsed by the World Meteorological Organization (WMO), builds upon the foundation of the 2021 WMO Sub-seasonal to Seasonal (S2S) AI Challenge.

## A global challenge for AI and meteorology experts

The AI Weather Quest invites participants from around the world to develop and submit AI/ML-based sub-seasonal forecasts. The competition aims to benchmark the performance of AI-based models, exploring how they can enhance operational forecasting.

Starting in March 2025, the competition will unfold in two phases:

- Initial Training Phase (March–August 2025): Participants will refine their models and familiarise themselves with the competition's submission and evaluation process in a non-competitive environment.
- Competition Phase (August 2025–at least September 2026): Participants will submit weekly,



real-time forecasts over four 13-week periods. They will be evaluated based on the Ranked Probability Skill Score (RPSS), comparing their forecasts to established benchmarks.

Forecasts will focus on three key variables: near-surface (2 m) temperature, mean sea level pressure, and precipitation. AI models must predict quintile probabilities for two lead times: days 19–25 and days 26–32, providing a probabilistic outlook crucial for decision-making in weather-sensitive industries. Forecasts will be required at a 1.5-degree resolution, with evaluations based on ECMWF's ERA5 reanalysis datasets.

Participants can submit up to three different AI models and provide up to 18 submissions per week (three variables across two lead times for up to three models). Models can be developed using any programming language and dataset, including observational data, reanalysis products, and existing physics-based sub-seasonal forecasts.

## Open participation and recognition

The competition is open to a wide range of participants, including AI/ML researchers, meteorologists, participants from startups, large technology companies, and public forecasting institutions. No prior experience in weather forecasting is required – only a strong interest in leveraging AI to tackle real-world forecasting challenges.

ECMWF will compute and publish weekly RPSS scores, with leaderboards displaying the best-performing models. At the end of each 13-week competition period, the top teams will be recognised in dedicated award celebrations, showcasing leading AI-driven forecasting approaches. To ensure inclusivity and fairness, ECMWF will also highlight exceptional models from diverse organisation types and those developed with limited computational resources.

## Resources for participants

To support model development, participants have access to a range of resources, including historical datasets such as ERA5. In addition, ECMWF's Open Data Catalogue offers valuable real-time and historical forecasting products. A dedicated AI Weather Quest (AI-WQ) Python package is available to facilitate forecast submission and evaluation, ensuring consistency and transparency throughout the competition.

## ECMWF's role and vision

As one of the world's leading weather forecasting institutions, ECMWF is uniquely positioned to explore the integration of AI into operational meteorology. By hosting this competition, ECMWF seeks to evaluate AI/ML potential for sub-seasonal forecasting and identify promising AI-driven approaches. The findings from this initiative could help guide the future of operational forecasting.

## Get involved

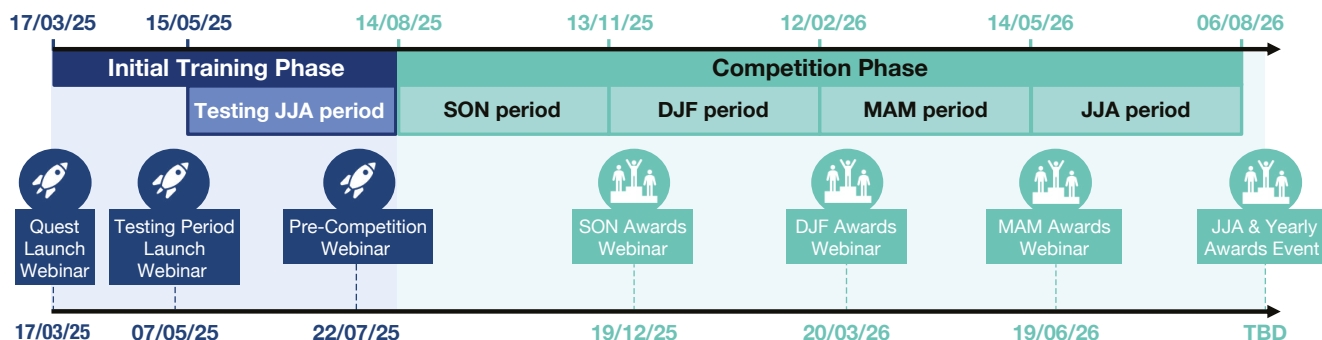
The AI Weather Quest is more than just a competition, it is a unique opportunity to contribute to the future of weather forecasting. Whether you

are an AI expert, a meteorologist, or a company interested in applying machine learning to forecasting, this challenge provides a platform to test ideas, gain visibility, and make

valuable connections.

Registration is now open. For more details, visit the AI Weather Quest website at <https://aiweatherquest.ecmwf.int/>.

## Timeline and structure



**The competition's structure.** The AI Weather Quest includes a five-month Initial Training Phase and a Competition Phase that will last for at least a year.

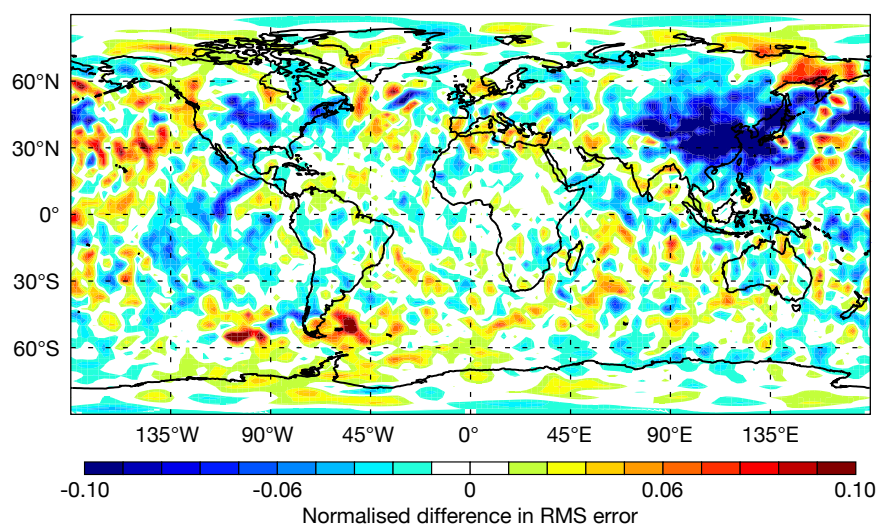
# Kilometre-scale modelling with the nonhydrostatic dynamical core

Jozef Vivoda, Inna Polichtchouk, Michail Diamantakis, Filip Váňa

At ECMWF, two nonhydrostatic dynamical cores are available: a nonhydrostatic spectral transform, which extends the hydrostatic spectral formulation used in the operational Integrated Forecasting System (IFS), and a grid-point dynamical core based on the finite volume method, designed for scalability at high resolutions. This study focuses on the spectral

nonhydrostatic dynamical core (NH-IFS), originally developed for the ALADIN model and later extended for the ARPEGE model by Météo-France. ECMWF shares this component of the model code with its Member States. NH-IFS uses finite element discretisation in the vertical to ensure full compatibility with its hydrostatic counterpart (HY-IFS).

It is generally accepted that nonhydrostatic (NH) effects become relevant at horizontal grid spacings finer than 10 km, and models at 1 km grid spacing or finer must use NH dynamics to maintain physical realism. The EU's Destination Earth (DestinE) initiative, which is partly implemented by ECMWF, aims to push global numerical weather prediction (NWP) toward these horizontal scales. Therefore, the significance of NH effects must be tested experimentally within the full framework of the spectral IFS to determine the grid



**Difference in error.** Normalised difference (NH-IFS vs HY-IFS) in root-mean-square (RMS) error for 96-hour wind forecasts at 100 hPa computed for 20 January to 29 February 2022. The dark blue colour represents a reduction of root-mean-square error by 10% and more. Computed with NH-IFS and HY-IFS at a horizontal grid spacing of 2.8 km.

spacing at which they become essential for accurate medium-range weather forecasting.

## Ensuring consistent comparisons

To ensure a fair and consistent comparison between HY-IFS and NH-IFS, the following key improvements have been introduced into NH-IFS:


- **Stabilizing model time-stepping and improving its accuracy:** NH-IFS required modifications to its time integration scheme to maintain numerical stability and improve its accuracy. This was achieved by ensuring full consistency between the nonlinear and linear parts of the dynamics, improving the coupling of the iterative time-stepping scheme with physics, adapting NH-IFS for single-precision computations, and suppressing numerical noise through slight first- or second-order decentering.
- **Ensuring NH dynamics is an extension of HY dynamics:** NH-IFS is now formulated as a direct extension of HY-IFS. To ensure consistency, all terms common to both models employ identical numerical methods. As a result, NH-IFS solutions match HY-IFS outputs at scales where the hydrostatic approximation holds.
- **Switch to a unified NH/HY vertical finite element (VFE) scheme:** ECMWF is the first centre performing such a comparison with a VFE scheme; other partners use a finite difference scheme version of NH dynamics.

## Case studies and systematic verification

We compare NH-IFS with HY-IFS at 9 km, 4.4 km, 2.8 km, and 1.4 km grid spacings using cubic spectral truncation. The goal is to assess the impact of nonhydrostatic effects across different scales, focusing on two phenomena where such effects are theoretically expected to appear:

- gravity wave propagation over complex terrain during strong near-surface flows
- tropical cyclones, where large-scale convection dominates.

Nonhydrostatic effects primarily



**Kinetic energy spectra.** The image shows kinetic energy spectra of NH-IFS and HY-IFS for runs with various horizontal grid spacings from 9 km up to 1.4 km for 36 h forecasts at a level of 50 hPa, computed for 21 December 2021. The differences appear at a resolution of around 2.8 km for waves with wavelengths shorter than 20 km.

influence the dispersion properties of small-scale gravity waves (GWs), affecting the characteristics of mountain-induced waves above major mountain ranges. While the HY-IFS model restricts wave propagation predominantly to the vertical, NH-IFS supports more dispersive wave dynamics, including significant horizontal propagation and the vertical trapping of small-scale modes. This increased dispersion tends to limit the upward transmission of wave energy, resulting in fewer waves reaching the upper troposphere and stratosphere, depending on background atmospheric conditions. These waves can generate clear-air turbulence, and momentum deposited by such waves can influence jet streams in the troposphere and stratosphere.

For this evaluation, January and February 2022 were selected – a period characterised by persistent strong near-surface winds above the Himalayas. The first figure shows that, under such conditions, NH effects are evident in +96 h forecasts, and the root-mean-square wind error at 100 hPa is significantly reduced. NH-IFS propagates less gravity wave energy associated with mountain waves into these regions of the atmosphere, consistent with established theory. Such signals can be seen only in the stratosphere above the Himalaya region in the winter period.

The influence of NH effects on tropical cyclones was also investigated. However, no significant impact was found when analysing maximum wind speeds and minimum surface pressure.

The second figure illustrates how model resolution affects the manifestation of NH effects. The kinetic energy spectra, representing energy associated with various horizontal scales at model level 50 hPa, show that differences begin to appear at a grid spacing of 2.8 km for horizontal scales shorter than 20 km.

## Conclusion

Recent advances in nonhydrostatic modelling at ECMWF have made the NH-IFS stable at horizontal grid spacings down to 1.4 km. This opens up opportunities to investigate the role of nonhydrostatic processes at the scales targeted by the DestinE initiative. However, studies have demonstrated that, at the 4.4 km grid spacing currently used in the DestinE operational suite, it is not essential to consider nonhydrostatic dynamics. Given that NH-IFS is approximately 2.5 times more expensive in terms of computational resources compared to HY-IFS, its use must be justified by compelling evidence. Systematic evaluation of its added value will continue, including verification against high-resolution observational data.



# EarthCARE data begin to make an impact

Mark Fielding, Marta Janisková, Shannon Mason, Robin Hogan, William McLean, Angela Benedetti, Richard Forbes

**A** rich stream of never-seen-before cloud and aerosol observations has begun flowing through the ECMWF data centre. Back in 2024, on a fittingly cloud-filled Californian May morning, the EarthCARE (Cloud, Aerosol and Radiation Explorer) satellite was launched into space from Vandenberg Space Force Base. Due to its size, a SpaceX Falcon 9 rocket was used to deliver EarthCARE into a sun-synchronous orbit at an altitude of 400 km. EarthCARE is the 8th European Space Agency (ESA) Earth Explorer mission and a joint venture between ESA and the Japan Aerospace Exploration Agency (JAXA). Equipped with four synergistic instruments – a Doppler radar, a high-spectral-resolution lidar, a multi-spectral imager and a broadband radiometer – EarthCARE provides unprecedented insights into the interactions between clouds, aerosols, precipitation and radiation. These processes play a key role in the Earth's weather and climate, yet they are some of the least understood. As greenhouse gas concentrations continue to rise, and severe weather is projected to become more frequent and intense, high-quality observations are essential for improving our understanding of the atmosphere and improving our climate and numerical weather prediction models.

ECMWF's forecast model, like many other models, exhibits regional biases in clouds and their radiative effects. These inaccuracies lead to biases in temperature and winds that affect the fidelity of ECMWF's medium-range forecasts. The biases become even more important for seasonal prediction, which requires a model to be as free of 'drifts' as possible. EarthCARE offers new observations to help pinpoint where our model is going wrong – not just in cloud microphysical properties, such as rain drop size and cloud water amount, but also something that has never been quantified from space before: the fall-speed of snow and rain. These much-anticipated observations are already being used to evaluate the cloud physics within ECMWF's Integrated Forecasting System (IFS).

For NWP, EarthCARE offers a second, more direct benefit through its use in data assimilation. The 4D-Var assimilation of humidity- and cloud-sensitive observations at ECMWF to improve the initial conditions

of forecasts has increased dramatically over the past decade. This has been driven by advances in 'all-sky' microwave radiance techniques, improvements in cloud representation within both the forecast model and the data assimilation system, and a growing number of meteorological satellites in orbit. Including active remote sensing measurements of clouds and precipitation from EarthCARE could further enhance the initial conditions produced from ECMWF's assimilation system.

In 2016, work began to enable the direct assimilation of space-borne radar and lidar observations within the ECMWF 4D-Var system. Once complete, using CloudSat and CALIPSO data as proxies for EarthCARE's radar and lidar, tests demonstrated a direct impact on forecast skill (Janisková & Fielding, 2020). Assimilating radar reflectivity and lidar backscatter improved large-scale temperature and wind forecasts, as well as the analysis of clouds and precipitation. While CloudSat and CALIPSO were never considered operational missions, their influence on data assimilation highlighted the potential value of EarthCARE. This is why EarthCARE's data is being made available in near real-time (often within two hours of observation), unlocking its potential for operational forecasting. In this article, we focus on the impact of EarthCARE on data assimilation in NWP.

Box A contains an overview of the instruments carried by EarthCARE and the measurements they provide.

## Taking all-sky assimilation to a new dimension

In contrast to all-sky microwave sounder observations, which have a broad sensitivity to the vertical profile of temperature and humidity, radar and lidar observations observe the atmosphere layer-by-layer with much higher resolution. The complex structure of clouds that is revealed is a mixed blessing when it comes to using the observations in data assimilation. While the new observations contain a huge amount of information on the atmospheric state, the 4D-Var data assimilation system at ECMWF fundamentally assumes a linear relationship between changes in the measurement (e.g. radar reflectivity) and changes in model state variables (such as temperature and humidity). If the model state is too far from the observations, this linearity assumption will be invalid and will lead to a sub-optimal minimisation and analysis. To mitigate these challenges, a raft of new developments was needed.

## a

### EarthCARE leaves clouds and aerosols with nowhere to hide

The mission's primary goal is to provide so-called 'radiation closure' for NWP and climate models. This means that the top-of-atmosphere radiative fluxes, which are measured by EarthCARE, can be predicted based on EarthCARE retrievals of the clouds, aerosols and precipitation underneath. To achieve this, four complementary sensors are used (see Figure 1). Firstly, the ATmospheric LIDar (ATLID), a UV high spectral resolution lidar (HSRL), is sensitive to aerosols and clouds, but is quickly attenuated in liquid cloud. By measuring the spectrum of backscattered laser to distinguish signals from particles and the air, it unambiguously retrieves the optical depth of clouds and aerosol within each lidar range bin.

Figure 2 showcases ATLID's capabilities with an example of the total attenuated particulate backscatter. Over northeast England, a distinct aerosol layer at 6 km resides above isolated cumuli at 2 km. Further south are multiple layers of altocumulus, with the lower layer being capped by a thin layer of supercooled liquid water that fully attenuates the lidar signal.

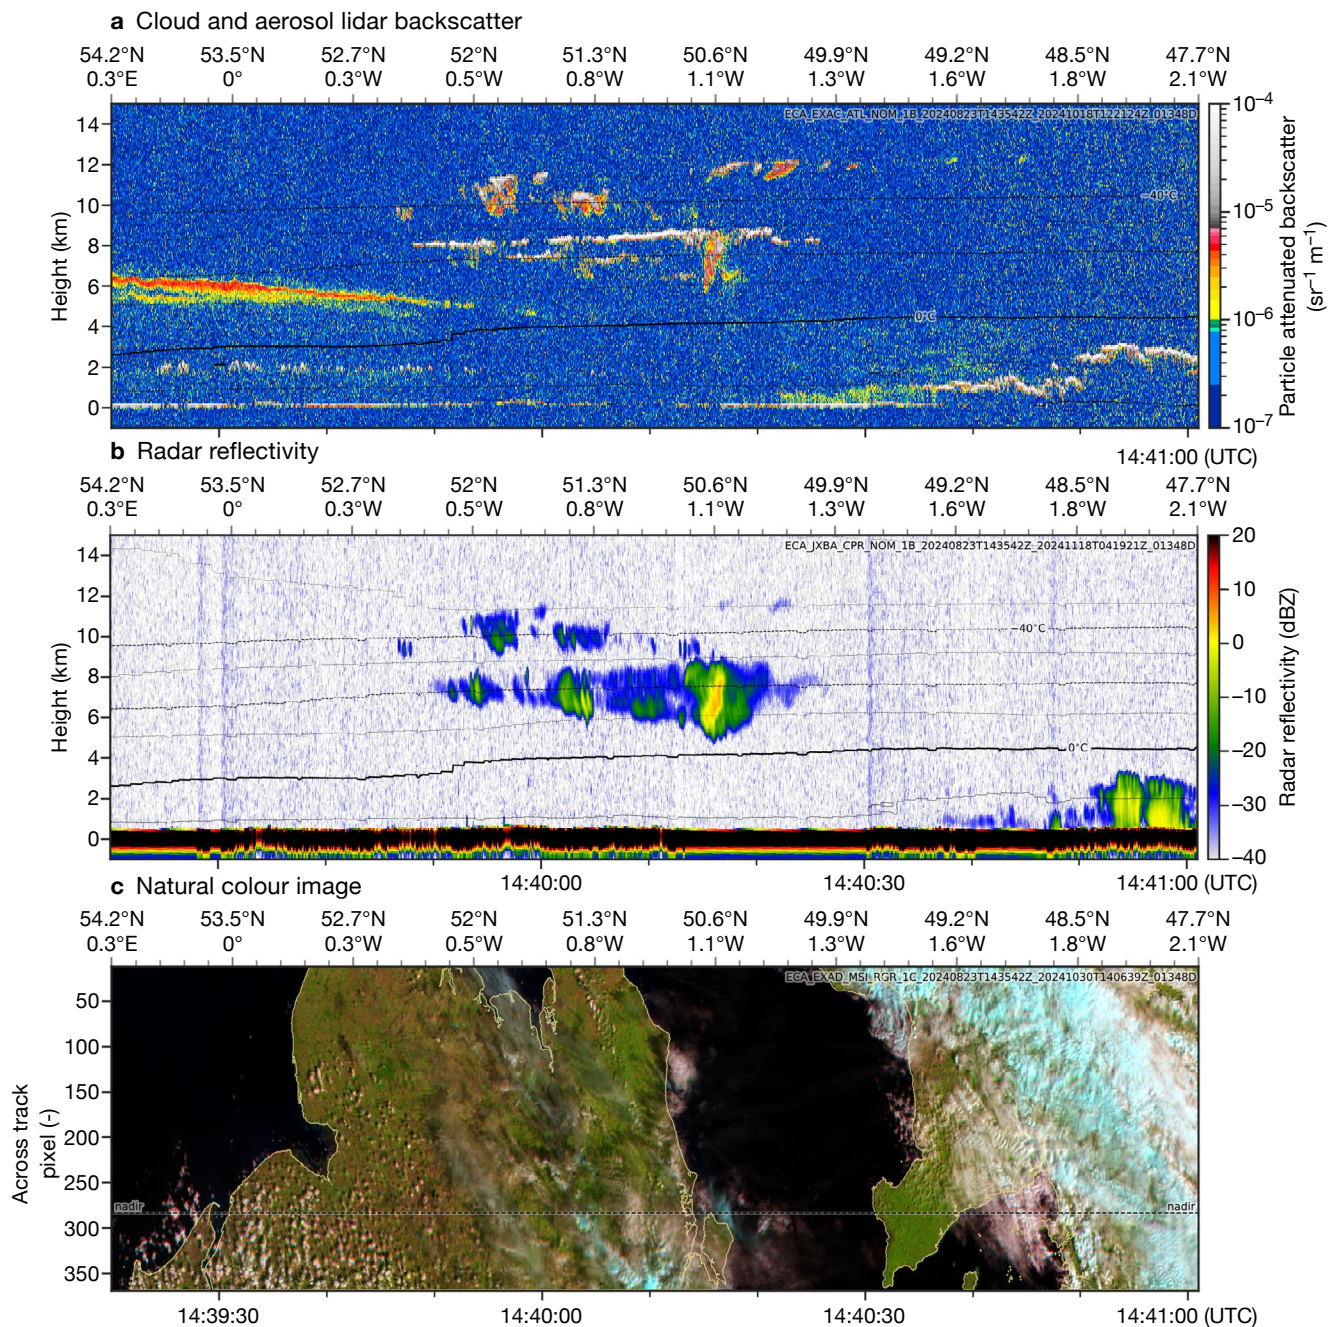
To reveal cloud structure further down, the W-band cloud profiling radar (CPR) is used. The CPR measures the backscatter from radar pulses that are sensitive to cloud and precipitation drops with a radius of around 10  $\mu\text{m}$  and larger. Areas of greater reflectivity indicate larger and/or more numerous particles. Looking again at the altocumulus in Figure 2, we can now see snow underneath the supercooled liquid layer, likely formed through rapid vapour deposition onto ice crystals where the air is super-saturated with respect to ice. Part of the cloud has exhausted its supercooled liquid supply, but has precipitation falling to an altitude of 4 km.

As an additional constraint on the particle size of clouds and aerosols, a Moderate Spectral-resolution Imager (MSI) observes radiances across several different wavelengths in the visible and near-infrared. The bottom panel in Figure 2 shows how the different wavelengths can also be used to separate liquid (white) and ice (blue) clouds. Finally, to verify the mission goal of radiation closure, a broadband radiometer (BBR) observes broadband radiative fluxes at three different viewing angles along the EarthCARE track.



**FIGURE 1** From launch to observation mode. Panels show (from top-left clock-wise), the launch of the Falcon 9 rocket carrying the EarthCARE payload; the EarthCARE satellite moments after being released into orbit; and a schematic showing the footprints of EarthCARE's four synergistic sensors (see Box A for an explanation of the acronyms). Bottom image credit: based on Illingworth et al. (2015), © American Meteorological Society, used with permission.





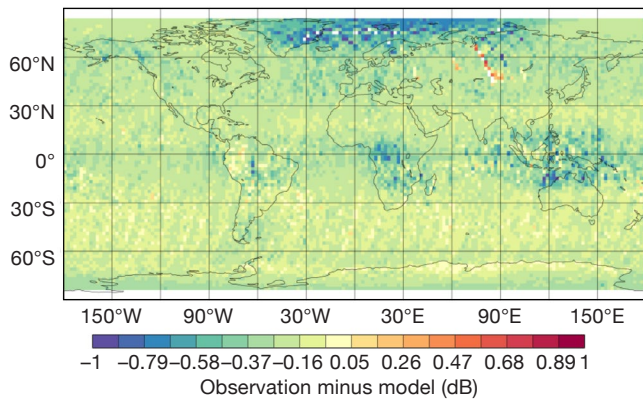
**FIGURE 2** Example EarthCARE L1 data from an overpass close to Reading (UK) on 23 August 2024. Plots show (a) ATLID particle attenuated backscatter, (b) CPR radar reflectivity, and (c) an MSI natural colour image (see Box A for details). The dashed line in the natural colour image indicates the path of the satellite.

A new forward model to map from model space to observation space was implemented. By taking a ‘triple-column’ approach, sub-grid variability in cloud cover and cloud amount is accounted for by simulating the radar and lidar signal passing between three sub-columns. For screening, observations that have large first-guess departures (large differences between short-range forecasts and observations) are discarded. This is also done where there is no cloud or precipitation in the model. There are also situations where we know our forward model does not account for multiple scattering of the radar or lidar signal, so cases where

this is suspected, such as in the cores of deep convection, are also excluded.

Another important aspect of data assimilation is quantifying observation errors, as they weight the influence of an observation. Observation errors of cloud measurements are highly regime dependent. To estimate how much weight to give to the EarthCARE observations within 4D-Var, we therefore take an ‘error inventory’ approach that adds different sources of error in quadrature. Often the sampling error resulting from the radar’s narrow footprint relative to the size of grid





**FIGURE 3** Average of ATLID co-polar Rayleigh attenuated backscatter first-guess departures, for model pressure between 0–400 hPa during January 2025.

boxes dominates. However, in some situations, such as stratiform ice cloud with relatively little turbulence, the uncertainties in certain microphysical assumptions can be larger.

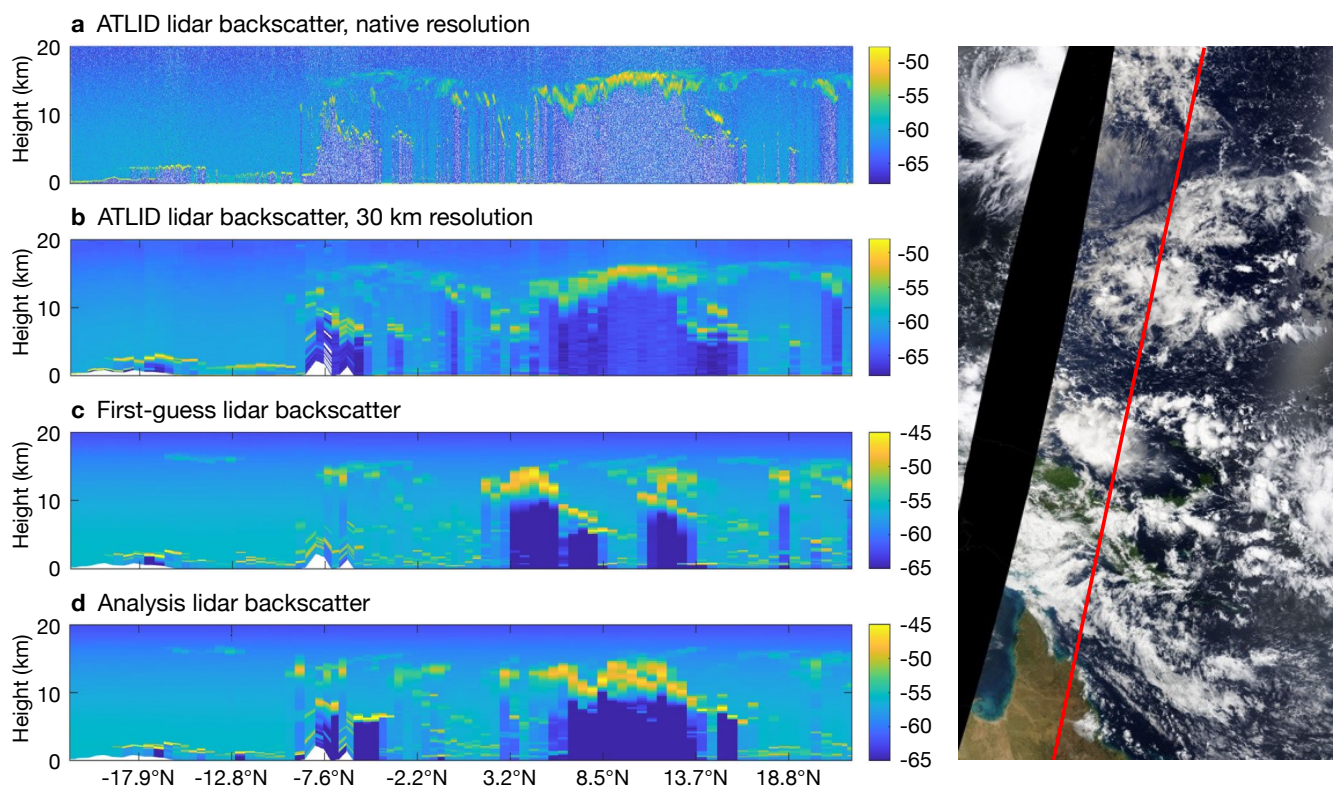
## ECMWF's role in the commissioning phase

Before incorporating new observations into ECMWF's assimilation system, the measurements must be carefully monitored for consistency and stability, while also checking for biases or artifacts that could degrade

the analysis. Observation monitoring at ECMWF is also invaluable for the calibration and validation of satellite instruments. Comparing observational and model data significantly enhances the detection of sudden jumps or drifts in bias.

For example, ECMWF's monitoring was instrumental in identifying and correcting mirror-temperature-dependent biases in ESA's seventh Earth Explorer, Aeolus. To support EarthCARE's calibration and validation activities, ECMWF began routine monitoring of CPR and ATLID data within days of the instruments being activated.

During EarthCARE's commissioning phase, NWP monitoring was particularly useful for tracking the radar calibration. Jumps in calibration due to changes in the signal processing were quickly reported to ESA and JAXA, which proved invaluable for helping engineers to troubleshoot problems. Traditional methods to calibrate space-borne radar rely on either direct overpasses with calibrated sensors on the ground, or a statistical analysis of sea-surface returns, both of which require many days to complete. Averaged statistics using ECMWF's OBSTAT plotting software (Figure 3) have also been useful for monitoring the radar and lidar observations. In Figure 3, the mean first-guess departures for Rayleigh backscatter reveal some contamination with spurious signals over



**FIGURE 4** Example for the assimilation of ATLID total backscatter on 3 August 2024, with EarthCARE passing the equator at 145°E. Plots show (a) ATLID total lidar backscatter at native resolution, (b) ATLID total lidar backscatter at 30 km horizontal resolution, (c) first-guess total lidar backscatter, and (d) analysis total lidar backscatter. Backscatter is shown in units of  $10 \log_{10} (\text{m}^{-1} \text{sr}^{-1})$ . The red line in the satellite image shows the path of the satellite.

Russia, the presence of Arctic polar stratospheric clouds not represented in the model, and upper tropospheric aerosols in the tropics (a negative bias due to the lack of attenuation in the model).

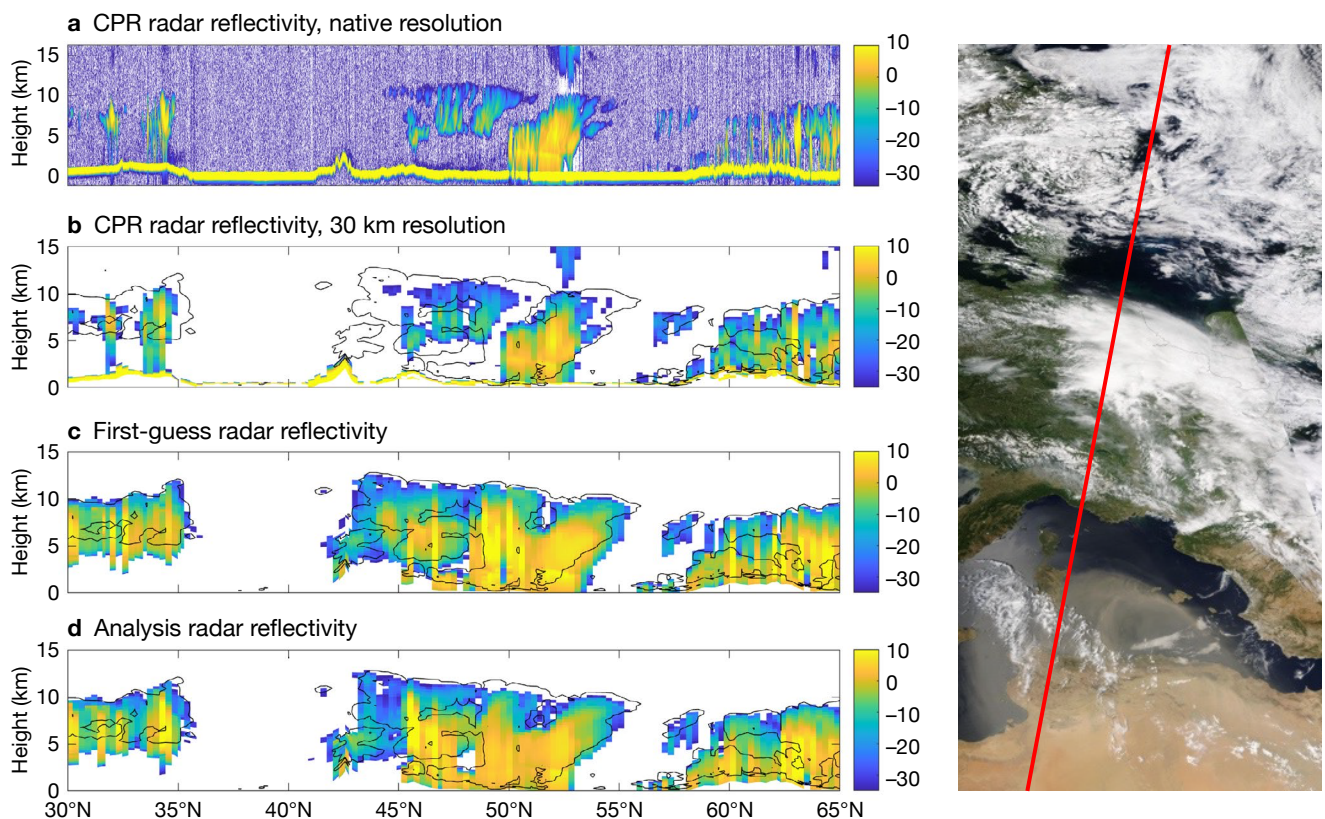
## Early assimilation tests and impact on forecasts

Preparations for radar and lidar assimilation during the joint ESA–ECMWF ‘PEARL Cloud’ project enabled EarthCARE assimilation experiments to begin soon after launch. Figure 4 shows an example of the assimilation of total attenuated backscatter for a section of an EarthCARE orbit over the tropical West Pacific, beginning over Australia and then passing Papua New Guinea. The frame is dominated by high ice cloud near the tropopause associated with the outflow from convection and a broad area of active deep convection at 10°N, which fully attenuates the lidar signal. Before assimilation (in the first-guess panel), much of the active convection in the model is displaced compared to the observations. However, after assimilation (in the analysis panel), the model represents the cloud positions much better. The mid-level convection over Papua New Guinea is also better represented.

A demonstration of the assimilation of CPR radar reflectivity is shown in Figure 5. In this case, EarthCARE

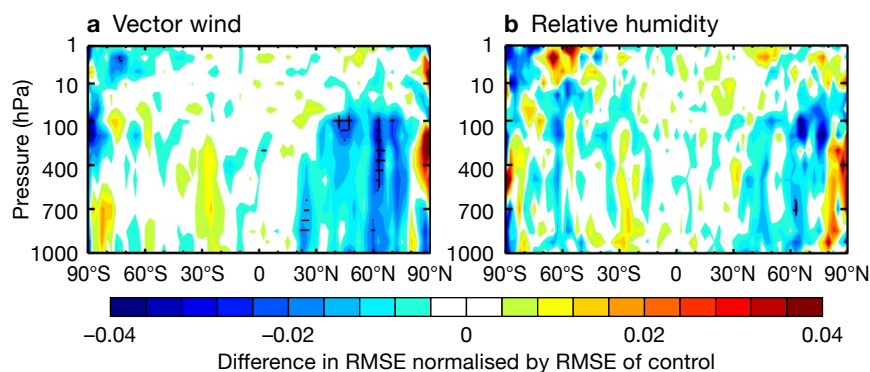
passed over Libya before traversing central Europe and Sweden. An occluded front lay across much of France and Germany, as seen by the heavy precipitation in the radar around 50°N. Assimilating the observations helps move the front further south in the analysis. Cloud amount in the southern part of the front is also improved by the assimilation.

In addition to assimilation case studies, the impact of assimilating EarthCARE on forecast skill has begun to be investigated. Two three-month cycling analysis experiments between August and October 2024 were performed, one where all regularly assimilated observations were used, and one where EarthCARE radar reflectivity and lidar attenuated backscatter observations were also included. Figure 6 shows that the change in root-mean-square error of vector wind relative to the control is improved in the northern hemisphere when assimilating EarthCARE, with some statistical significance beginning to appear at a forecast lead time of five days. Elsewhere, the impact is generally neutral, but these results are extremely encouraging for a new sensor, and further testing and tuning of observation errors and screening is likely to improve the impact. Longer experiments are also required to reach statistical significance (typically six-month experiments are required to detect significant impacts for single sensors).




**FIGURE 5** Example for the assimilation of CPR radar reflectivity (dBZ), with EarthCARE crossing central Europe at approximately 10°E. Plots show (a) CPR radar reflectivity at native resolution, (b) CPR radar reflectivity at 30 km horizontal resolution, (c) first-guess radar reflectivity, and (d) analysis radar reflectivity. The red line in the satellite image shows the path of the satellite.






**FIGURE 6** Normalised change in the root-mean-square error (RMSE) of five-day forecasts, measured against own analysis, of (a) vector wind and (b) relative humidity when assimilating EarthCARE observations in addition to operationally assimilated observations, relative to control experiments without EarthCARE observations. The experiments were performed on forecasts from 8 August 2024 to 31 October 2024. Cross-hatching indicates 95% confidence.

**a** ATLID backscatter on CAMS and simulated imagery



**b** ATLID backscatter on GOES-E composite RGB image



**FIGURE 7** Example of EarthCARE ATLID lidar backscatter overlaid on (a) the CAMS aerosol optical depth 18-hour forecast and simulated visible imagery from the ECMWF 4.4 km Destination Earth (DestinE) model, and (b) a GOES-E composite RGB image. According to EarthCARE, the smoke plumes are more extensive than in the CAMS forecast.



## The future is bright

EarthCARE data will be included in the operational data assimilation as soon as possible, subject to the usual checks for including new observations. If all goes well, our forecasts will begin directly benefiting from EarthCARE this year. However, the use of EarthCARE at ECMWF will not end there. In a similar way to the cloud observation, aerosol signals from ATLID are planned to be included in operations of the EU's Copernicus Atmosphere Monitoring Service (CAMS), implemented by ECMWF, within the next year. The HSRL's capability to directly measure layer-by-layer aerosol radiative properties will offer a paradigm shift in constraining the vertical distribution of aerosols, as demonstrated in Figure 7.

While this article has focused on the data assimilation of EarthCARE observations, the greatest impact is likely to come from using these observations to improve the physical processes within models. To assist with this, a unified retrieval to simultaneously retrieve aerosol, cloud and precipitation properties using three EarthCARE sensors has been developed at ECMWF. The products from this retrieval will be used by scientists in ECMWF

Member States and worldwide to evaluate and improve NWP and climate models. The ESA-funded EarthCARE DISC (Data Innovation and Science Cluster) brings together scientists from many European institutes, including the Dutch national meteorological service (KNMI), the German Aerospace Center (DLR) and the National Observatory of Athens (NOA), to monitor, validate and improve this and other retrieval products.

For ECMWF, EarthCARE is perfectly aligned with our new Strategy 2025–2034, which aims for advances in km-scale modelling and Earth system development. The high-resolution observations of updrafts in convective cores have potential for evaluating the next generation of km-scale models. Using the Doppler radar, the parametrizations of ice and snow fall speeds in our model can be constrained on a global scale for the first time. The myriad interactions between aerosols and clouds can be elucidated, regional radiation biases deciphered, and drizzle evaporation rates quantified. All this will help drive down systematic model errors for improved forecasts. In short, watch this space for further exciting science from EarthCARE.

## Further reading

**Fielding M. D. & M. Janisková**, 2020: Direct 4D-Var assimilation of space-borne cloud radar reflectivity and lidar backscatter. Part I: Observation operator and implementation, *Q J R Meteorol Soc.*, **146**, 3877–3899. <https://doi.org/10.1002/qj.3878>

**Illingworth, A.J., H.W. Barker, A. Beljaars, M. Ceccaldi, H. Chepfer, N. Clerbaux et al.**, 2015: The EarthCARE Satellite: The Next Step Forward in Global Measurements of Clouds, Aerosols, Precipitation, and Radiation, *Bull. Amer. Meteor. Soc.*, **96**, 1311–1332, <https://doi.org/10.1175/BAMS-D-12-00227.1>

**Janisková, M. & M.D. Fielding**, 2020: Progress towards assimilating cloud radar and lidar observations, *ECMWF Newsletter No. 162*, 20–26. <https://doi.org/10.21957/72xdt8mj03>

**Janisková, M. & M.D. Fielding**, 2020: Direct 4D-Var assimilation of space-borne cloud radar and lidar observations. Part II: Impact on analysis and subsequent forecast. *Q J R Meteorol Soc.* **146**, 3900–3916. <https://doi.org/10.1002/qj.3879>

**Mason, S.L., R.J. Hogan, A. Bozzo & N.L. Pounder**, 2023: A unified synergistic retrieval of clouds, aerosols, and precipitation from EarthCARE: the ACM-CAP product, *Atmos. Meas. Tech.*, **16**, 3459–3486. <https://doi.org/10.5194/amt-16-3459-2023>

**Wehr, T., T. Kubota, G. Tzeremes, K. Wallace, H. Nakatsuka, Y. Ohno et al.**, 2023: The EarthCARE mission – science and system overview, *Atmos. Meas. Tech.*, **16**, 3581–3608. <https://doi.org/10.5194/amt-16-3581-2023>

## Reintroducing the analysis of humidity in the stratosphere

Nouredine Semane, Massimo Bonavita

**A**n analysis of humidity in the stratosphere to help determine the initial conditions of weather forecasts was introduced in 1999 but removed again a few months later because of excessively high stratospheric humidity, which resulted in significant systematic forecast errors. We have now developed a new way of analysing humidity in the stratosphere, which improves forecasts. It will be introduced in the second half of this year in Cycle 50r1 of ECMWF's Integrated Forecasting System (IFS). The work described here is part of the Copernicus Atmosphere Monitoring Service (CAMS) EvOluTion project (CAMEO) funded by the EU.

It is widely recognised that properly initialising the humidity field in numerical weather prediction is crucial for a realistic representation of the hydrological cycle, convection, precipitation, and clouds. Additionally, it is important to note that water vapour is a long-lived tracer in the stratosphere, influencing atmospheric radiative processes and directly affecting the stratospheric temperature profile, as well as the shape and strength of the tropopause. Consequently, an accurate analysis of the stratosphere's humidity field is important for controlling and reducing systematic errors in the IFS. However, experience at ECMWF illustrates that analysing humidity in the stratosphere can present challenges. The ECMWF operational analyses produced from 12 October 1999 to April 2000 using IFS Cycle 21r4, which permitted stratospheric humidity increments, revealed increased moisture in the lower stratosphere at high and mid-latitudes, leading to significant systematic forecast errors (Jakob et al., 2000). To address this systematic model drift, stratospheric humidity increments were disabled in the subsequent IFS Cycle 22r1 by zeroing background errors above the tropopause. Essentially, the stratospheric humidity analysis is only derived from the short-range forecast of the preceding analysis, except in regions of supersaturation caused by a temperature drop from the analysis. As a result, the stratospheric humidity primarily evolves according to the model's dynamics and the parametrization of physical processes, remaining largely independent of observations.

Bland et al. (2021) examined the biases in IFS forecasts for the lower stratosphere and their underlying causes. Their research reveals that humidity levels in the

lowermost stratosphere are significantly overestimated: they reach up to 150% of the observed values compared to various independent data sources, including radiosondes, satellite limb sounders, lidar, and aircraft measurements. The moist bias in the lower stratosphere is also identified as the root cause of an increasing cold bias in temperature forecasts, with a cooling rate of  $-0.2$  K per day. This cold bias results from excessive longwave radiative cooling due to the moist bias. These systematic errors are also discussed in Polichtchouk et al. (2021). This article describes the investigation into reintroducing the humidity analysis in the stratosphere and how this change addresses the warm and cold bias issues in IFS forecasts for the lower stratosphere.

### Humidity assimilation in 4D-Var

Starting with IFS Cycle 26r3 (October 2003), the ECMWF 4D-Var assimilation system incorporates a pseudo-relative humidity control variable, defined as the humidity mixing ratio scaled by the saturation mixing ratio of the background. The primary advantage of this approach is that the error statistics for this control variable closely resemble those of a Gaussian distribution and are, therefore, more suitable for 4D-Var. The 4D-Var system adjusts the values of pseudo-relative humidity using various observations, which mainly concern the troposphere. These include radiosonde humidity profiles, GNSS-RO bending angles, and radiances from satellite microwave and infrared sounders. Radiosonde humidity data are not assimilated above the model-diagnosed tropopause (typically between 300 hPa and 100 hPa; see box for an explanation). Despite the absence of radiosonde humidity data above the tropopause, 4D-Var is still

**a**

#### The tropopause

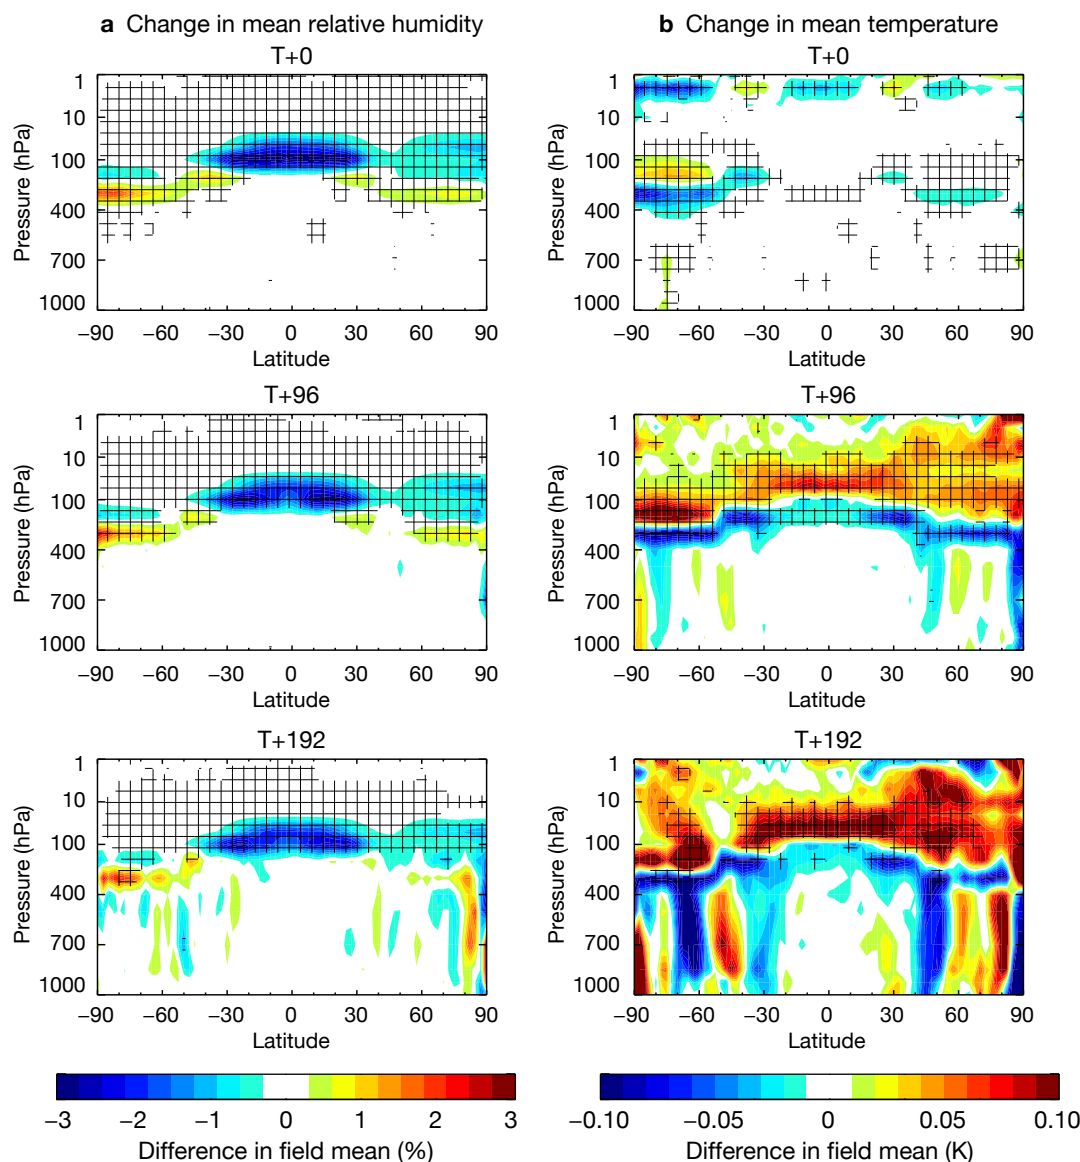
In addition to its association with a minimum temperature and a sharp change in static stability, the tropopause level (hygropause) is linked to the model level at which the specific humidity exceeds  $3 \text{ mg kg}^{-1}$  and the specific humidity two model levels below is greater than  $5 \text{ mg kg}^{-1}$ , when examining model levels from 70 hPa down to 500 hPa.

able to produce useful analysis increments in the lower stratosphere due to the spatial structure of the vertical correlations and the standard deviation of the background errors used in the analysis update. Additionally, humidity-sensitive channels from satellite-based radiometers, which peak in sensitivity within the upper troposphere, often exhibit an extended tail of sensitivity that reaches into the stratosphere, where humidity values are two orders of magnitude smaller than those in the troposphere. In summary, multiple pathways within the 4D-Var analysis can alter the lower stratospheric humidity field, even when it is not directly observed.

So, what has changed since the mid-2000s that has made the reintroduction of stratospheric humidity analysis feasible? One factor is that 4D-Var has benefited from the introduction of ensemble-based techniques for uncertainty estimation through the Ensemble of Data Assimilations (EDA); progressive improvements in the fidelity of the IFS model, including refinements to its

physical parametrizations; and concomitant increases in both vertical and horizontal resolution. EDA-derived covariances have drastically improved the situation compared to the mid-2000s, as they provide sharper and more localised error estimates than the previously used method (see Bonavita et al., 2016). Furthermore, the increase in the number of vertical levels in the IFS (60 in IFS Cycle 21r4 versus 137 in the current IFS Cycle 49r1), with a more than proportional increase around the tropopause, has also enhanced the data assimilation (DA) system's ability to represent sharp error structures, thus yielding more localised analysis increments.

The combination of these enhancements in DA and the model is already sufficient, as demonstrated in the next section, to significantly improve the situation regarding the systematic errors of the analysis and forecast in the lower stratosphere. As anticipated, incorporating additional humidity observations with sufficient vertical resolution to resolve the tropopause enhances the accuracy and fidelity of the resulting analyses.



**FIGURE 1** Change in (a) the forecast mean relative humidity and (b) the forecast mean temperature in a stratospheric humidity analysis experiment compared to the control over two and a half months of experimentation (13 December 2020 to 28 February 2021) at forecast times of 0, 96 and 192 hours. Areas marked with cross-hatching are statistically significant at the 95% level.

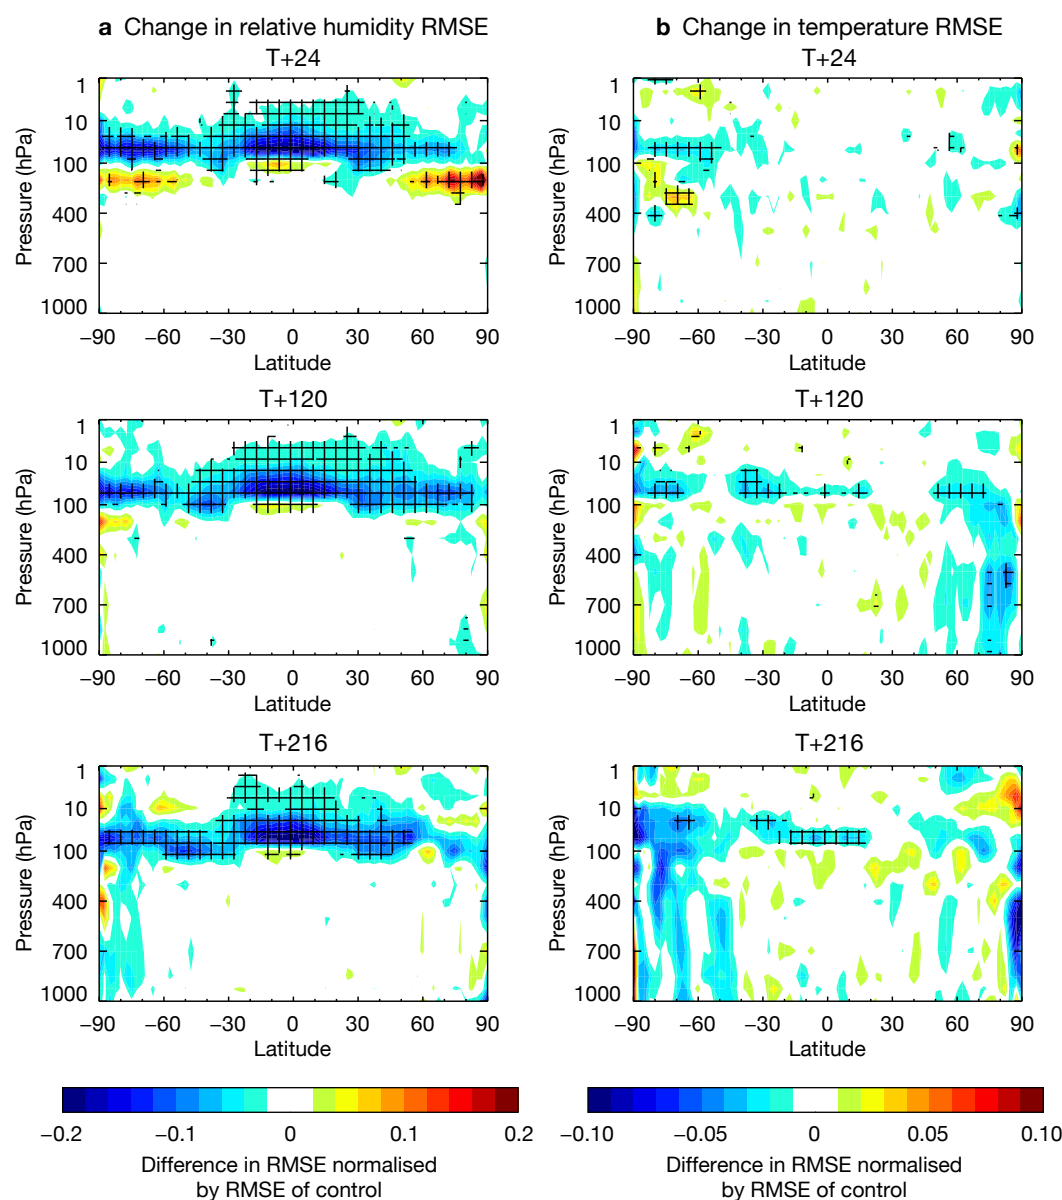


## Humidity data assimilation in the stratosphere

Beyond retaining the stratospheric humidity increments produced by 4D-Var, we assimilate two additional types of observations sensitive to stratospheric humidity in the experiments discussed here. One source is the EOS-Aura Microwave Limb Sounder (MLS) moisture retrievals, which have a vertical resolution of approximately 3–5 km. The second additional source of stratospheric humidity information involves extending the vertical usage of the humidity observations from RS41-type sondes up to 60 hPa, as they are recognised for accurately measuring humidity across a wide range of atmospheric conditions. This will provide relevant measurements even when MLS data are no longer available, which is anticipated for mid-2026. Because of the limited availability of MLS data, they will not be assimilated in IFS Cycle 50r1, but the benefits of

enabling stratospheric humidity increments are preserved even without MLS data.

Figure 1 presents zonal average plots of the mean forecast change in relative humidity (left) and temperature (right) between an experiment incorporating stratospheric humidity analysis using MLS and sonde humidity observations and a control. The primary direct effect of assimilation is the significant drying of the lower stratosphere, as illustrated on the left-hand side of Figure 1, alongside the corresponding moistening below the tropopause. Both effects mitigate the known systematic deficiencies of the IFS. A secondary, more indirect effect is observed on the right-hand side of Figure 1. Due to radiative effects, the temperature of the lower stratosphere becomes warmer, while the area under the tropopause cools, with this effect growing as the forecast time increases. Again, these changes



**FIGURE 2** Normalised change in root-mean-square error (RMSE), in the stratospheric humidity analysis experiment compared to the control, of (a) relative humidity forecasts and (b) temperature forecasts, verified against own analysis, over three months of experimentation (3 December 2020 to 27 February 2021), at forecast times of 24, 120 and 216 hours. Areas marked with cross-hatching are statistically significant at the 95% level. Blue cross-hatched areas indicate a significant reduction in forecast error and hence an improvement in forecast quality.

significantly reduce the known IFS analysis and forecast biases. Additionally, the changes in the humidity field and those in the temperature field strengthen and sharpen the tropopause. This is also a step in the right direction, as a weak and diffuse tropopause is another well-documented systematic deficiency of ECMWF analyses and forecasts (see for example Krüger et al., 2024). Figure 2 displays the normalised change in forecast root-mean-square error (RMSE) of relative humidity (left-hand side) and temperature (right-hand side) from the stratospheric humidity analysis experiment, compared to the control and verified against its own analysis. Reintroducing the analysis of stratospheric humidity evidently improves forecasts across all lead times. The enhancements are more pronounced in the Upper Troposphere Lower Stratosphere (UTLS) layer, but there are indications that they propagate into the troposphere with increasing lead time. This is very encouraging because the interaction between tropospheric and stratospheric dynamics is intermittent, and these results suggest that the improved characterisation of the tropopause may lead to a better representation of the troposphere–stratosphere interaction.

## Future directions

The results presented here clearly demonstrate the benefits of reintroducing the analysis of stratospheric humidity in the ECMWF operational assimilation cycle after a hiatus of over 20 years. This will be one of the main contributions to the upcoming IFS Cycle 50r1. They also highlight the importance of limb sounding observations that are sensitive to humidity and possess sufficient vertical resolution to resolve the tropopause, such as the proposed Changing-Atmosphere Infra-Red Tomography Explorer (CAIRT). If implemented, with high vertical resolution and the capability to measure key atmospheric components like water vapour, ozone, and aerosols, CAIRT would provide invaluable data to enhance climate models, improve weather forecasting, and deepen our understanding of atmospheric processes in the UTLS.

Humidity, like other trace gases, functions as a long-lived tracer in the stratosphere. Accurate initialisation directly influences extended-range predictability through wind-tracing effects and indirectly through radiative effects in the model and the observation operators of nadir-sounding satellite instruments. Therefore,

enhancing the accuracy of initial conditions for trace gases in the stratosphere can substantially improve forecast skill in the medium and extended range. Current efforts aim to unlock this potential in future IFS operational cycles.

*The CAMEO project (grant agreement No 101082125) is funded by the European Union. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the Commission. Neither the European Union nor the granting authority can be held responsible for them.*



**Funded by the  
European Union**

## Further reading

**Bland, J., S. Gray, J. Methven & R. Forbes**, 2021: Characterising extratropical near-tropopause analysis humidity biases and their radiative effects on temperature forecasts. *Q.J.R. Meteorol. Soc.*, **140**, 3878–3898. <https://doi.org/10.1002/qj.4150>

**Bonavita, M., E. Hólm, L. Isaksen & M. Fisher**, 2016: The evolution of the ECMWF hybrid data assimilation system. *Q.J.R. Meteorol. Soc.*, **142**, 287–303. <https://doi.org/10.1002/qj.2652>

**Jakob, C., E. Andersson, A. Beljaars, R. Buizza, M. Fisher, E. Gerard et al.**, 2000: The IFS cycle CY21R4 made operational in October 1999. *ECMWF Newsletter* **No. 87**, 2–9. <https://www.ecmwf.int/en/elibrary/78244-newsletter-no-87-spring-2000>

**Krüger, K., A. Schäfler, M. Weissmann & G.C. Craig**, 2024: Influence of radiosonde observations on the sharpness and altitude of the midlatitude tropopause in the ECMWF IFS, *Weather Clim. Dynam.*, **5**, 491–509. <https://doi.org/10.5194/wcd-5-491-2024>

**Polichtchouk, I., P. Bechtold, M. Bonavita, R. Forbes, S. Healy, R. Hogan et al.**, 2021: Stratospheric Modelling and Assimilation. *ECMWF Technical Memorandum*, **No. 877**. <https://doi.org/10.21957/25hegfoq>

## ECMWF contributes to Swiss supercomputing project in weather and climate science

Thomas Geenen, Ahmed Benallegue, James Hawkes, Simon Smart, Tiago Quintino, Umberto Modigliani (all ECMWF), Anurag Dipankar (ETH Zurich), Oliver Fuhrer (MeteoSwiss)

**T**he SwissTwins project, led by the Swiss National Supercomputing Centre (CSCS), started in December 2022 and was established to support the Swiss weather and climate science community (<https://www.cscs.ch/about/collaborations/swisstwins>). One of the primary objectives is to ensure that Swiss research infrastructures remain well integrated in Europe. As part of this wider initiative, ECMWF has collaborated closely with CSCS and MeteoSwiss to establish efficient access to ECMWF's forecast products and meteorological archive for the training of artificial intelligence (AI)/machine learning (ML) models. These models have already shown great potential in weather forecasting and are increasingly being explored for climate modelling.

### Bridging data and computational power

Earth system sciences and the study of climate change require substantial computational power and access to large volumes of high-quality data. This is especially true for ML models and particularly for emerging trends in ML models trained on vast datasets, called foundation models. Compared to task-specific ML tools, foundation models need new training methods, and they require significantly more complex neural network architectures that are trained from massive datasets. Therefore, their development needs substantially more resources than that of task-specific models.

This understanding triggered the initial design concepts for the SwissTwins project and resulted in the design of the DataHypercube. The DataHypercube is a high-performance storage cluster that is deployed in ECMWF's data centre in Bologna (Italy), close to the datasets that are key to train both task-specific ML models as well as foundation models. ECMWF has one of the largest meteorological archives in the world and provides also other datasets widely used for ML model training, such as ERA5 meteorological reanalysis data. The DataHypercube is connected through an ultra-large bandwidth network connection to the Alps supercomputing system of the CSCS in Lugano (Switzerland). Alps hosts the MeteoSwiss weather prediction model and is also used by groups from ETH Zurich for Earth system sciences and climate change

research. Leveraging ECMWF data handling software services, deployed on the DataHypercube, specific data can be extracted, pre-processed and staged for training purposes. The Alps supercomputer is one of the fastest AI supercomputers in Europe. Connecting the AI and ML capability of the Alps infrastructure with the unique data products and archive hosted by ECMWF in Bologna will open up unprecedented research opportunities. This project pilots a proof of concept, informing an outline methodology for possible future initiatives with other ECMWF Member States to leverage ECMWF data and services in a similar fashion.

### Enabling AI/ML for weather and climate research

A significant part of the project is to set up the DataHypercube system in Bologna and to establish the high-performance network connection between the DataHypercube system and Alps in Lugano, but there is more. To allow for data access with low latency and the extraction of specific datasets to train AI/ML models for the benefits of the Swiss and international numerical weather and climate prediction community, the deployment of key ECMWF data handling services on this system is essential. ECMWF has been developing these tools and services over the last few years and further evolved them in the context of the EU's Destination Earth initiative (<https://destine.ecmwf.int>). The architecture of the software and services deployed on the DataHypercube will closely follow the developments of the Digital Twin Engine of Destination Earth and the architecture of the Destination Earth data bridges. In that sense, it is yet another example of using the ECMWF Software Engine (López Alós, 2024). In particular, it will leverage the Fields Database (FDB – Smart et al., 2017) for data storage and the Polytope service (Hawkes et al., 2020) for efficient data delivery. It uses the GribJump functionality to efficiently locate data in larger GRIB files (<https://github.com/ecmwf/gribjump>). FDB is ECMWF's semantic object store technology designed for efficient data handling of hierarchical, high-dimensional datacubes. FDB will be distributed across the cluster, ready to ingest data from ECMWF's forecasts. Its high-bandwidth communication protocols can be used directly for bulk consumers of this data, including by AI/ML models.



Polytope is ECMWF's data delivery service with the unique ability to efficiently extract features, such as point data (e.g. time-series, vertical profiles) or area data (e.g. polygon extractions), from the vast FDB data store without any intermediate copies or transposition of the data.

This will bring a new level of accessibility to the data, accelerating the work of Earth system and climate scientists. On Alps, an additional ECMWF data handling software component called earthkit will be deployed (Russell et al., 2024). This is to provide users with a rich API to extract and process the data from the DataHypercube on Alps.

The combination of a dedicated high-performance storage cluster with a high-performance network connection on Alps and the deployment of advanced data handling capabilities makes this a potent tool in the hands of weather and climate scientists (Figure 1). ETH Zurich's exascale simulation platform, EXCLAIM, will be benefiting from faster on-site access to ECMWF data through Alps. This improvement supports use cases on EXCLAIM, which aim to understand the evolving climate. These include kilometre-scale global climate simulations and limited-area large-eddy simulations.

The climate studies rely on various ECMWF products, such as reanalysis, daily operational analysis, and forecasts, to provide the necessary initial and boundary conditions for the modelling system used at ETH Zurich. These datasets are crucial for model validation and verification.

Enhancing the data cube on Alps with advanced tools like

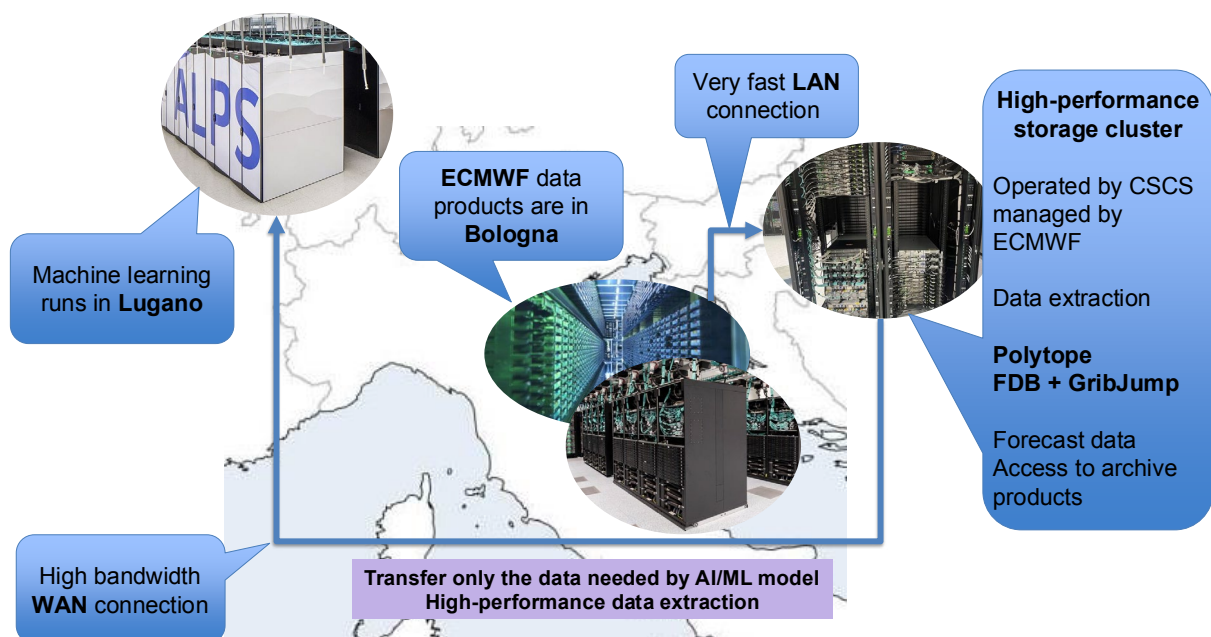
Polytope and earthkit is a big plus. These tools enable faster and more selective access to relevant data, which is expected to significantly improve user productivity.

## A dedicated multi-domain service

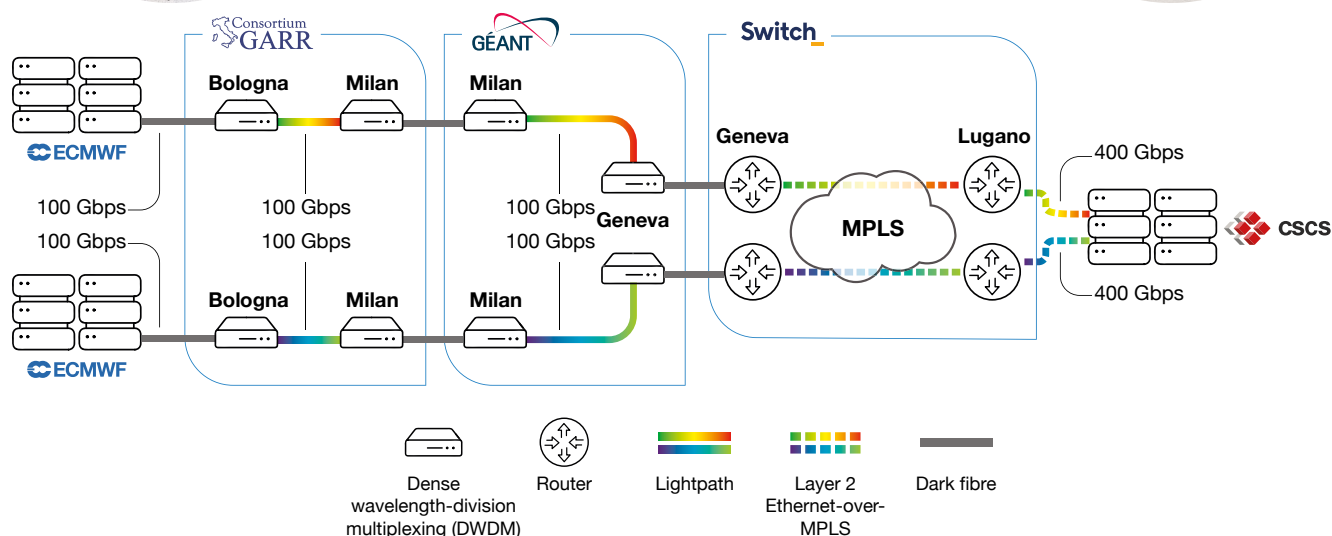
An end-to-end, ultra-large bandwidth network infrastructure is required to enable massive data transfers from ECMWF's meteorological data products and archives to CSCS supercomputing facilities. It is foreseen that during the ramp-up period and subsequent phases of the project, multiple AI/ML models in the weather and climate domain will be consuming data from the DataHypercube simultaneously, requiring a further increase of the total aggregate throughput as the number of users grows.

Thanks to collaboration between GARR (Gruppo per l'Armonizzazione delle Reti della Ricerca, the Italian national computer network for research and education), SWITCH (the Swiss national research and education network) and GÉANT (Gigabit European Academic Network, the pan-European data network for the research and education community), the two centres are now interconnected by a state-of-the-art, ultra-large bandwidth, dedicated connection that seamlessly traverses the three network domains. In designing this link, the requirements included not only extremely high capacity and scalability but also significant resilience, reliability, security, and direct user access to the optical infrastructure.

This dedicated connection comprises two separate 100 Gbps links, independently managed by the three



**FIGURE 1** This diagram illustrates how the high-performance storage cluster and network will provide specific datasets from the ECMWF data holdings towards the Alps system in Lugano to be used by the Swiss weather and climate science community. (Copyright of the Alps photo: Marco Abram, CSCS)



**FIGURE 2** A resilient, fast, dedicated, secure connection has been created between ECMWF and CSCS. This diagram shows the details of the high-performance network connection between the ECMWF data holdings in Bologna and CSCS Alps in Lugano. See <https://www.garrnews.it/caffe-scientifico/tre-reti-al-servizio-della-terra>. (Copyright of the Alps photo: Marco Abram, CSCS)

operators. To ensure the security and reliability of the service, new resources have been deployed. At the ECMWF site in Bologna, two new fibre links have been installed by Lepida, GARR's local partner in the Emilia-Romagna region, to access the backbone. These complement the existing GARR fibre connections that are already used for operational forecast product dissemination at ECMWF.

From there, two optical light paths have been configured to reach GARR PoPs (Points of Presence) in Milan. Here, the Italian network connects with the GÉANT network, which carries the traffic to its PoP located at CERN in Geneva. From Geneva, SWITCH takes over and delivers the data to its final destination at CSCS in Lugano.

The implementation methods differ between the networks: GARR and GÉANT have chosen an optical domain approach, while SWITCH uses a Layer 2 Ethernet-over-MPLS (Multi-Protocol Label Switching) service (Figure 2).

"This type of service needs a high level of collaboration and trust among operators and would not be possible outside the research networks ecosystem," said Sabrina Tomassini (Head of Network Department at GARR). When it comes to implementation, management, and

monitoring, each network is responsible for its domain, therefore a high level of communication is necessary. National Research and Education Networks (NRENs), such as GARR, GÉANT and SWITCH, offer multi-domain services, transcending their domains to harmonise services across networks. This capability is unique to NRENs, as commercial operators are confined to their domains of competence.

To ensure a high level of security and resilience, a deep technical analysis was conducted, and diversification, shelters, and robust pathways were employed to identify and address potential faults, thereby ensuring the highest possible resilience of the service.

## In Bologna as in Lugano

"This service is one of the first to leverage the capabilities of the new next-generation GARR-T network for optical infrastructure access," stated Sabrina Tomassini. "The advantage is that the entire 100 Gbps channel capacity is wholly dedicated to the user, without statistical multiplexing."

"Moreover, a peculiarity of this interconnection is that the sites are linked at the layer 2 level. This enables CSCS to have full visibility of the DataHypercube machines located

at ECMWF, as if they were located in its local network in Lugano. This is a feature enabled by the new GARR-T network,” said Alessandro Inzerilli (head of the Network Operations Centre and Operations groups at GARR).

The DataHypercube platform itself also plays an important role in enhancing the data-proximity experience for the users on Alps. The interactive, user-facing data handling services that serve time-critical data are hosted from a Kubernetes cluster (a three-node cluster) with access to a fast flash storage pool of 500 TiB (11 storage nodes), which hosts the FDB services. In addition, the DataHypercube contains a second tier of capacity storage (2 PiB ClusterStor Lustre appliance) for specific historical datasets that can be pushed to the system based on user request. The DataHypercube hardware specifications are provided in Table 1.

## Going forward

The project, currently planned for two years, serves as a pilot. If successful, it could lead to further developments, including the possibility of increasing capacity to 400 Gbps in the future. One of the advantages of the optical link with GARR is that any potential upgrade to 400 Gbps could be performed relatively quickly with minimal technical effort from ECMWF.

This high-performance connection, together with the deployment of a high-performance storage cluster and data handling services both on Alps and in the Bologna data centre, opens new opportunities for research in meteorology and scientific computing, enabling closer interaction between ECMWF and the Swiss numerical weather and climate prediction community. It facilitates the sharing of large volumes of data and a tight integration of analysis software using earthkit on Alps and backend services like polytope and FDB, which are deployed on the DataHypercube in Bologna. The project potentially accelerates progress in climate modelling and weather forecasting and is an example of close collaboration with our Member States and international organisations to generate value through innovation.

Hardware component	Number	Specifications
Compute and login nodes	3	2x AMD EPYC 7713 64c 512 GB RAM 30 TB local SSD
Storage/compute nodes	11	1x AMD EPYC 9534 64c 384 GB RAM 61 TB SSD
ClusterStor appliance (Lustre)	1	HPE Cray E1000 2 PB usable space
High speed network switch	6	200 Gbit/s HPE Cray Slingshot-11
External network switch	2	100 Gbit/s Aruba router

**TABLE 1** The hardware specification of the DataHypercube installed in the ECMWF data centre in Bologna.

## Further reading

**López Alós, A.**, 2024: The modernisation of the Data Stores at ECMWF, *ECMWF Newsletter No. 181*, 9–10. <https://www.ecmwf.int/en/newsletter/181/news/modernisation-data-stores-ecmwf>

**Smart, S.D., T. Quintino & B. Raoult**, 2017: A Scalable Object Store for Meteorological and Climate Data, *PASC '17*, 13, 1–8. <https://doi.org/10.1145/3093172.3093238>

**Hawkes, J., N. Manubens, E. Danovaro, J. Hanley, S. Siemen, B. Raoult et al.**, 2020: Polytope: Serving ECMWFs Big Weather Data, *EGU General Assembly 2020 session entry*. <https://doi.org/10.5194/EGUSPHERE-EGU2020-15048>

**Russell, I., T. Quintino, B. Raoult, S. Kertész, P. Maciel, J. Varndell et al.**, 2024: Introducing earthkit, *ECMWF Newsletter No. 179*, 48–53. <https://www.ecmwf.int/en/newsletter/179/computing/introducing-earthkit>



## ECMWF publications

(see [www.ecmwf.int/en/research/publications](http://www.ecmwf.int/en/research/publications))

### 50th anniversary publications

**Corporate document:** Fifty years of Earth system modelling at ECMWF. *April 2025*

**Corporate document:** Fifty years of data assimilation at ECMWF. *April 2025*

**Corporate document:** Ten years of Copernicus at ECMWF. *April 2025*

### ESA or EUMETSAT Contract Reports

**Salonen, K., P. Weston & P. de Rosnay:** Quarter 4 2024: Operations Service Report. *January 2025*

## ECMWF Calendar 2025

Apr 28	Advisory Committee for Data Policy (virtual)	Oct 6–8	Scientific Advisory Committee
Apr 28–29	Finance Committee	Oct 6–9	Training course: Use and interpretation of ECMWF products
Apr 30	Policy Advisory Committee (virtual)	Oct 20–21	Technical Advisory Committee (virtual)
May 7	AI Weather Quest: Testing period launch webinars	Oct 27–31	Training course: Machine learning for weather prediction
May 14	Webinar: OpenStack Infrastructure as Code in the European Weather Cloud	Oct 28	Finance Committee (virtual)
Jun 2–4	8th C3S General Assembly	Oct 27	Policy Advisory Committee (virtual)
Jun 16–19	Online training course: Atos HPC training	Nov 10–14	Training course: A hands-on introduction to Numerical Weather Prediction Models: Understanding and Experimenting
Jul 3–4	Council (virtual)	Dec 4–5	Council
Sep 3–5	9th CAMS General Assembly		
Sep 15–19	ECMWF's 50th anniversary events in Bologna, Italy		
	Sep 15–18 Using ECMWF's Forecasts (UEF2025)		
	Sep 15–19 21st workshop on high-performance computing in meteorology		
	Sep 16 Code for Earth 2025		
	Sep 16–17 ECMWF DestinE Annual Meeting		

## Contact information

ECMWF, Shinfield Park, Reading, RG2 9AX, UK

Telephone National 0118 949 9000

Telephone International +44 118 949 9000

ECMWF's public website [www.ecmwf.int/](http://www.ecmwf.int/)

E-mail: The e-mail address of an individual at the Centre is [firstname.lastname@ecmwf.int](mailto:firstname.lastname@ecmwf.int). For double-barrelled names use a hyphen (e.g. [j-n.name-name@ecmwf.int](mailto:j-n.name-name@ecmwf.int)).

For any query, issue or feedback, please contact ECMWF's Service Desk at [servicedesk@ecmwf.int](mailto:servicedesk@ecmwf.int). Please specify whether your query is related to forecast products, computing and archiving services, the installation of a software package, access to ECMWF data, or any other issue. The more precise you are, the more quickly we will be able to deal with your query.



**Newsletter | No. 183 | Spring 2025**

European Centre for Medium-Range Weather Forecasts

[www.ecmwf.int](http://www.ecmwf.int)