ECMWF
European Centre for Medium-Range Weather Forecasts

Newsletter

No. 186 | Winter 2026

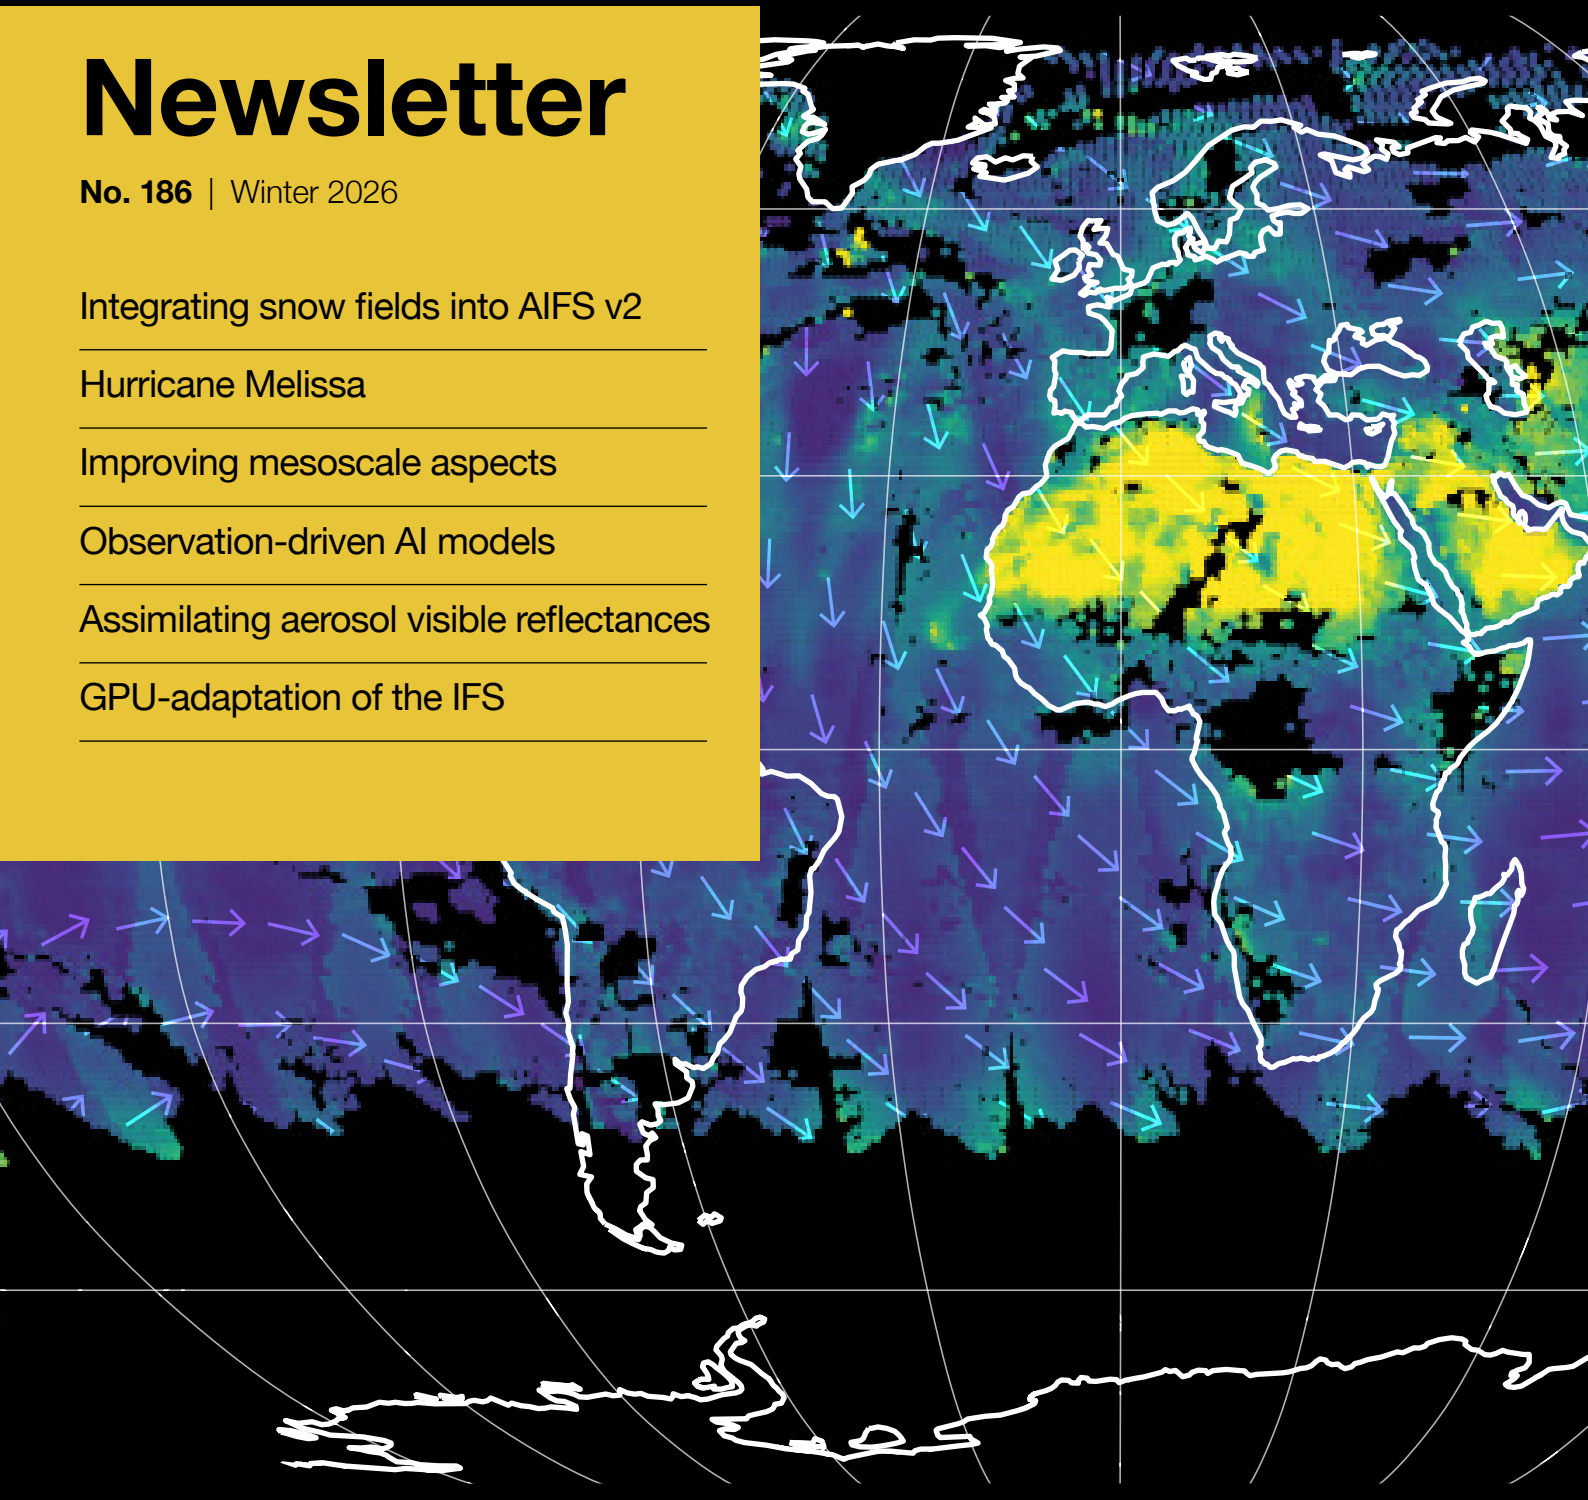

European Centre for Medium-Range Weather Forecasts, Shinfield Park, Reading, RG2 9AX, UK

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The ECMWF Newsletter is not peer-reviewed.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 Chelsea.Snell@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*

Building the foundations ahead

Taking up the role of Director-General at ECMWF at the start of this year is both a privilege and a responsibility. We enter 2026 with a strong scientific, technical and operational legacy, trusted partnerships across Europe and beyond, and a clear purpose: to deliver forecasts and information that help societies prepare for and respond to severe weather. My focus now is to build on these foundations to meet the demands of a rapidly evolving world.

This newsletter reflects a central theme for the years ahead: turning excellence across the entire forecast chain into robust, reliable and impactful services for our national hydrometeorological services (NMHSs). From observations and data assimilation, through software engineering and high-performance computing, to product delivery and user engagement, we are strengthening the pipeline end-to-end so that innovation translates into measurable benefits.

Several feature articles illustrate this clearly. Advances in medium-range forecast skill, improvements in ensemble initial conditions, and the assimilation of new observation types all contribute directly to forecast reliability. At the same time, we are adapting our Integrated Forecasting System (IFS) to next-generation high-performance computing architectures to ensure that our forecasts can be delivered efficiently and sustainably on Europe's future supercomputers. These efforts are not ends in themselves; they enable faster development cycles, higher-resolution modelling, and more timely information for users.

Machine learning (ML) is now a stable part of this landscape. Integrating ML-based systems with deep physically based knowledge – and incorporating them consistently into operations – is one of our most exciting challenges. Our work on ML-driven forecasting in co-development with our NMHSs, as well as the integration of new components such as snow fields into the Artificial Intelligence Forecasting System (AIFS), reflects our commitment to innovation that delivers impact. Our aim is not novelty for its own sake but to push the boundaries where it delivers added forecast value: enhancing skill, reliability, resilience and usability.

That extends beyond our forecasting systems. Initiatives such as ClimAfrica and the conclusion of the EO4EU project show how ECMWF supports early warning, climate resilience and decision-making, in close partnership with our Member and Co-operating States and users. Our new single access point for learning resources is another step towards making knowledge easier to find and apply.

Following ECMWF's 50th anniversary year, it is worth reflecting on what underpins our success: people and partnerships – across disciplines, institutions and borders. Looking ahead, my priority is to strengthen these collaborations, listen to user needs, and ensure our forecasts continue to deliver tangible benefits to our national hydrometeorological services. This newsletter offers a snapshot of that journey. I look forward to continuing it together.

Florian Pappenberger
Director-General

Contents

Editor Chelsea Snell • **Typesetting and graphics** Chris Matthews • **Cover** MODIS aerosol visible reflectance

doi: 10.21957/sagx-yf02

Integrating snow fields into AIFS v2

Nina Raoult, Ewan Pinnington, Gabriele Arduini, Christoph Rüdiger, Patricia de Rosnay, Matthew Chantry

Accurate forecasts of snow depth and snow cover are essential for anticipating hazards such as avalanches, transport disruption, and impacts on water resources. Snow also influences surface–atmosphere interactions, affecting water and energy fluxes and playing a key role in albedo feedbacks. As part of ongoing work to extend machine-learning (ML) capabilities within the Artificial Intelligence Forecasting System (AIFS), ECMWF has developed a new representation of snow fields for AIFS v2.

AIFS v1 became operational in 2025, using the Anemoi framework developed in collaboration with our Member States. While the first version focused primarily on atmospheric variables, work is under way to provide a more complete representation of the Earth system using machine learning. This includes integrating components such as land, hydrology, oceans, waves, and sea ice. The addition of snow in AIFS v2 is a step toward that broader objective.

Representing snow with machine learning

The physical Integrated Forecasting System (IFS) model represents snowfall as snow water equivalent (SWE) – the amount of water stored in the snowpack. From this and snow density, the model calculates snow depth and continually updates it as snow falls, melts, or sublimates. Snow cover fraction is diagnosed from depth and density, with a non-linear transition from partial to full coverage as snow accumulates.

In the AIFS, these relationships are not hard-coded. Instead, the ML system is trained to predict snow depth directly, while snow cover fraction is learnt as a diagnostic variable using snow depth and related fields. This allows the model to infer accumulation and coverage patterns from the underlying data rather than from predefined equations.

Two approaches are being explored in parallel for extending ML-based Earth system modelling. One is to integrate additional variables directly into the AIFS, training a single joint model for

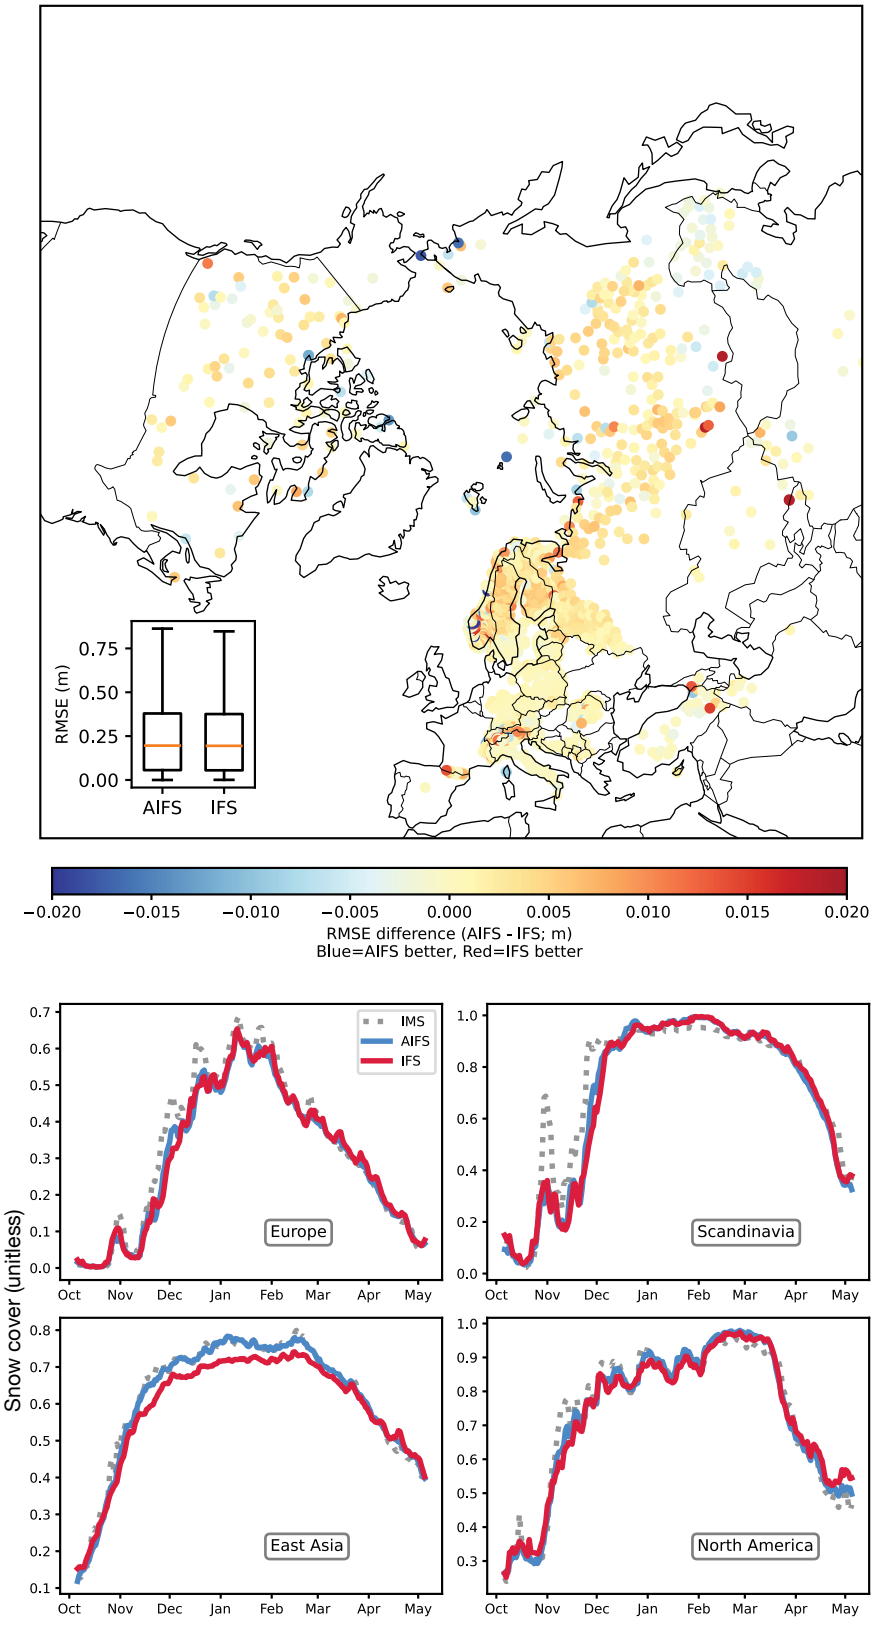

Snow depth and snow cover verification. Comparison of AIFS and IFS performance against (top) in-situ snow-depth measurements from SYNOP stations and (bottom) satellite-retrieved snow cover from the Ice Mapping System (IMS) product.

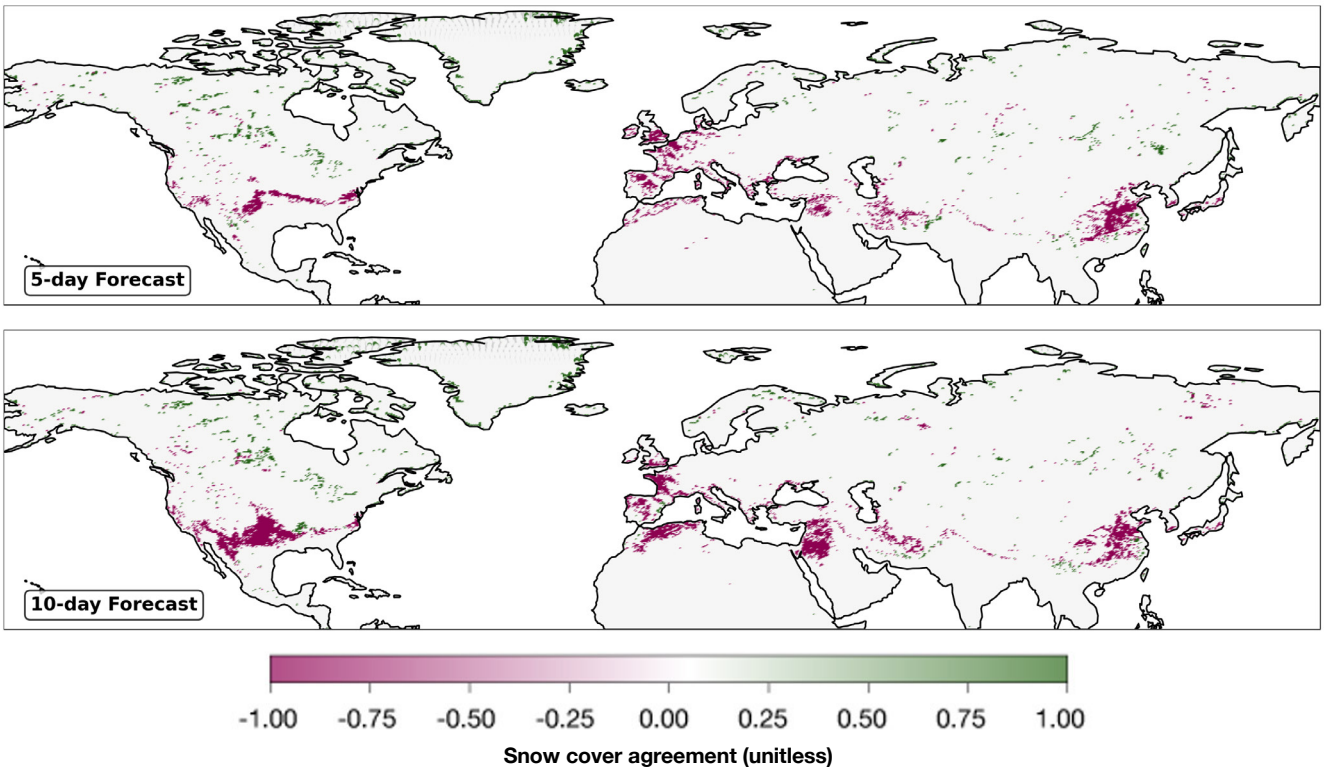

Snowline comparison in the AIFS and IFS. Five-day (top) and ten-day (bottom) forecast (FC) snow cover predicted by the IFS and AIFS, compared with IFS analysis (AN). Purple areas show where the IFS predicts snow not supported by analysis or the AIFS.

atmosphere–land–ocean–ice fields. The second is to develop separate ML components for different domains, which are then coupled, similar to traditional physics-based modelling. The work on snow in AIFS v2 follows the first approach, while the EU Destination Earth (DestinE) initiative is testing the second.

Modelling challenges

Snow processes evolve more slowly than atmospheric variables and show strong spatial heterogeneity. Permanently snow-covered regions exhibit stable conditions, whereas regions with seasonal snow cover are characterised by shallow snow depths and highly variability. These regions are the most difficult for ML systems to capture.

ML models typically rely on normalised inputs, but for snow depth this proved less effective. Better performance was achieved when the model was trained on raw values, allowing it to learn first from persistent snow in glacier regions and then adapt to the more variable behaviour of seasonal snow. This approach helped the system capture both stable and rapidly changing conditions.

Performance compared to the physical model

The AIFS has learnt to forecast snow depth and snow cover using autoregressive training, in which the model predicts future values based on its own previous output. Its performance has been evaluated against ground-based measurements and satellite-derived snow cover.

Overall, the machine-learnt model shows skill comparable to that of the physical IFS. For snow depth, the IFS retains a small advantage, with root-mean-square errors differing by less than half a centimetre. For snow cover, the AIFS shows slightly better domain-averaged performance, particularly over East Asia (see the first figure). These differences reflect, in part, the contrasting resolutions of the models: the IFS operates at around 9 km, while the AIFS currently uses a 0.25° grid (approximately 28 km), favouring regional placement over local detail.

Improved snowline representation

A known limitation in the physical model is the tendency to retain snow for too long. Operational forecasts rely on data assimilation to correct these biases. The AIFS, however, is trained on the ERA5 reanalysis and operational snow-depth analysis data, which already include assimilation corrections. As a result, the ML system can reproduce realistic snowlines, even in forecasts extending beyond the period used for training (see the second figure).

Next steps

The integration of snow fields into AIFS v2 is an important step towards representing the entire Earth system with ML. Future developments will explore the use of additional observational data sources, expand static inputs such as vegetation and soil type, and investigate whether the model can learn snow density or predict snow water equivalent directly.In parallel, work within DestinE continues to develop and couple ML components for waves, ocean, land, sea ice, and hydrology. Together, these efforts contribute to ECMWF's mission of advancing data-driven forecasting systems and to DestinE's objective of developing high-resolution digital twins to support climate adaptation and resilience.

ClimAfrica 2025: from ideas to impact in early warning and climate resilience

Katherine Egan, Karolin Eichler

ClimAfrica 2025 brought together governments, public institutions, scientists, innovators, and development partners to explore how technology and data can strengthen early warning systems and climate resilience across the continent.

Held in Rabat, Morocco, and co-organised and hosted by the General Directorate of Meteorology of Morocco, the two-day forum was designed to deliver solutions, scalable innovations, strategic partnerships and tangible political commitments to tackle the pressing climate risks intensifying across Africa.

For the first time, ClimAfrica featured an Innovation Challenge under the theme "From Ideas to Impact: Tech & Data Innovation for Climate Resilience and Early Warning in Africa". The challenge invited applicants from across Africa to present technological, digital, organisational, or data-driven solutions that help anticipate climate risks, inform decision-making, and protect lives in sectors including agriculture, water, health, energy, and urban infrastructure.

The competition reflects ClimAfrica's broader mission: to empower local innovation, elevate youth voices, and connect promising ideas with the institutions and resources needed to scale their impact. ECMWF, as an Innovation Partner through the Strengthening Early Warning in Africa (SEWA) programme, sponsored the challenge, highlighting the Centre's commitment to supporting early warning capacities and fostering Africa–Europe collaboration in climate science.

By framing the challenge around SEWA, the ClimAfrica Innovation Challenge linked innovation with the practical ambition of improving actionable early warning systems across the continent. The competition encouraged solutions that are locally adapted, scalable, and capable of integrating observations, forecasts, and climate data into decision-relevant insights. At the same time, it promoted opportunities for African innovators to interact with international experts and institutions, building networks that extend beyond the forum itself.

This first edition of the Innovation Challenge received 22 submissions across five prize categories:

- Weather and climate-induced early warning services and tools;
- Demonstration of the use of ECMWF NWP, Copernicus and/or Destination Earth (DestinE) data;
- Developments on impact-based forecast services and tools;
- Applications on the use of Cloud services, AI/ML and other digital infrastructure in scope of early warning;
- Enhancement of partnerships between African and European communities on the application of NWP data for early warning activities.

The winning projects were selected by a panel of representatives from ECMWF, the Joint Research Centre (JRC), Meteo Maroc, the Institut de Recherche pour le Développement (IRD) and the International Water Research Institute (IWRI). Each project showcased how smart, data-driven approaches can address urgent challenges and highlighted not only technical ingenuity but also a deep understanding of the societal and environmental context in which early warning systems operate.

ClimAfrica 2025 winning projects

Ground sensing – Mohamed El Garnaoui from Sultan Moulay Slimane University, Morocco

Uses a network of affordable sensors to provide real-time groundwater monitoring, helping communities anticipate water stress and protect agricultural livelihoods.

Area of interest (AOI)-based Saharan dust alerts – Mohamed Smouni from Sultan Moulay Slimane University, Morocco

Delivers personalised early warnings via mobile platforms for Saharan dust events, supporting health, agriculture, and transport decision-making.

Remote sensing for drought resilience – Emmanuel Hanyabui from University of Cape Coast, Ghana

Combines remote sensing and climate data to guide smallholder pineapple farmers in Ghana, providing actionable advice to manage drought risks and protect crops and livelihoods.

Low-cost analog ensemble forecasting – Badreddine Alaoui from Direction Générale de la Météorologie, Morocco

Applies machine learning (ML)-enhanced analog ensemble forecasting to improve low-visibility and surface weather forecasting at Moroccan airports, offering affordable probabilistic guidance.

Leveraging the Global Drought Observatory – Sara Moutia from Direction Générale de la Météorologie, Morocco

Integrates Moroccan and regional drought data into European monitoring platforms, creating actionable, borderless early warning information.

Strengthening Early Warning in Africa (SEWA)

Jointly implemented by the African Union Commission, ECMWF and EUMETSAT as part of the Africa – EU Space Partnership Programme, SEWA aims to develop space-based services and applications to strengthen early warning systems for hazardous weather and climate-related events in sub-Saharan Africa. ECMWF is leading one of the five SEWA activities, which is to enable the co-design and delivery of impact-based forecast services and tools in collaboration with other European stakeholders.

More details of the initiative can be found online: https://www.ecmwf.int/en/about/what-we-do/environmental-services-and-future-vision/strengthening-early-warning-africa

Benefits of assimilating clear-sky GEO radiances at higher spatial and temporal density

Josef Schröttle, Cristina Lupu, Chris Burrows, Angela Benedetti

ECMWF assimilates clear-sky radiance observations from several key geostationary satellites (GEOs), including Meteosat, GOES, and Himawari. These satellites provide continuous coverage from the tropics to the mid-latitudes, observing Earth's atmosphere across a relatively wide range of spatial scales – from planetary-scale patterns stretching across thousands of kilometres down to structures of 1 km in size, with updates every 10 minutes.

Currently, only clear-sky observations are assimilated from the infrared part of the spectrum. These observations provide information on water vapour in atmospheric columns and on winds via wind tracing in the four-dimensional variational data assimilation system (4D-Var). Although clear-sky observations represent only a fraction of the all-sky observations globally, they can reveal mesoscale and synoptic-scale water vapour gradients in troughs of extratropical cyclones, dry air in anti-cyclones or heatwaves, as well as enhance the representation of the wind in the mid-latitude jet-stream region.

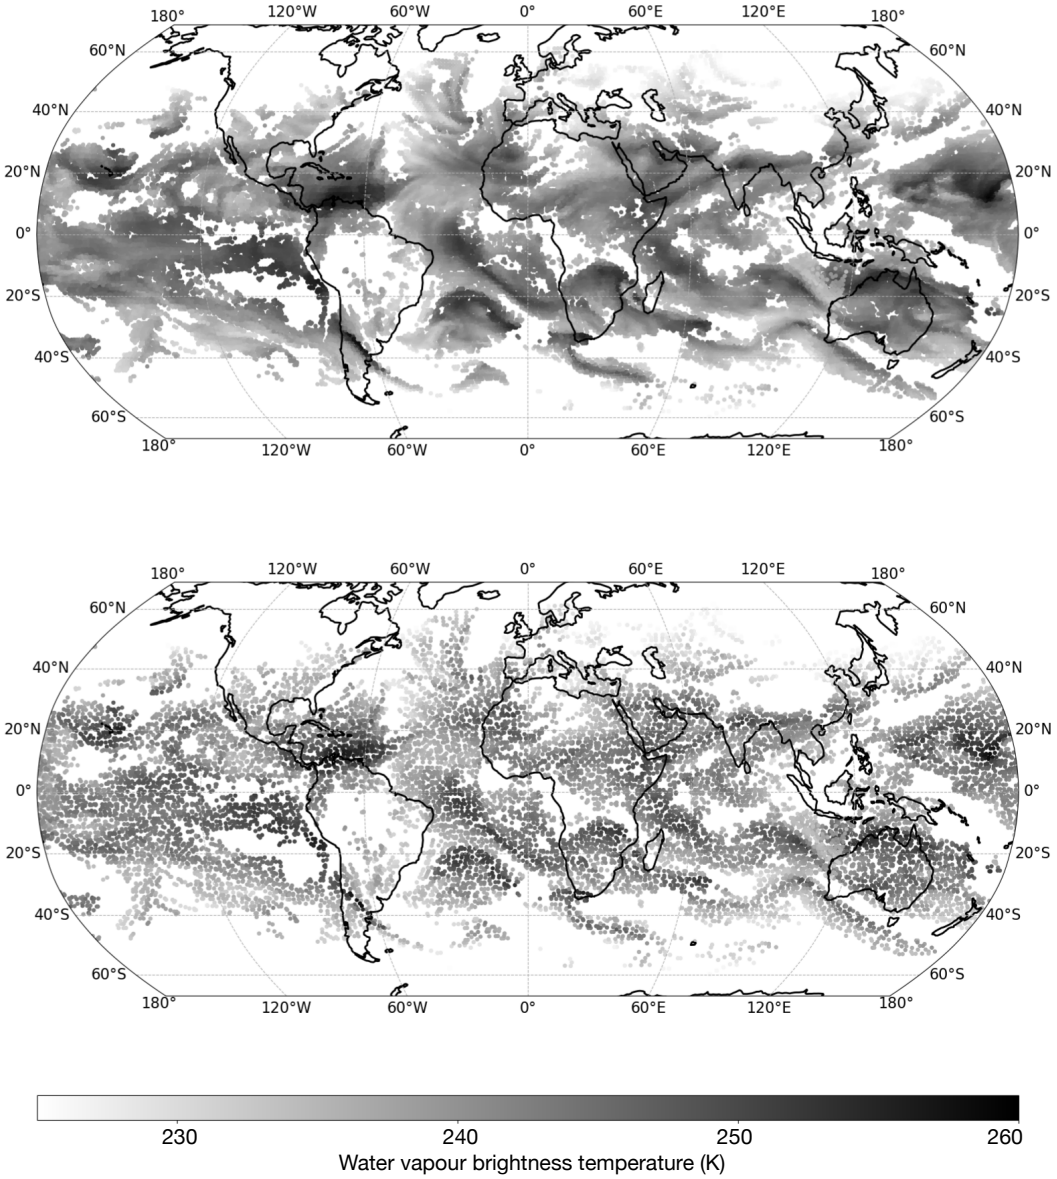

Comparison of thinning strategies for clear-sky water vapour observations. These global maps show assimilated clear-sky radiances from the water vapour 6.2 µm band on 2 December at 07 UTC, during strong jet activity in the southern hemisphere. The top map shows globally denser radiances every 75 km, while the bottom map shows the operational 125 km thinning.

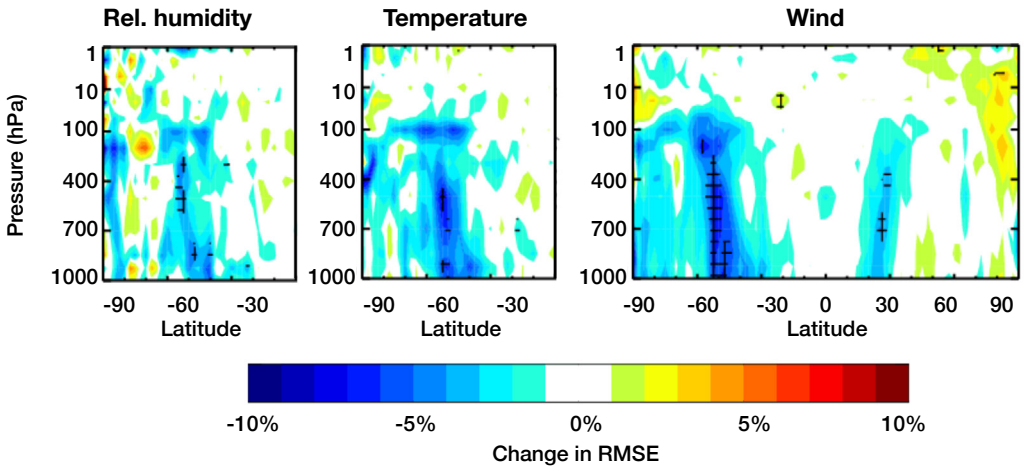

Impact of denser thinning on forecast accuracy. Percentage change in root mean squared error (RMSE) for relative humidity (left), temperature (middle), and wind (right) in five-day forecasts (T +120 h) when assimilating GEO observations with a 75 km thinning compared to the operational 125 km thinning in the control experiment. Improvements (negative values, blue) mainly occur in the southern hemisphere near the jet stream. Both the experimental and control assimilate these observations hourly. Forecast scores are verified against own analysis.

Meteorological impact of denser observations

In Cycle 50r1 of the Integrated Forecasting System (IFS), ECMWF's 4D-Var system will assimilate a denser set of GEO radiance observations – hourly and every 75 km globally, compared to the 125 km configuration (see the first figure). This improvement, achieved by adjusting spatial thinning parameters, leads to measurable gains in forecast skill.

Experiments conducted at the operational resolution of 9 km (TCo1279) for the period 1 December 2022 to 15 January 2023 showed clear improvements in temperature, humidity, and wind forecasts up to five days ahead (see the second figure). The largest impact occurred in the southern hemisphere, where satellite observations play a larger role due to the relatively sparse network of conventional observations and possibly also due to the southern hemispheric summer period, which offers more clear-sky observations. Increased observation density led to significant improvements at the short range, as well as in the troposphere in medium-range forecasts. The fact that temperature, humidity and wind improvements reach throughout the full atmospheric column in the troposphere indicates that there are adjustment processes in the numerical model that are better represented in the 4D-Var system with higher observation density.

Higher frequency sampling

Research experiments have also been carried out assimilating clear-sky water vapour radiances at 30-minute intervals with higher temporal frequency as observed from GOES-16, -17 and -18, Meteosat-12, and Himawari-9. Forecasts based on assimilating higher-frequency observations are compared with a control experiment using an hourly update frequency. Improvements at the short range are consistent for all datasets tested, particularly for the wind field across tropospheric altitudes where GEO sensors have the highest sensitivity to the water vapour signal. Focusing on the assimilation of Himawari-9 observations at 30-minute frequency in TCo1279 experiments, the strongest impact is found in regions where the southern hemispheric jet stream occurs. This improvement is accompanied by a positive impact on the medium-range forecast of the 500 hPa geopotential for days 4 and 5 (see the third figure).

Outlook towards incorporating higher-resolution observations

Future work will exploit the exceptional spatial and temporal resolution of observations from geostationary satellites, such as the Meteosat Third Generation (MTG) series. These satellites will allow us to incorporate a greater portion of water vapour sensitive atmospheric features at finer spatial resolution, further improving the short- and medium-range forecast skill of the IFS down to the km-scale.

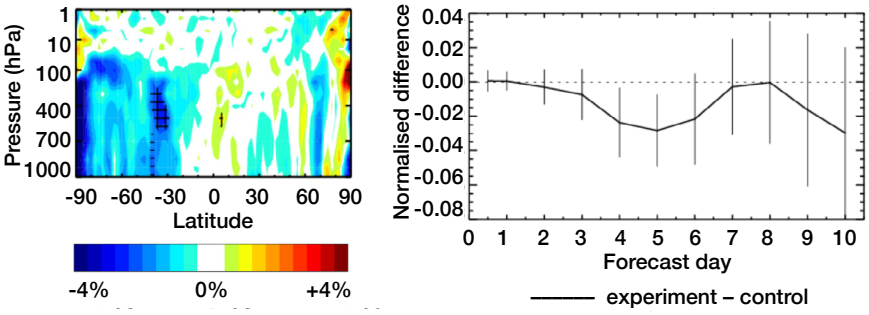

Effect of increased observation frequency on medium-range forecasts. The left panel shows change in RMS wind error at T +120 h when clear-sky Himawari-9 observations are assimilated every 30 minutes instead of hourly. Errors are significantly reduced in the southern hemisphere. The right panel shows corresponding improvements in 500 hPa geopotential in the southern hemisphere on days 4 and 5. All results are verified against own analysis.

Celebrating 50 years of ECMWF

Celebrating 50 years. Member States and esteemed guests at the ECMWF 50th anniversary event in Reading, UK.

In 2025, ECMWF marked its 50th anniversary, celebrating five decades of collaboration, innovation, and impact in numerical weather prediction and related sciences. The anniversary programme brought together scientists, technical experts, and users from our Member and Co-operating States across three event weeks in Bonn (Germany), Bologna (Italy), and Reading (UK), reflecting on ECMWF's achievements and looking ahead to future advances.

The anniversary year opened in April with a week of events in Bonn, combining the Annual Seminar with four scientific workshops. Discussions spanned forecasting in a changing climate, advances in data assimilation, improved representation of surface–atmosphere interactions, and the growing importance of next-generation ancillary datasets. Machine learning featured prominently throughout, highlighting its increasing role alongside physics-based approaches. The breadth of participation and disciplines provided a snapshot of a field undergoing rapid transformation.

In September, the celebrations continued in Bologna, where more than 400 delegates took part in scientific and technical workshops. A central theme was high-performance computing (HPC) and its transformative role in weather and climate prediction. Delegates explored past achievements, current capabilities, and future directions, including the integration of artificial intelligence into forecasting systems and the development of digital twin technologies through the EU Destination Earth (DestinE) initiative.

The European Weather Cloud (EWC) User Workshop was also held in Bologna for the first time in person. Participants from academia, national meteorological services, and industry exchanged ideas and shared experiences on accessing and using ECMWF's growing open datasets. The launch of the EWC Community Hub and the demonstration of new tools for data interaction highlighted ECMWF's commitment to supporting a connected and collaborative research community.

The anniversary programme concluded in December in Reading, with events focused on machine learning for weather forecasting. These included a meeting of the Member State Machine Learning Pilot Project and a showcase of European capabilities, highlighting how machine learning is being integrated alongside physics-based models through close collaboration across ECMWF's Member States. The celebrations were closed with a Gala evening that brought together Member States, partners, and current and former colleagues to mark the end of the 50th anniversary year.

Alongside the events, ECMWF celebrated the anniversary with the release of six publications highlighting key scientific and technical achievements in forecasting, data assimilation and high-performance computing. These were brought together in a commemorative book, *50 years of ECMWF*, reflecting on the Centre's history and its future role in the global meteorological community (available to download at https://www.ecmwf.int/en/elibrary/81684-50-years-ecmwf).

The 50th anniversary year was an opportunity to reflect on ECMWF's journey – from its first operational forecasts in 1979 to today's ensemble Earth-system approaches, machine learning applications and digital twin initiatives – and to thank the staff, partners and Member and Co-operating States who have shaped its success. The collaborations strengthened and insights shared during 2025 form a strong foundation for the next 50 years of advancing weather and environmental prediction.

A new single access point for ECMWF learning resources

Chris Stewart, William Becker

ECMWF has launched a new learning catalogue which provides a single access point for training events and self-paced learning materials produced across the Centre. The catalogue brings together resources from ECMWF's training activities in meteorology and computing, training in the context of ECMWF's support to the EU's Copernicus programme and Horizon Europe-funded projects, as well as upcoming courses and materials from the Destination Earth (DestinE) and Strengthening Early Warning in Africa (SEWA) initiatives. The resources cover topics from meteorology and climate to machine learning (ML).

The new catalogue responds to community feedback and the growing need for clear, unified access to ECMWF's training resources. Users can now browse, filter and access everything from upcoming courses to expert tutorials in one easy-to-navigate place.

One hub for events and learning materials

The catalogue encompasses all of ECMWF's training events, where users can follow links to register for in-person courses and online webinars. It also provides easy access to a rapidly growing library of on-demand training resources across a range of formats:

- Recorded tutorials, webinars and training videos.

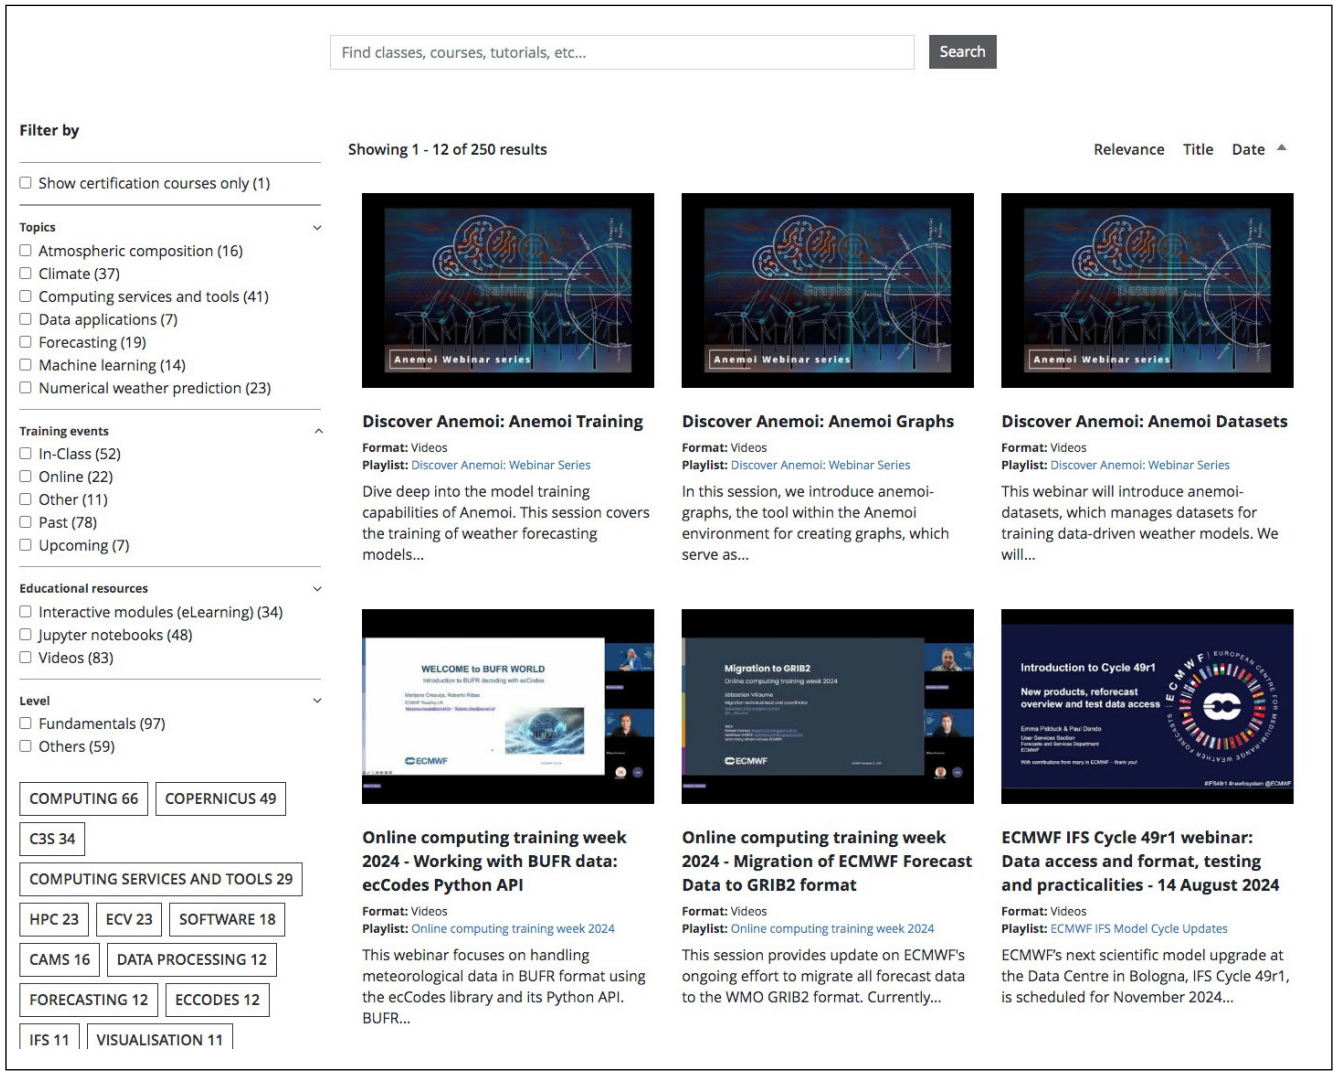

The new ECMWF learning catalogue. The screenshot shows the online training catalogue, with streamlined navigation for users to filter and select resources (https://www.ecmwf.int/en/learning/training/search).

- Interactive eLearning modules, allowing learners to explore concepts at their own pace, from the fundamentals of numerical weather prediction to advanced data assimilation or ML workflows.
- Hands-on Jupyter notebooks, which can be run directly in ECMWF's new JupyterHub environment, enabling users to interact with datasets.

Together, these materials support users ranging from newcomers to experienced practitioners across weather, climate and Earth system science.

Structured navigation with flexible filters

A key improvement is the catalogue's streamlined navigation (see the figure). Users can filter resources by:

- Topic: including various thematic categories
- Format: videos, Jupyter notebooks and interactive modules (eLearning)
- Tags: specific tools, datasets or topics
- Training events: can be filtered to upcoming or past events, and in-class or online
- Level: materials are sorted into "Fundamentals" (introductory material) and "Other" (for more advanced materials).

Each entry includes a description of the resource and links to launch and explore the content.

The learning catalogue reflects the diversity of ECMWF's scientific and operational work, with content organised across multiple thematic areas, including numerical weather prediction, climate, atmospheric composition, forecasting, computing services and tools, data applications and ML.

Supporting hands-on learning

The integration of Jupyter notebooks with ECMWF's Data Store Service JupyterHub allows users to run workflows immediately in the browser, removing setup barriers and enabling practical training with data. This capability strengthens ECMWF's focus on accessible, hands-on learning.

Catalogue expansion

The learning catalogue will continue to grow as new resources and programmes are developed and added. Ongoing maintenance and updates are essential to ensure the content remains relevant and trusted: this will be carried out in collaboration with experts, with the aim of updating materials in an efficient and timely manner. The catalogue provides a long-term home for training resources across ECMWF and its programmes, reinforcing ECMWF's commitment to provide advanced training to its Member and Co-operating States and beyond.

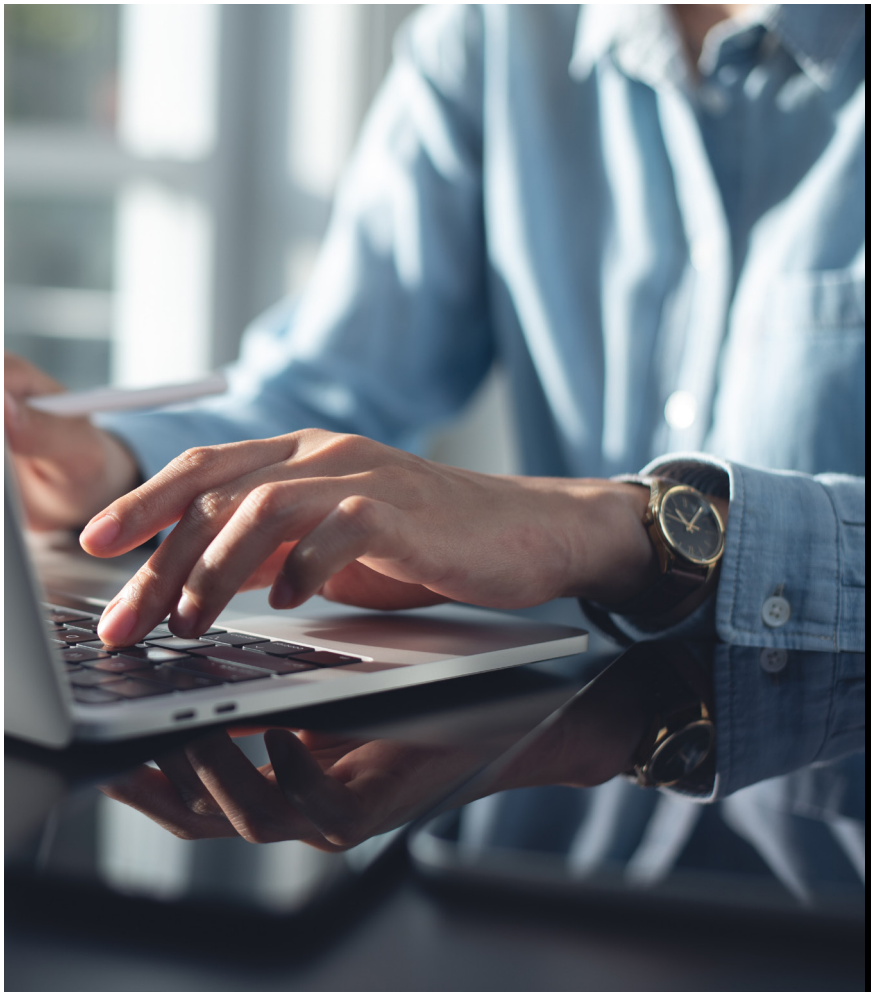

Browse our learning materials:
https://www.ecmwf.int/en/learning/training/search

Hurricane Melissa

Linus Magnusson, Fernando Prates

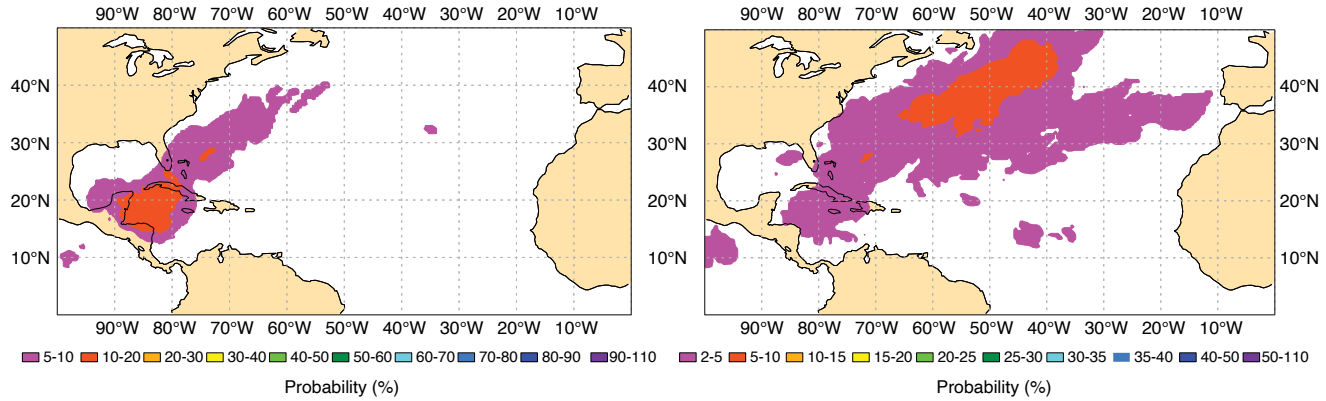

Early signal of increased tropical cyclone risk. Weekly probability of a tropical storm passing within 300 km for the week starting 27 October, based on the forecast from 13 October (left), compared with the climatological probability for this lead time and time of the year derived from ECMWF re-forecasts (right).

At the end of October 2025, Hurricane Melissa brought destructive winds and rainfall to Jamaica as it made landfall as a Category 5 hurricane. Melissa formed on 21 October over the Caribbean Sea from an African Easterly Wave (AEW) that had crossed the Atlantic. The AEW affected the eastern Caribbean on 19 October with moderate rainfall. As it later slowed down over the warm Caribbean Sea, the system developed more organised convection. After becoming a tropical cyclone on 21 October, it continued to slowly move to the north-west for almost a week. The cyclone then rapidly intensified on 26–27 October while moving westwards. On 28 October, the cyclone, now a Category 5 hurricane, turned to the north and made landfall in south-western Jamaica later that day, and in Cuba the day after. The system also brought torrential rainfall to Haiti, leading to landslides and flooding.

Forecasting challenges

As early as 13 October, ECMWF forecasts started to indicate the risk of a cyclone in the Caribbean Sea in the week beginning 27 October (see first figure). This signal was visible in the probability map of weekly tropical cyclone activity produced by the sub-seasonal forecasting system. Over the western Caribbean Sea, the probability for a tropical cyclone (defined as winds exceeding 17 m/s in the forecast) was 10-20%, albeit somewhat to the west of the actual storm. Based on the ECMWF re-forecasts, the climatological probability for a storm during a week at the end of October is 2–5%, indicating enhanced tropical cyclone activity in the forecast. The predictability of the tropical cyclone genesis was linked to the prediction of the AEW, which at the time of forecast initialisation was located just west of Africa.

After the cyclone formed, its westward movement followed by a rapid turn to the north caused large uncertainties regarding the timing and location of the landfall in Jamaica. These uncertainties persisted (see second figure) from the medium-range into the short-range forecast. This is shown by the evolution of the longitude and timing of the cyclone crossing the latitude of 18°N (latitude of southern Jamaica; see third figure). For forecasts issued between 20 and 22 October, the Artificial Intelligence Forecasting System (AIFS) ensemble turned the cyclone northwards too early. This resulted in predictions crossing 18°N that were too far east and too early. This was also predicted by the Integrated Forecasting System (IFS) ensemble. While the ensemble mean error decreased with shorter lead time, the forecast from around 25 October still showed an early northward turn in both systems, putting the eastern part of Jamaica under threat. However, both ensembles had large ensemble spread for this case, indicating large uncertainties.

Intensity prediction issues

Large uncertainties were also present in the prediction of the intensity of the cyclone. Some members in the IFS ensemble predicted the rapid intensification in the forecast between 23 and 25 October. However, they did not reach a pressure below 900 hPa as was the estimated intensity from observations. The AIFS ensemble did not capture the intensification at all. After the rapid intensification of the storm, the ECMWF analyses, and subsequent forecasts, struggled to capture the intensity, with an error of around 60 hPa. This occurred despite the cyclone being well observed by dropsondes and aircraft measurements. One dropsonde measured the strongest verified hurricane wind speed on record, according to the National Center for Atmospheric Research (NCAR). This raises the question of whether the cyclone was too small-scale to be properly handled by the data assimilation and whether the available observations could be used more effectively.

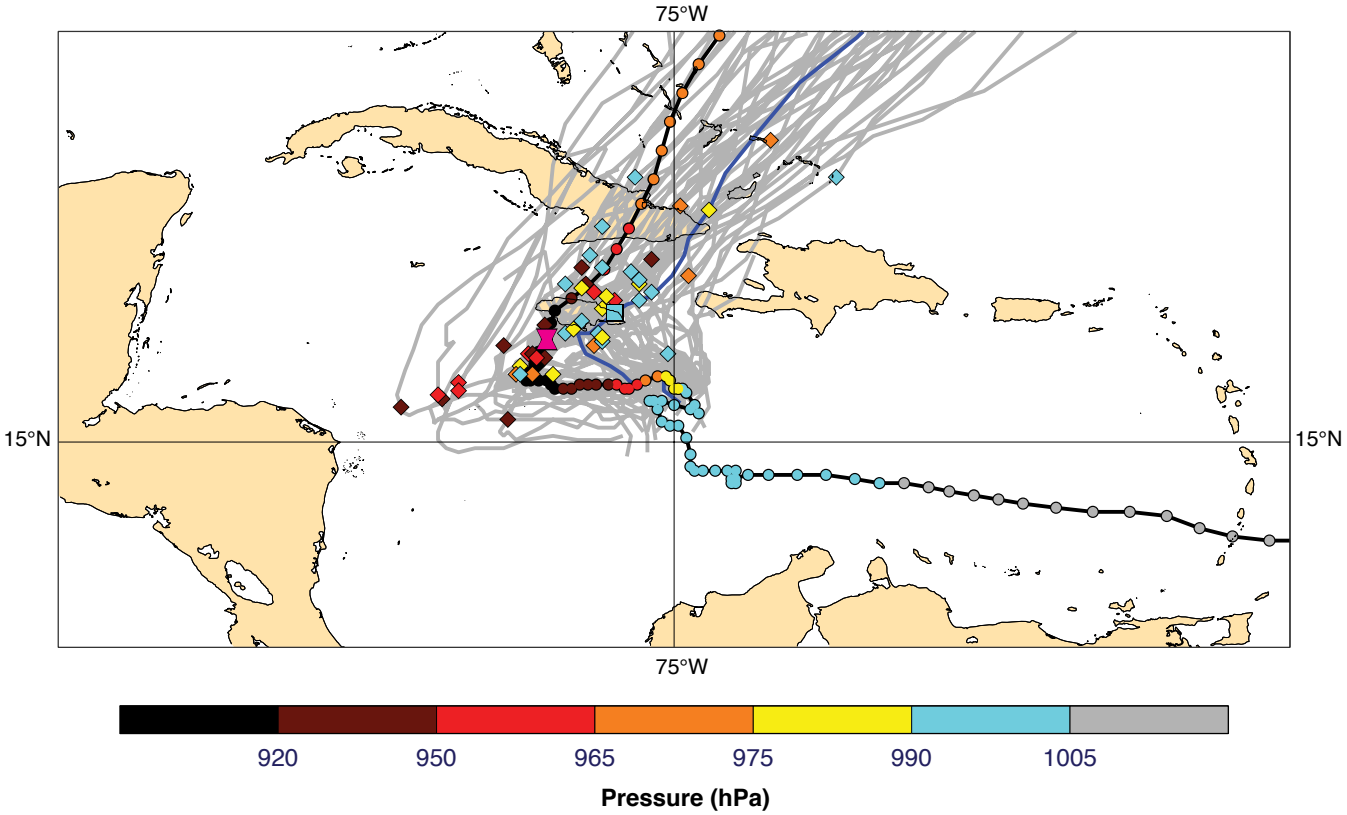

Forecast uncertainty in track and intensity close to landfall. Ensemble forecast of tropical cyclone tracks from ECMWF's Integrated Forecasting System (IFS) initialised at 00 UTC on 25 October. Grey lines show individual ensemble member tracks with the control forecast in blue. Coloured symbols indicate position and central pressure on 28 October 12 UTC. Observations of position and central pressure based on the International Best Track Archive for Climate Stewardship (IBTRACS) are shown as coloured circles, with the position on 28 October 12 UTC as hourglass.

In summary, Hurricane Melissa posed different challenges at different time scales. ECMWF is working together with the University of Miami and NCAR to better understand the predictability of tropical cyclone genesis that originate from AEWs in the medium and sub-seasonal ranges. In the medium range, the presence of the uncertainties in landfall position and timing was a challenge. In the short range, predictions were affected by large analysis errors for the intensity. Further work is under way within the EU Destination Earth initiative to improve the use of observations and the resolution of data assimilation.

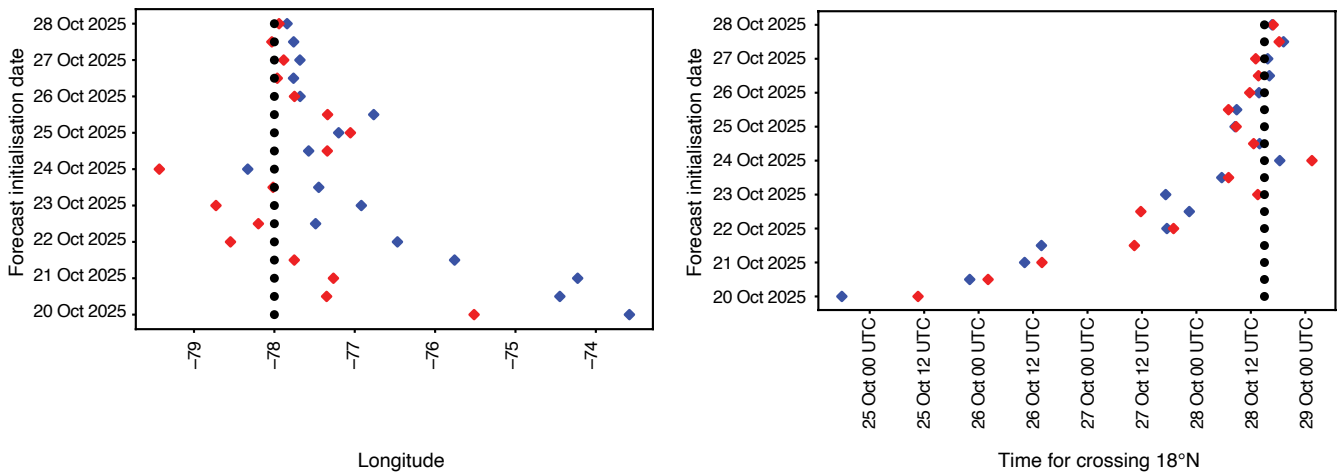

Evolution of forecasts for the cyclone crossing 18°N. Ensemble mean of longitude (left) and timing (right) based on forecasts from the IFS ensemble (red) and AIFS ensemble (blue), compared with observations from IBTRACS (black).

ecWAM 1.5: waves now support two types of GPUs

Ahmad Nawab, Josh Kousal, Jean Bidlot, Michael Lange

To make full use of the diverse range of computing hardware available to the EU Destination Earth (DestinE) initiative across various EuroHPC systems, the ECMWF Integrated Forecasting System (IFS) has been adapted for graphics processing unit (GPU) accelerators. The latest version of ECMWF's open-source operational wave model, ecWAM 1.5 (https://github.com/ecmwf-ifs/ecwam), is now GPU-ready, as part of an ongoing collaborative effort that includes ECMWF, CINECA, Nvidia and AMD.

Scientific significance

Ocean surface waves, generated primarily by the action of wind, are an integral component of the Earth system due to their role in modulating momentum, heat and moisture fluxes between the ocean and atmosphere. Third-generation spectral wave models are the best tools we have for modelling wave generation, transformation, and decay across the global oceans. These models explicitly resolve the full two-dimensional wave action spectrum and evolve it through physically based source terms representing wind input, nonlinear wave–wave interactions, and dissipation, without imposing constraints on spectral shape or wave growth. They also describe key propagation processes such as advection, refraction, and shoaling. ecWAM is ECMWF's state-of-the-art third-generation spectral wave model and a core component of the IFS Earth system.

GPU-adaptation strategy

Because ecWAM contains a diverse set of computational patterns in the model, a mixed strategy was used to adapt ecWAM to GPUs. The wave propagation kernel uses a bespoke low-order computational stencil, and was adapted for GPU via manual, tailored code optimisations and insertion of pragma directives.

Source-term computation forms the other major scientific component of ecWAM. It represents the physical processes that generate, dissipate and redistribute wave energy. This is a large and scientifically fast evolving part of the code, and a manual GPU-adaptation strategy would thus be unsustainable. Here, ECMWF's in-house source-to-source translation

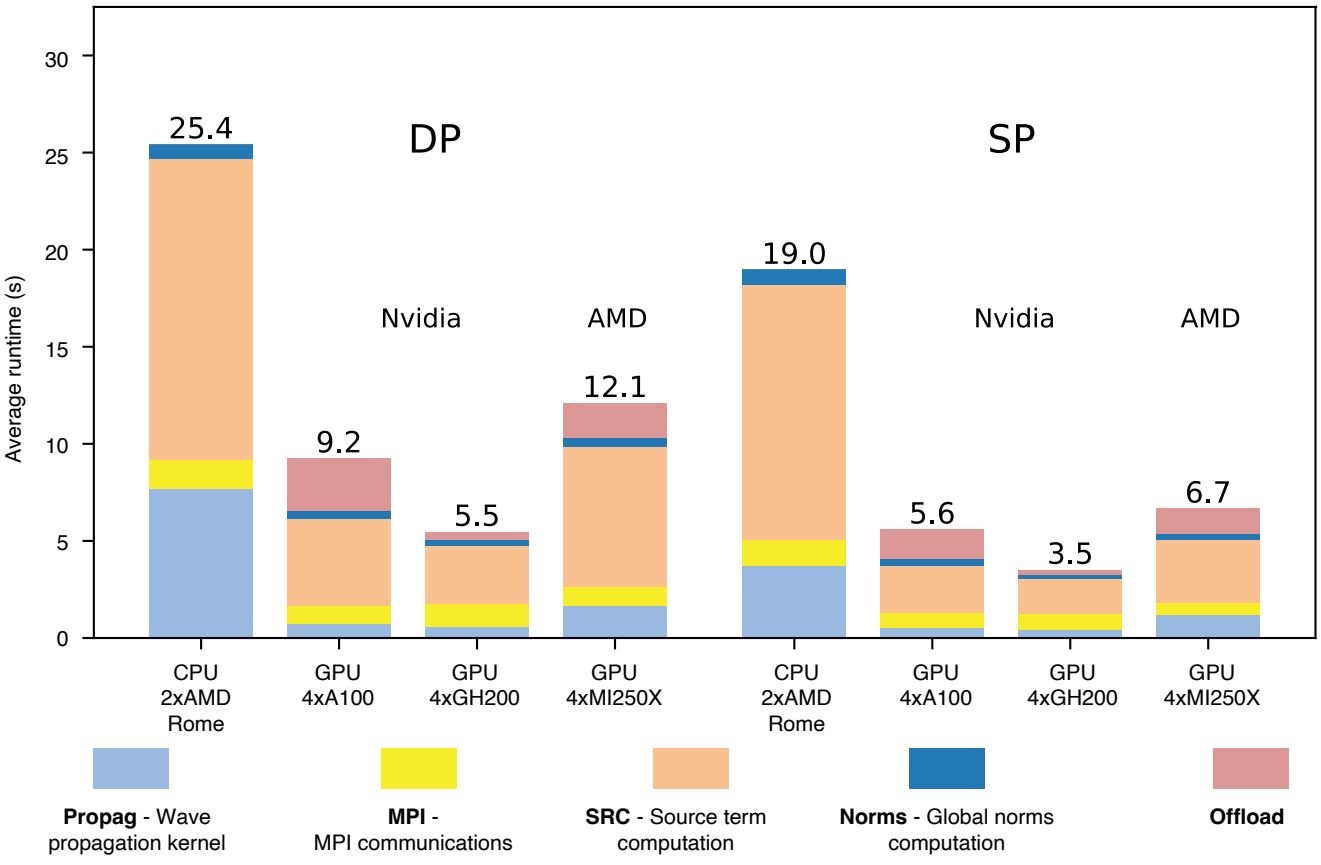

Runtime (model execution time) for a six-hour standalone ecWAM forecast on an O320 grid. Timings were collected on ECMWF's Atos HPC, JUPITER and LUMI. Results are shown for both double precision (DP) and single precision (SP), which is used more commonly in operational configurations. Note that the performance comparisons are across different node configurations and do not necessarily represent the same number of CPUs and GPUs.

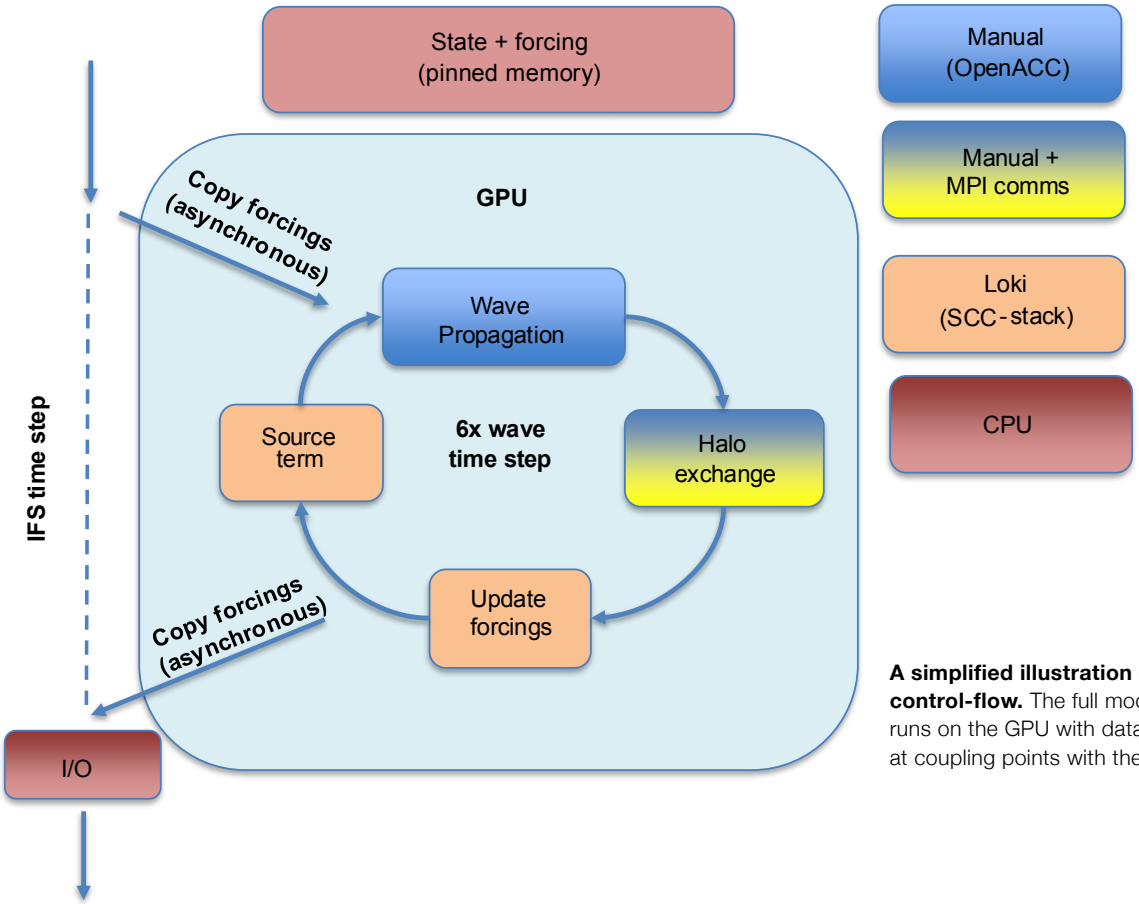

A simplified illustration of ecWAM's control-flow. The full model timestep now runs on the GPU with data exchanged only at coupling points with the IFS.

tool Loki (https://github.com/ecmwf-ifs/loki) is used to automatically generate the GPU code from the central processing unit (CPU) code as a preprocessor step in the ecWAM build procedure. Moreover, as Nvidia and AMD GPUs use distinct programming models, generating the GPU optimised code using Loki is essential for targeting both architectures without making highly intrusive code modifications. As the ecWAM source-term computation kernel features the same "single-column" (with each column comprised of one grid point) memory access pattern as the grid-point computations in the IFS dynamics and physics, the same Loki GPU-adaptation recipes can be used.

Data transfers between the separate memory spaces of a CPU and a GPU represent one of the key challenges of adaptation. To address this, the ecWAM data structures have been rewritten around Field-API (https://github.com/ecmwf-ifs/field_api), an array data management abstraction co-developed with Météo-France for the highly bespoke memory data layouts used in the IFS. Field-API provides vendor-agnostic support for highly optimised CPU-to-GPU data transfers and enables this functionality to be transparent to scientific developers behind an intuitive and non-intrusive abstraction. More details about Loki, Field-API and the wider IFS GPU-adaptation strategy can be found in the accompanying Newsletter article titled 'GPU-adaptation of IFS for EuroHPC machines' (see article by Lange et al. in this Newsletter).

Modularisation and open sourcing

Modularisation and open sourcing have also played a critical role in the GPU-adaptation of ecWAM. Modularisation, specifically the ability to build and test ecWAM independently of the IFS, significantly shortens the development cycle, and open sourcing greatly simplifies collaboration with external contributors.

The initial GPU-adaptation of the wave propagation kernel was performed by CINECA under a contracted DestinE activity. Both modularisation and open sourcing were instrumental in enabling this work to be completed in a timely manner. ecWAM has also benefited immensely from direct vendor contributions, initially from Nvidia and, more recently, AMD, which again was facilitated by open sourcing. Perhaps even more importantly, having a modular and open-source ecWAM also greatly lowers the barrier to continued optimisation, both by internal and external developers.

Computational performance

Whilst there is still room for optimisation, significant speedups have been achieved across a variety of GPU architectures as shown in the first figure. The entire ecWAM timestep is now executed on GPU, and data is only exchanged between CPU and GPU at the coupling points with the IFS when forcings are exchanged either way (see the second figure). Limiting the CPU–GPU data transfers in such a way allows the GPU computational performance gains to translate directly into reduced overall model-execution time.

The conclusion of the EO4EU project

Claudio Pisa, Vasileios Baousis, Mohanad Albughdadi, Tolga Kaprol, Federico Fornari, Marica Antonacci

"AI-augmented ecosystem for Earth Observation data accessibility with Extended reality User Interfaces for Service and data exploitation" (EO4EU) was a Horizon Europe-funded project, which aimed to make Earth observation (EO) data more accessible and usable for environmental, governmental, and business applications by leveraging artificial intelligence (AI) tools and extended reality interfaces. Running from June 2022 to November 2025, the project involved 19 organisations from 10 countries, with ECMWF responsible for the technical coordination.

Earth observation (EO) data holds enormous potential, but its use is often hindered by fragmentation across multiple sources and platforms. Researchers, policymakers, and businesses therefore face significant barriers in accessing, integrating, and exploiting these datasets efficiently. The EO4EU Platform (https://www.eo4eu.eu/platform) was conceived to overcome these obstacles by creating a unified, AI-augmented cloud-native system that simplifies data discovery, processing, and visualisation. With advanced technologies such as knowledge graphs, machine learning (ML), and immersive interfaces, EO4EU is making EO data more accessible and actionable for a wide range of stakeholders. The work aligns closely with European priorities for sustainable environmental governance, while supporting initiatives like Copernicus and Destination Earth (DestinE) and fostering innovation across sectors.

Key capabilities of the EO4EU Platform include:

- **Knowledge graph-based search engine:** facilitates access to heterogeneous EO data sources from Copernicus, DestinE, EUMETSAT and Sentinel-2.
- **Data processing engine:** allows the execution of workflows designed through a graphical interface.
- **AI/ML marketplace:** provides AI and ML models, algorithms, techniques and documentation.
- **A wide range of user interfaces:** graphical, command-line, APIs, and augmented/extended reality, to interact both with the platform functionalities and with data.

The EO4EU Platform employs a wide range of cloud computing and infrastructure-as-code technologies, supporting a geographical multi-cloud deployment on top of the CINECA and WEkEO Data and Information Access Services (DIAS) cloud infrastructures. Its architecture and features underpin its flexibility and ability to serve different user groups (see the figure).

Practical impact

From ECMWF's perspective, the EO4EU project has yielded several technical achievements. Its multi-cloud architecture and infrastructure-as-code approach have acted as a precursor to the design of data bridges within Destination Earth. The project has also contributed to improvements in secret management, cybersecurity, infrastructure monitoring, DNS record and certificate management and application containerisation.

EO4EU has made EO data easier to access and use for both technical and non-technical users. The platform streamlines data selection and retrieval from sources such as Copernicus services, DIAS, and DestinE, reducing complexity and accelerating the development of EO applications. EO4EU also supports the integration of in situ measurements with EO datasets, enabling advanced correlation and fusion techniques. Its

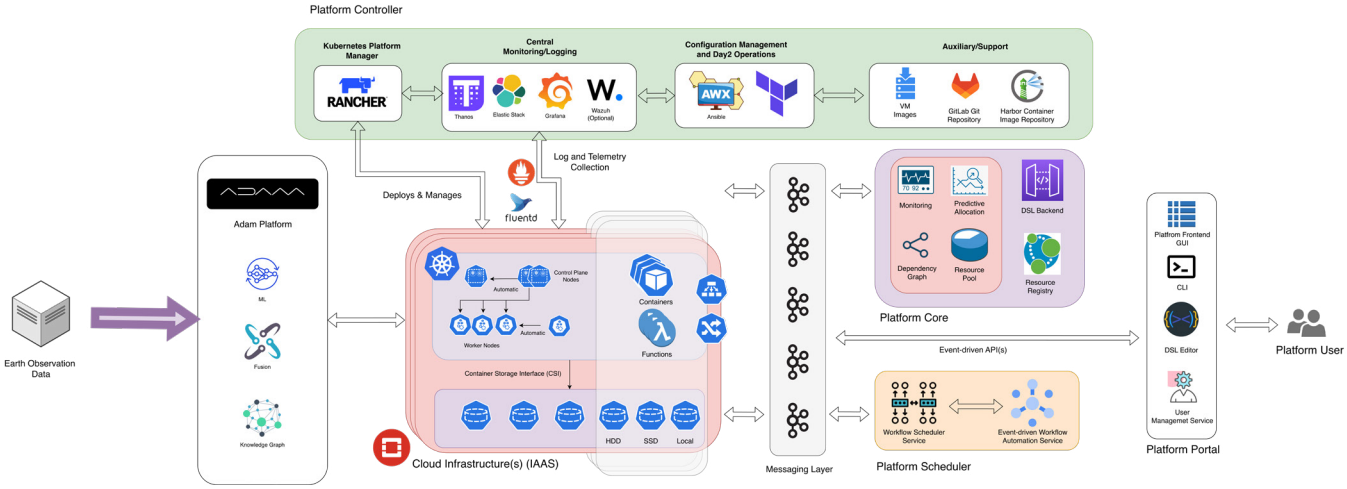

The EO4EU Platform architecture. EO4EU relies on a distributed multi-cloud infrastructure, based on OpenStack and Kubernetes, on top of which cloud-native tools and EO4EU platform components are deployed. Communication between components, observability, orchestration of workflows and machine learning play key roles behind the scenes, while users interact seamlessly with EO4EU through a plethora of interfaces.

visualisation capabilities allow users to explore original, spatio-temporally constrained, and transformed data, creating opportunities for deeper insights into environmental processes.

Use cases

To demonstrate its practical impact and versatility, EO4EU implemented seven use cases across environmental, social, and economic domains:

- Ocean Monitoring for marine health and resource management.
- Food Security to support agricultural planning.
- Forest Ecosystems for biodiversity and conservation.
- Soil Erosion monitoring for land management.
- Environmental Pests detection to protect crops and ecosystems.
- Personalised Healthcare Services using EO data for health-related insights.
- Civil Protection to improve disaster preparedness and response.

Next steps

EO4EU is being onboarded into the DestinE Core Service Platform (DESP), to ensure continued visibility and long-term use beyond the project's official end. The consortium is also exploring opportunities to position EO4EU within emerging frameworks such as AI Factories, leveraging its advanced AI and ML capabilities to support new use cases. These efforts aim to maximise the impact of EO4EU, fostering innovation and enabling broader adoption of Earth observation data for science, policy, and industry.

Funding acknowledgement

The EO4EU project (grant agreement No 101060784) was 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.

Machine learning services at ECMWF: new capabilities for users

Meghan Plumridge, Gert Mertes, William Becker, Xavier Abellan, Milana Vučković, Victoria Bennett, Umberto Modigliani

Machine learning (ML) is becoming increasingly integrated into ECMWF's activities, offering a growing suite of operational services that deliver data, software, and computing capability to the user community. Developments in this area are fast-paced and highly collaborative; teams across ECMWF and European national meteorological and hydrological services (NMHSs) are contributing to these shared tools and services.

In 2025, ECMWF continued to expand and refine the tools enabling this shift. Operational data-driven forecasts, open-source modelling frameworks, graphics processing unit (GPU)-enabled computing environments and new training materials have significantly lowered the barrier to engagement. A broad range of ML services is now available to support users at every stage, from first-time exploration to the development of advanced research and operational workflows.

AIFS: data-driven forecasting now in operations

A major milestone for the meteorological community was reached on 25 February 2025, when ECMWF began producing operational forecasts using the deterministic

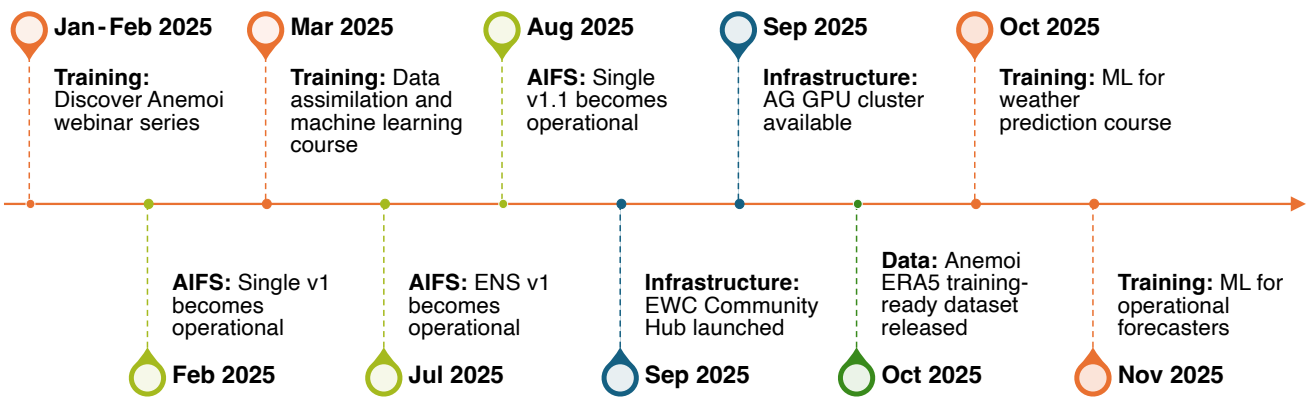

Machine learning activities and developments at ECMWF in 2025. Key milestones in training, infrastructure, data and operations.

version of the Artificial Intelligence Forecasting System (AIFS Single). This marked the first time a global, data-driven prediction model became part of an operational forecast service at an international meteorological centre.

AIFS Single produces a full 15-day forecast in minutes on ECMWF's high-performance computing facility (HPCF), showcasing the efficiency of AI-based weather prediction.

On its first full day of operations, AIFS Single disseminated 130 gigabytes of customised forecast data to 46 remote sites around the world. In addition to these bespoke products, the complete set of AIFS Single forecast products was made freely available on open data platforms, democratising access to forecast data.

The operational launch of the AIFS ensemble (AIFS ENS) model followed on 1 July 2025, adding uncertainty estimation and enabling users to explore probabilistic data-driven prediction. Together AIFS Single and AIFS ENS generate approximately 6 terabytes of forecast data each day, providing a robust data-driven complement to the operational Integrated Forecasting System (IFS).

Access AIFS data and products

All global real-time forecast products produced by the AIFS are freely available to the public.

- **AIFS open data** are published immediately after production on a variety of open data platforms.
- **AIFS forecast charts** showing a subset of the most popular products are published on the OpenCharts platform each day. These charts allow users to interact with AIFS weather forecast products without downloading data.

Bespoke real-time and historical AIFS forecast products are also available for registered users.

For users wishing to run the AIFS and generate data themselves, model weights for AIFS Single and AIFS ENS are also openly available on HuggingFace (https://huggingface.co/ecmwf).

Software and data services

Anemoi: the framework powering data-driven innovation

Co-developed by ECMWF and NMHSs across Europe, Anemoi is an open-source, Python-based framework that provides a complete toolkit for developing data-driven weather forecasting models. Since its launch in 2024, the Anemoi framework has grown in functionality and now forms the backbone of ECMWF's operational data-driven forecasting system, the AIFS. The Anemoi ecosystem is also evolving and actively being developed. A key achievement in 2025 was the release of an open Anemoi dataset – a training-ready version of ECMWF's ERA5 reanalysis. ERA5 has been widely used to train data-driven weather forecasting models, including the AIFS. The release of the ERA5 Anemoi dataset enables users to train their own data-driven models using the Anemoi framework.

Anemoi's impact across the community was recognised in 2025 with two major international awards: the EMS Technology Achievement Award 2025 and the HPCwire Readers' and Editors' Choice Award 2025.

Get started using Anemoi

The Anemoi ecosystem includes comprehensive documentation, example notebooks and introductory webinars covering:

- dataset creation and management;
- model configuration and training;
- evaluation and forecast generation.

To generate your own AIFS forecast using the **anemoi-inference** package, follow the notebook examples for AIFS Single and AIFS ENS (https://huggingface.co/ecmwf).

To download the Anemoi ERA5 dataset and get started training your own model, refer to the anemoi-training package documentation (https://anemoi.readthedocs.io/projects/training/en/latest/user-guide/download-era5-o96.html).

ai-models

In addition to developing its own models, ECMWF runs several third-party data-driven weather forecasting systems alongside the AIFS using the ai-models Python package. The models include DeepMind's GraphCast, Huawei's Pangu-Weather, Microsoft's Aurora and NVIDIA's FourCastNet. These models are run daily for verification and comparison with the AIFS.

Interact with third-party models

Although ai-models is experimental, with limited support and no plans for further expansion, the package and its output are available to users:

- Graphical forecast products and performance scores are publicly available on the OpenCharts platform (https://charts.ecmwf.int/catalogue/packages/ai_models/), enabling comparison with the AIFS and IFS.
- Users can also generate their own forecasts from these models using the experimental ai-models Python package (https://github.com/ecmwf-lab/ai-models), which handles the retrieval of input data required to initialise the models.

Computing resources

GPU access

Many ML workflows require GPU acceleration and specialised software environments. In 2025, ECMWF expanded both its high-performance computing and cloud-based offerings to support this growing demand. Access is available to users affiliated with NMHSs in ECMWF's Member and Co-operating States.

ECMWF high-performance computing facility

In 2025, a new GPU cluster was introduced to the HPCF. The cluster, called AG, more than doubles ECMWF's existing GPU capacity. The AG cluster consists of 30 accelerated nodes, each containing four Grace Hopper GH200 Superchips, for a total of 120 GPUs. To facilitate access, ECMWF's JupyterHub service (https://jupyterhub.ecmwf.int/hub/home) allows users to run interactive GPU sessions in a browser.

European Weather Cloud

GPUs are also available to users of the European Weather Cloud, which currently offers 32 Nvidia Ampere A100 GPUs.

In addition to hardware, the EWC also provides ready-to-use ML software environments. In September 2025, the European Weather Cloud Community Hub (https://www.europeanweather.cloud/community-hub) was launched. The Hub offers a centralised platform where users can discover, evaluate, select, and deploy items and services tailored to the European meteorological community (see the screenshot). In the context of ML, this includes ready-made templates that provide pre-configured software stacks for the AIFS and Anemoi.
These resources lower the barrier to entry for users who want to experiment with ML.

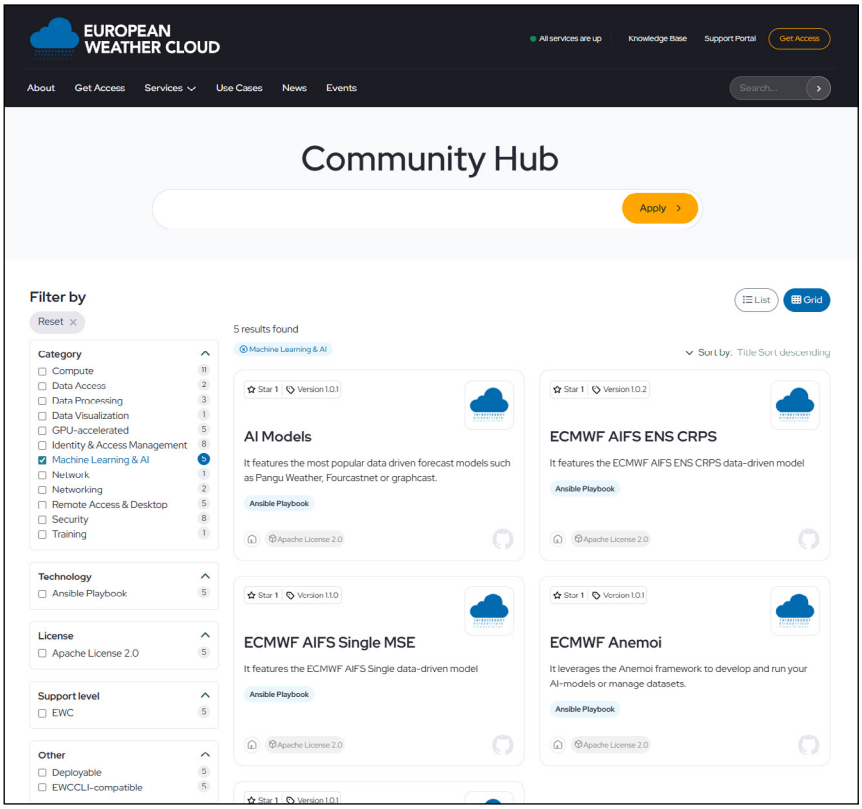

Screenshot of the European Weather Cloud Community Hub. Users can find ready-made ML templates.

Use computing resources

Users who wish to access GPUs on the HPC or EWC should:

- Refer to the access guidelines (https://www.ecmwf.int/en/computing/access-computing-facilities), to check eligibility.
- If you are a Member or Co-operating State user, contact your Computing Representative to arrange access.
- You may also apply for access as part of a European Meteorological Infrastructure (EMI) R&D project.

Training and community engagement

Training and knowledge exchange remain essential to supporting the growth of ML. Throughout 2025, ECMWF delivered a number of training courses, webinars and training resources, with materials available on the ECMWF website:

- In January, the six-part Discover Anemoi webinar series demonstrated how to use key components of the Anemoi ecosystem, covering dataset creation, graph construction, model training, inference and contributing as a developer.
- In March, the five-day Data assimilation and machine learning training course introduced meteorologists to a range of ML techniques applicable to data assimilation workflows through lectures and hands-on exercises.
- In October, the five-day Machine learning for weather prediction training course brought together meteorologists and ML scientists, providing an overview of ML in Earth system sciences and introducing software and hardware frameworks available at ECMWF.
- In November, the three-part Machine learning for operational forecasters webinar series attracted operational forecasters from more than 100 countries, offering guidance on accessing, using and interpreting AIFS forecast products.

Looking ahead to 2026

The landscape of ML at ECMWF has evolved significantly over the past five years.The milestone of the first operational data-driven forecast with the AIFS marks only the beginning of exploring ML's potential.

In 2026, the AIFS Single and AIFS ENS models will be upgraded in parallel with the implementation of IFS Cycle 50r1. The upgrades will feature improved forecast scores and new output variables, including snow.

The year ahead will also see the release of the first AIFS sub-seasonal model. More than 60 sub-seasonal models are currently being developed by international teams as part of the AI Weather Quest competition. ECMWF is prototyping three models ahead of the 2026 release. The competition is ongoing and welcomes both individual participants and teams of up to ten members.

Finally, further training and outreach activities are planned for 2026. ECMWF will host the fifth iteration of the joint ECMWF-ESA ML workshop in Bologna, Italy, in April 2026. Under Destination Earth, a course on machine learning will take place in Bonn in February 2026, followed by a series of online courses on ML in Earth Systems Modelling beginning in March, which will run over 2026. The data assimilation and machine learning course will also run again in March 2026 in Reading. Users may register for these training courses and workshops on the ECMWF website and by following the Destination Earth page (https://destine.ecmwf.int/ml-training/).

Further resources

To keep up with the rapid pace of ML developments at ECMWF, users can refer to the website and user guide for ML services (https://confluence.ecmwf.int/display/UDOC Machine+Learning+ Services+and+Support).

Users are also invited to join our communication channels to receive updates about future AIFS cycle upgrades and other ML activity:

- Subscribe to our mailing lists by emailing forecast_changes-request@lists.ecmwf.int (subject: *subscribe*) for AIFS cycle updates, and ml_training-request@lists.ecmwf.int (subject: *subscribe ml-training*) to receive information about upcoming training events.
- Join our user forum (https://forum.ecmwf.int/) and watch the announcements in the 'IFS, AIFS and OpenIFS' category.
- Follow our LinkedIn channel for users (https://www.linkedin.com/showcase/ecmwf-users/).

Changes in Directors

A number of changes to ECMWF's senior leadership took effect on 1 January 2026.

Florian Pappenberger succeeded Florence Rabier as Director-General at the start of the year. Florian was previously Director of Forecasts and Services.

Andy Brown stepped down as Director of Research at the end of December 2025. Stephen English, formerly Deputy Director of Research and Lead Scientist, is acting as Director of Research. Following Florian Pappenberger's move to Director-General, Umberto Modigliani, previously Deputy Director of Forecasts in the Forecasts and Services Department, became acting Director of Forecasts and Services.

These acting appointments provide stability and continuity while recruitment processes for the roles are carried out in due course.

Improving mesoscale aspects in the ensemble forecast initial conditions

Martin Leutbecher, Sarah-Jane Lock, Elias Hólm, Aristofanis Tsiringakis, Marieke Plesske, Ioan Hadade

The increase in horizontal resolution of the Ensemble of Data Assimilations (EDA) from 18 km to 9 km in Cycle 49r1 of the Integrated Forecasting System (IFS) has sharpened the depiction of mesoscale features, enabling better representation of weather phenomena – especially in areas where observational data is sparse. While this has clear benefits, it has also exposed a challenge in how ensemble forecast initial conditions are constructed.

The methodology for constructing ensemble initial conditions involves re-centring the EDA onto the deterministic analysis. However, the re-centring can introduce spurious structures in the smaller scales for individual ensemble members. These artefacts have become more noticeable with the increased mesoscale variability in Cycle 49r1 and are particularly evident for tropical cyclones, which may appear deformed or may have double cores. To address this, a new method has been developed to re-centre the EDA onto the deterministic analysis in a scale-dependent way, improving the realism of the mesoscales in the ensemble initial conditions. This article outlines the motivation for the new approach, explains how scale-dependent re-centring is implemented, and summarises its meteorological impacts and operational integration in Cycle 50r1.

As well as perturbations from the EDA, the ensemble initial conditions include singular vector perturbations. The method for generating these remains unchanged, so we do not discuss them further, but they are nonetheless part of the experiments reported below.

Re-centring of the EDA

The EDA was introduced in 2010 as a method for representing initial uncertainties in ECMWF ensemble forecasts (Buizza et al., 2008). In addition, it plays a key role in estimating situation-dependent background error covariances in variational data assimilation. Each ensemble forecast begins with slightly different initial conditions that reflect possible errors in the analysis. These are constructed by combining the deterministic analysis with perturbations from the EDA. Specifically, the difference between an EDA member and the EDA mean (i.e. an EDA perturbation) is added to the deterministic analysis (see Box, equation 1). The combination of the EDA with the deterministic analysis can be viewed as a shifting, or 're-centring', of the EDA with an increment consisting of the difference between the deterministic analysis and the EDA mean (Box, equation 2).

Re-centring is an important step as perturbed EDA members use a computationally cheaper configuration and are based on observations available six hours earlier than those used in the deterministic analysis (Lang et al., 2015). Aligning the ensemble members with the deterministic analysis therefore increases the probabilistic skill of the ensemble forecasts.

However, this approach has limitations. At the higher resolution in Cycle 49r1, the EDA contains richer mesoscale variability, and the re-centring occasionally produces unrealistic small-scale features in the ensemble initial conditions, in particular for those that are poorly constrained by observations (Lang et al., 2015). The variability in smaller scales in the EDA increased considerably with the horizontal resolution increase to TCo1279 (9 km) in Cycle 49r1 (Roberts et al., 2024).

Scale-dependent re-centring of the EDA

Motivated by work on re-centring perturbed EDA members within the EDA itself (Hólm et al., 2022), the idea emerged to spatially filter the increment used to re-centre the ensemble initial conditions on the deterministic analysis. The filtering operation is implemented in spectral space and applied to the upper-air fields of the increment (Box, equation 3). The filter scales are controlled by specifying a filter wavenumber, N. Waves with a total wavenumber larger than N (representing the smaller scales) are removed by the filter, while those with a total wavenumber up to and including N (representing the larger scales) remain unchanged. In this way, each ensemble member keeps the overall structures from the corresponding EDA member, while only the larger scales are shifted towards the more accurate and more recently updated deterministic state. Sensitivity tests on the filter wavenumber were carried out to determine the impact on ensemble skill and the occurrence of spurious structures in the ensemble initial conditions.

When N is small, spurious structures disappear from the initial conditions. However, when N is too small, too little information is retained from the more recent deterministic analysis, leading to degraded ensemble skill compared to using the standard EDA re-centring approach that uses all scales of the increment. Noticeable skill degradations start to emerge for values of N of less than about 100. Conversely, when N is larger than about 250, significant spurious structures remain, particularly around tropical cyclones. Based on these results, a spectral filter with N = 159 was chosen for implementation. This corresponds to waves with a wavelength of about 250 km.

doi: 10.21957/jcme-9t91

Re-centring methodology

Standard re-centring

The initial condition for each ensemble member j is constructed by adding that member's EDA perturbation to the deterministic analysis:

$$x_j = x_c + x_{EDA,j} - \overline{x_{EDA}} \quad (1)$$

Meaning:

Initial condition = deterministic analysis + deviation of that EDA member from the EDA mean.

See the schematic (a) for an illustration of the terms.

Note, the EDA does not run in the time-critical path, so its analyses are not yet available when the ensemble forecasts are started. Instead, 6-hour EDA forecasts from the most recently available EDA analyses must be used as a proxy.

Figure 1 illustrates examples of x_j for mean sea level pressure. Figure 2 represents x_c, and examples of $x_{EDA,j}$ can be seen in Figure 3. Not illustrated: $\overline{x_{EDA}}$ would be the mean field constructed from all EDA members, $x_{EDA,j}$ $(j=1,\ldots, 50)$.

Scale-dependent re-centring

Equation (1) can be rewritten as:

$$x_j = x_{EDA,j} + x_c - \overline{x_{EDA}} \quad (2)$$

Here, the increment $x_c - \overline{x_{EDA}}$ shifts the EDA member towards the deterministic analysis.

In the new approach, the increment is passed through a spectral filter F_N, so that only large-scale components are retained.

$$x_j = x_{EDA,j} + F_N(x_c - \overline{x_{EDA}}) \quad (3)$$

Wavenumbers up to N are preserved; higher wavenumbers are removed. This keeps the smaller-scale patterns from the EDA while aligning only broad-scale features with the deterministic analysis.

Figure 4 illustrates examples of x_j for mean sea level pressure, with spectral filter F_N applied for $N=159$.

a Standard re-centring

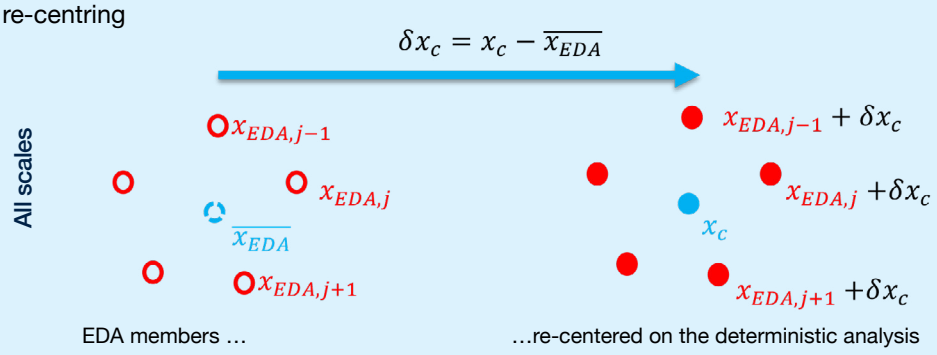

b Scale-dependent re-centring

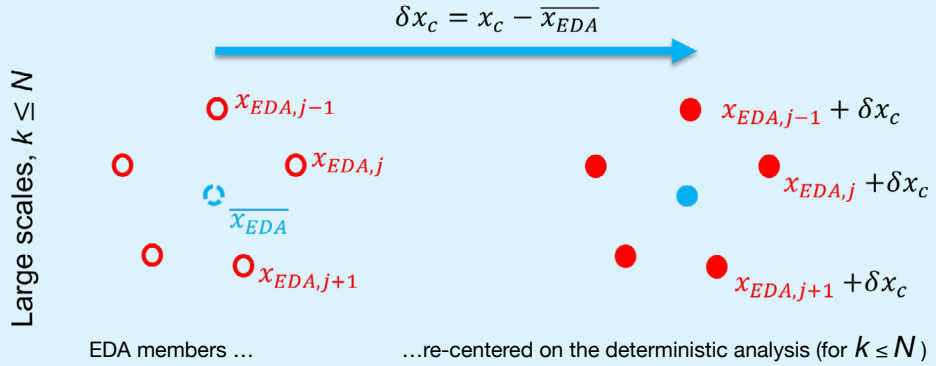

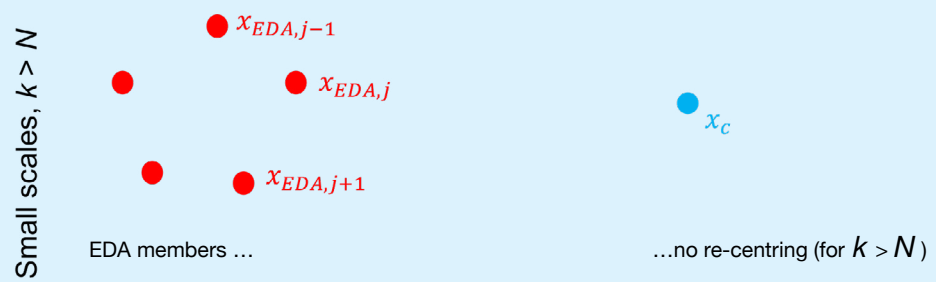

case study Tropical Cyclone Freddy

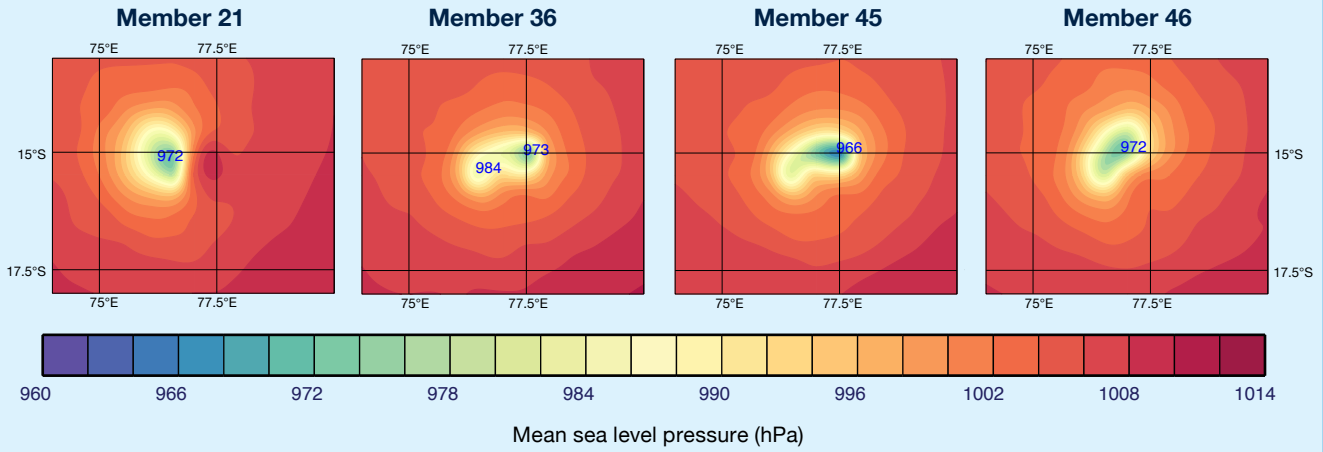

FIGURE 1 Mean sea level pressure (hPa) of the initial conditions for Tropical Cyclone Freddy (valid 17 February 2023, 00 UTC) of four ensemble members constructed with the standard EDA re-centring operational in Cycle 49r1. Members display unrealistic structures such as distorted or double-centre vortices.

The benefits of scale-dependent re-centring of the EDA are particularly evident in the ensemble initial conditions for tropical cyclones.

Tropical Cyclone Freddy was an intense and long-lasting tropical cyclone that tracked westward across the Indian Ocean in February 2023, passing north of La Réunion and later making landfall in Madagascar and Mozambique. Using the standard re-centring in Cycle 49r1, some ensemble members display distorted, elongated or double-core vortex patterns. Examples of ensemble members with artefacts in the initial conditions are shown in Figure 1:

- Member 21 has an anticyclonic anomaly located east of the centre of the cyclone leading to a zone of enhanced pressure gradients and anomalously large curvature of the isobars in the SE quadrant of the storm.
- Member 36 has a double centre and is elongated in the zonal direction.
- Member 45 is also elongated in the zonal direction with a hook-shaped appearance and anomalous curvature of the isobars south of the centre.

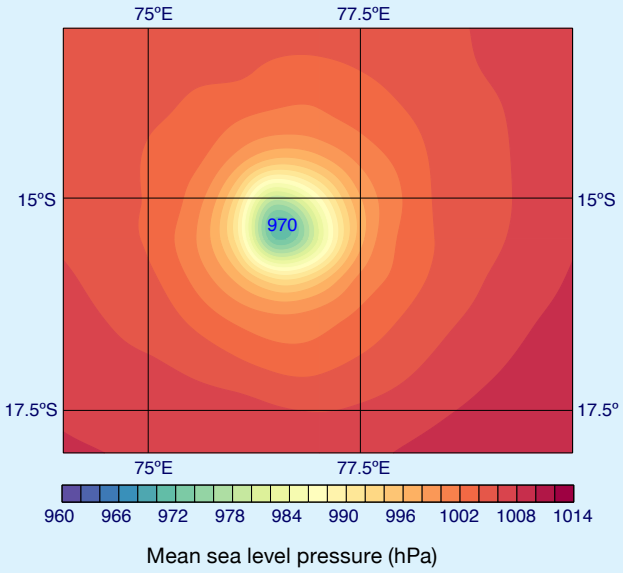

FIGURE 2 Mean sea level pressure (hPa) of the deterministic analysis for Tropical Cyclone Freddy (valid 17 February 2023, 00 UTC).

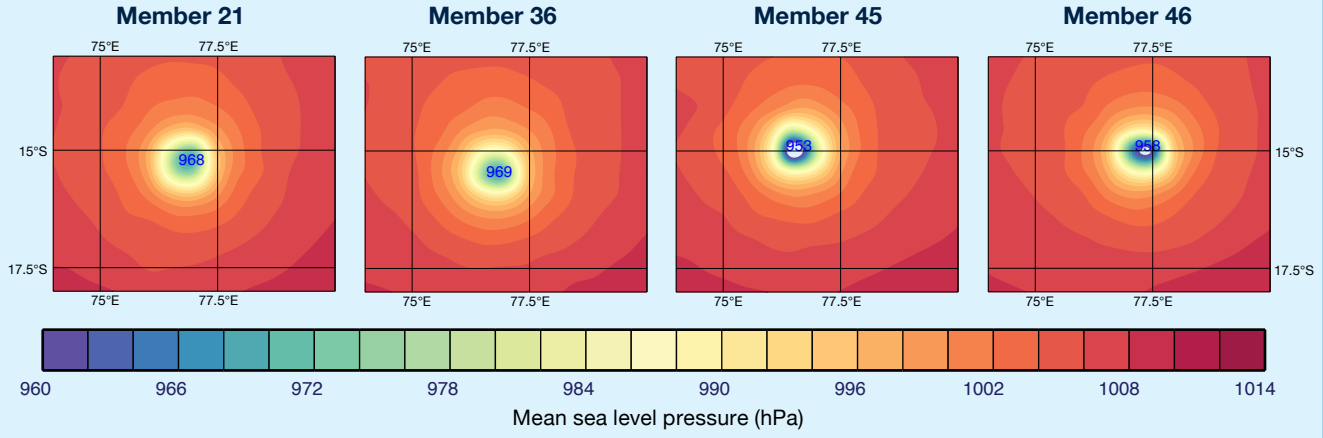

FIGURE 3 Mean sea level pressure (hPa) of 6-hour forecasts from EDA members 21, 36, 45 and 46 for Tropical Cyclone Freddy (valid 17 February 2023, 00 UTC).

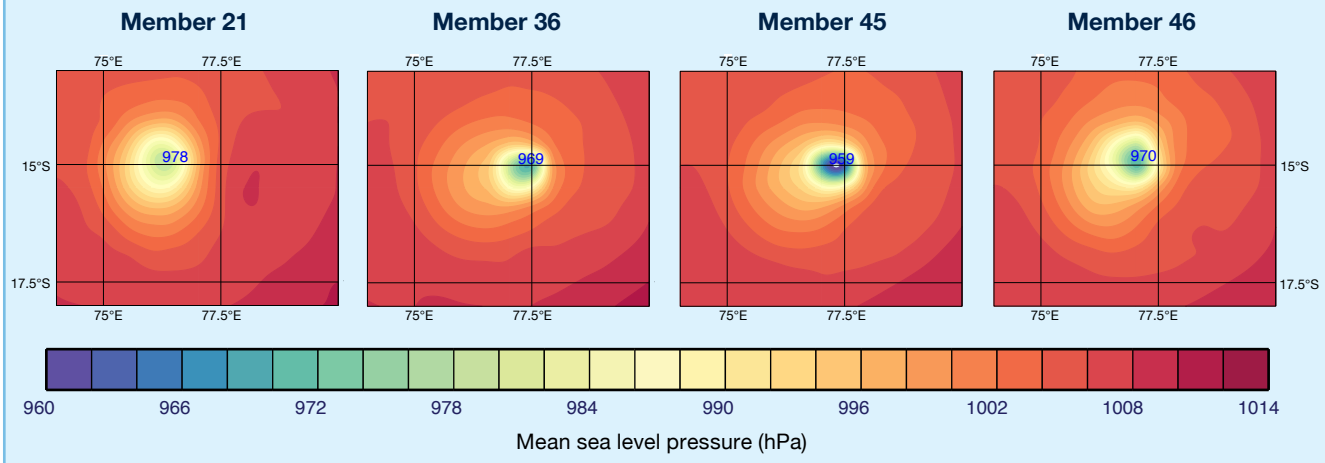

FIGURE 4 Mean sea level pressure (hPa) of the initial conditions for Tropical Cyclone Freddy (valid 17 February 2023, 00 UTC) of four ensemble members constructed with the scale-dependent EDA re-centring that will be operational in Cycle 50r1. The unrealistic vortex distortions seen in Figure 1 are no longer present.

- Member 46 is elongated in the south-west to north-east direction and has anomalously curved isobars on the south-east side.

These unrealistic structures were absent in the deterministic analysis and the EDA members used for constructing the initial conditions (Figures 2 and 3). Importantly, these structures disappear during the first 24 hours of the forecast, which is consistent with them being artefacts of the re-centring process rather than real atmospheric phenomena.

When scale-dependent re-centring of the EDA was applied (N = 159), all four members showed much more axisymmetric, realistic cyclone structures (Figure 4).

Across the 50-member ensemble for this case, roughly a third of the members showed major deviations from axisymmetry in the Cycle 49r1 initial conditions, but none exhibited major deformations when applying the spectral filter to the increment.

Systematic evaluation of the impact of the revised ensemble initial conditions

The impact of scale-dependent re-centring on ensemble forecasts has been evaluated for a larger sample of cases as is usual for any meteorologically active change that is planned to enter operations. Medium-range ensemble forecast experiments have been run for two periods at operational resolution (TCo1279, approx. 9 km): June–August 2022 (85 cases) and December 2022–February 2023 (83 cases). These experiments included ten perturbed members, initialised from analyses and an EDA at the same resolution. The introduction of the scale-dependent re-centring (N = 159) leads to a small increase in mean tropical cyclone core pressure of about 1–2 hPa which quickly decreases with forecast lead time. In terms of root mean square (RMS) error of the ensemble mean, the impact on core pressure is neutral. The revision of the ensemble initial conditions does not lead to a statistically significant change in the error of the ensemble mean position of tropical cyclones beyond initial time. The spread in the tropical cyclone position shows a moderate increase in the boreal summer period at lead times up to four days, while no change is observed in the boreal winter period.

The medium-range ensemble forecasts have also been used

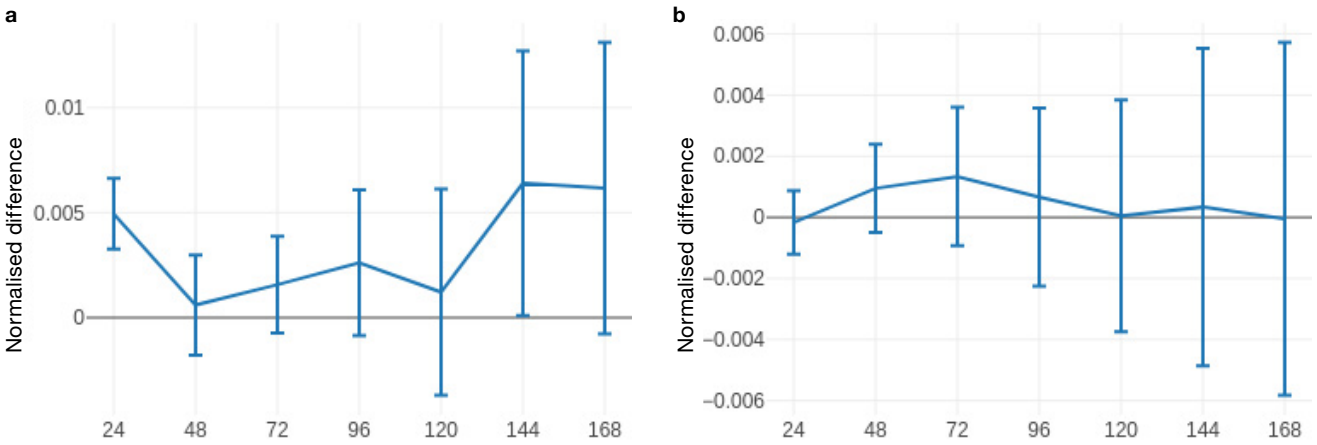

FIGURE 5 Relative change in fair Continuous Ranked Probability Score (fCRPS) versus lead time (hours) for 500 hPa geopotential in (a) the northern extratropics and (b) the southern extratropics during the period June–August 2022. Positive values imply higher skill for the experiment with scale-dependent re-centring of the EDA. Vertical bars are confidence intervals at the 99.7% level.

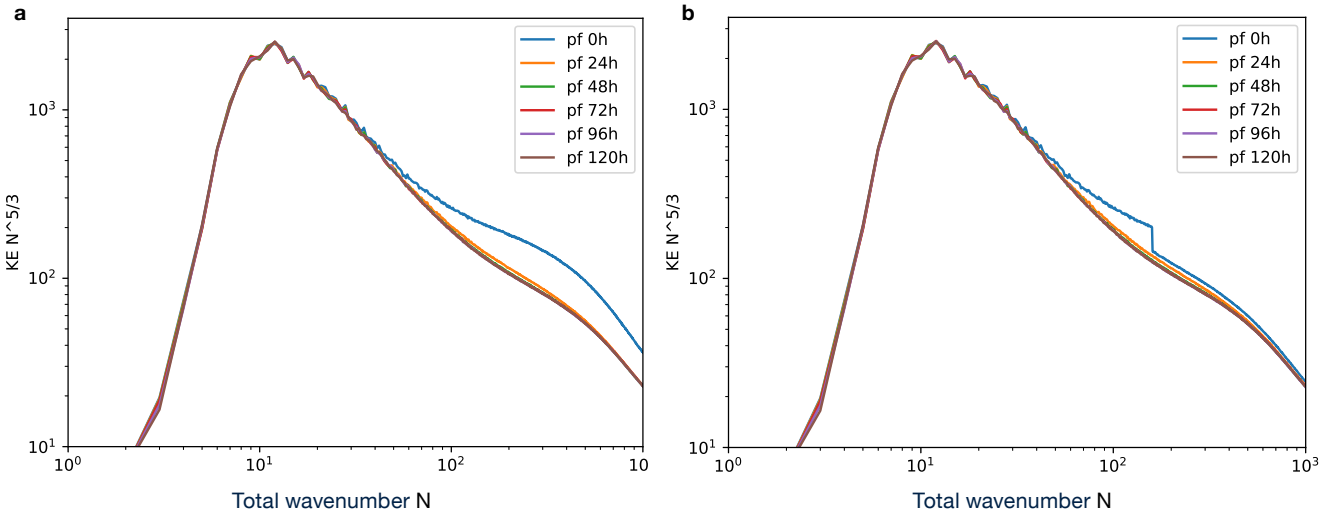

FIGURE 6 Kinetic energy spectra for 250 hPa wind anomalies for (a) an ensemble with initial conditions as in Cycle 49r1 and (b) an ensemble with initial conditions using the scale-dependent EDA re-centring. Curves correspond to lead times from 0 to 120 hours. The kinetic energy is scaled so that a -5/3 spectrum would appear as a flat line. Data are aggregated for June 2022 and all ten perturbed members. Anomalies are computed with respect to the monthly mean wind.

to quantify the impact of the scale-dependent re-centring for probabilistic skill of the ensemble. Figure 5 shows the fair Continuous Ranked Probability Score (fCRPS) of 500 hPa geopotential for the period June–August 2022 (92 start dates). Relative improvements in fCRPS were small (<0.5%) and consistent with slightly increased spread (not shown), at early lead times. Overall, the impact on probabilistic skill is very close to neutral.

Atmospheric winds exhibit variability that depends on the spatial scales and this can be quantified by looking at kinetic energy spectra. These show variances of wind anomalies as a function of wavenumber (i.e. inverse wavelength). With standard re-centring, the variance at initial time is considerably larger (up to 75% more) than at later lead times for all wavenumbers exceeding 50 (Figure 6). Most of the excess variance in the small scales dissipates relatively quickly during the first day of the forecast. In contrast, forecasts starting from initial conditions obtained with the scale-dependent EDA re-centring (with $N = 159$) exhibit large excess variance only in a band of wavenumbers between 50 and 160.

Implementation in operations

The construction of the ensemble initial conditions depends on the availability of the deterministic analysis and is therefore time critical. With the scale-dependent EDA re-centring, additional time is required for the new tasks that perform the spectral filtering. Without attention, this would have reduced the available time for running the ensemble forecast in the operational schedule. However, the additional time has been absorbed by other optimisations in the auxiliary code that is used to construct the ensemble initial conditions – optimisation work that has considerably eased the implementation of scale-dependent re-centring in Cycle 50r1.

Conclusion

The resolution upgrade of the EDA to TCo1279 (9 km) in Cycle 49r1 has exposed known limitations in the generation of ensemble forecast initial conditions. These manifest themselves in spurious mesoscale structures that can be seen for instance around tropical cyclones or in excess small-scale variance at initial time. Scale-dependent re-centring of the EDA on the deterministic analysis addresses these deficiencies in the mesoscale. Results from a case study for Tropical Cyclone Freddy in 2023 illustrate the marked improvement in the realism of the initial conditions and broader evaluation shows minimal impact on the probabilistic skill. The method also integrates efficiently into the operational schedule, facilitating its implementation in Cycle 50r1.

This development illustrates the continuous evolution of ensemble forecasting techniques, ensuring that higher-resolution data translates into more realistic, reliable forecasts without compromising integrity. As the EDA advances (e.g. if it could be run in real time and therefore have seen the latest observations), it might be possible to allow even more selective re-centring, further improving mesoscale realism while preserving ensemble forecast skill.

Further reading

Buizza, R., **M. Leutbecher & L. Isaksen.**, 2008: Potential use of an ensemble of analyses in the ECMWF Ensemble Prediction System. *Q. J. R. Meteor. Soc.,* **134**, 2051–2066. *https://doi.org/10.1002/qj.346*

Hólm, E., **M. Bonavita**, & **S. Lang**, 2022: Soft recentring Ensemble of Data Assimilations. *ECMWF Newsletter* **No. 171**.

Lang, S. T. K., **M. Bonavita** & **M. Leutbecher**, 2015: On the impact of re-centring initial conditions for ensemble forecasts. *Q. J. R. Meteor. Soc.,* **141**, 2571–2581. *https://doi.org/10.1002/qj.2543*

Roberts, C., **B. Ingleby**, **A. Geer**, **E. Hólm**, **M. Janousek**, **F. Prates** & **M. Rodwell**, 2024: IFS upgrade improves near-surface wind and temperature forecasts. *ECMWF Newsletter* **No. 181**, 16–25. *https://doi.org/10.21957/crx2bn4is8*

AI-DOP: an update on medium-range forecast scores

Eulalie Boucher, Mihai Alexe, Peter Lean, Ewan Pinnington, Patrick Laloyaux, Simon Lang, Tony McNally

Over the past year, ECMWF has continued its research on an end-to-end Artificial Intelligence (AI)-based Direct Observation Prediction (AI-DOP) model. In previous Newsletter articles (McNally et al., 2024 and McNally et al., 2025), we reported the first demonstration that medium-range forecasts using an AI-DOP concept could be produced. This update summarises the progress made during 2025 and outlines the next steps for this rapidly evolving research track.

Machine-learned (ML) weather prediction models typically derive their skill from reanalysis datasets such as ERA5. These datasets, however, rely on traditional approaches of incorporating observations into numerical weather prediction (NWP). Although modern data assimilation (DA) systems are remarkably sophisticated, they still struggle to extract the full information content from the global observing system. Simplified forward operators between measured radiances and geophysical variables, limitations in representing uncertainties, and the volume of satellite observations mean that some of the potential value in the observations remains under-utilised.

As a result, ML models trained solely on reanalysis inherit the constraints of the upstream DA system, rather than the richness of the observations themselves. AI-DOP was initiated to explore whether AI models could learn directly from observations, either alongside reanalysis, in place of parts of the DA pipeline, or in an end-to-end system and thereby unlocking information that current systems cannot fully exploit. Over the past year, it has become increasingly clear that integrating observations more directly into AI-based forecasting systems may be central to advancing forecast skill. This makes the AI-DOP research track an important step toward the next generation of observation-aware weather prediction models.

Current model architecture

GraphDOP (Alexe et al., 2024) is an AI-DOP model based on a graph neural network with an encoder-processor-decoder architecture (Figure 1). While its design shares

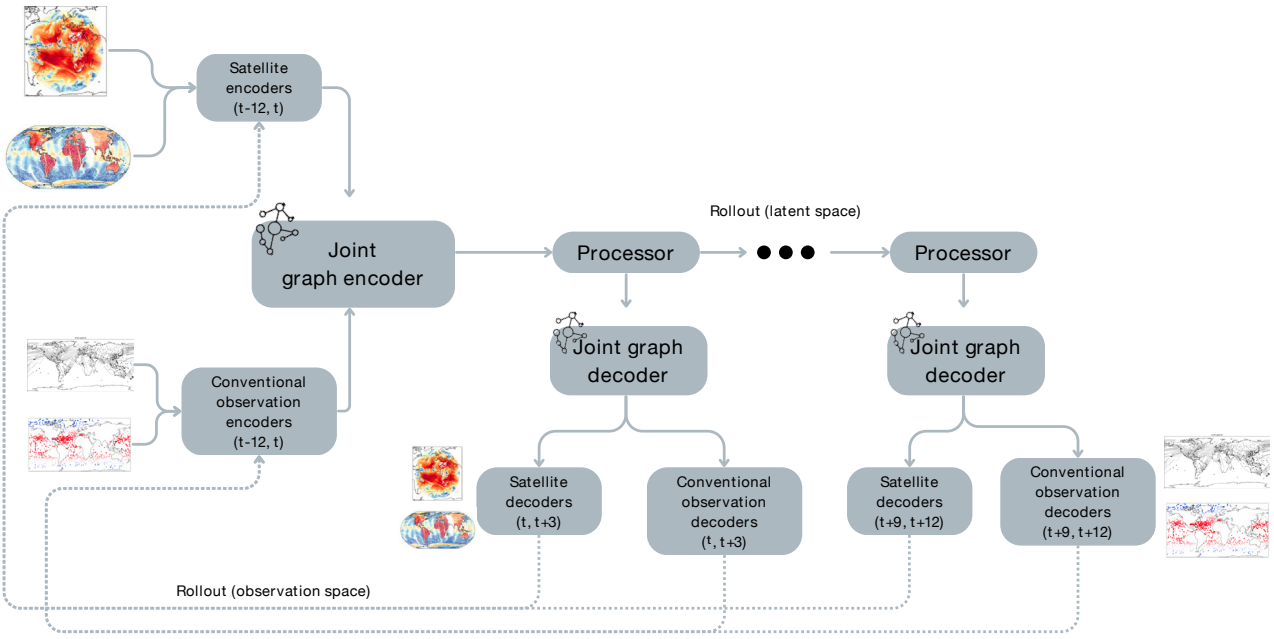

FIGURE 1 GraphDOP architecture with three-hour autoregressive time-stepping (rollout) in the latent space. During finetuning, the forecasted observations are fed back as inputs to produce the next forecast (rollout in observation space).

doi: 10.21957/6jvc-5m45

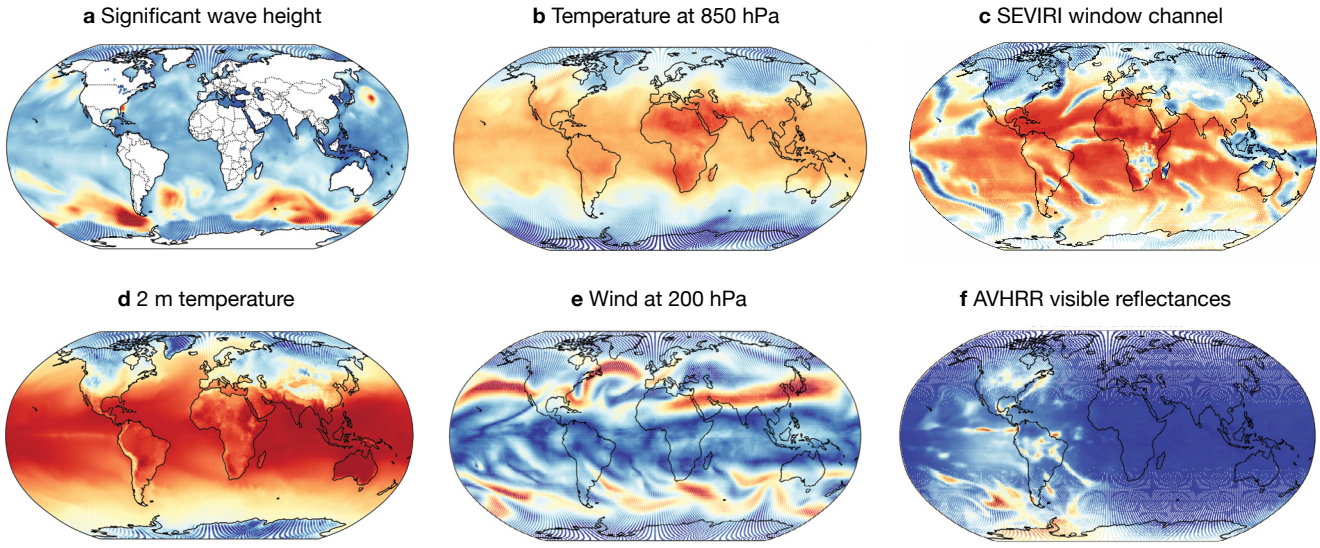

FIGURE 2 Example of gridded forecasts at 12-hour lead time of different parameters and satellite channels produced from the GraphDOP model.

many aspects with ECMWF's Artificial Intelligence Forecasting System (AIFS) (Lang et al., 2024), GraphDOP's architecture is adapted specifically for observation prediction, operating solely in observation space, without gridded climatology or NWP (re)analysis inputs or feedback.

A graph-based encoder projects the 12 hours of input observations onto a latent space. Subsequently, the processor (a sliding-window transformer) evolves the latent space representation throughout a 12-hour target observation window. Finally, a decoder projects the latent representation back into observation space to produce the model forecasts. The latent representation is defined on an O96 reduced Gaussian grid (a spatial resolution of ca. 1 degree), with 1,024 features. The encoder graph mapping is constructed on the fly from the input observations available in each training and validation batch to their nearest grid points in the latent space representation, a k-nearest neighbour graph, with k=1. Conversely, the decoder graphs connect each target observation to its three nearest neighbour nodes on the latent mesh.

The latest version of the model forecasts 12-hour observation windows by auto-regressively processing and decoding target observations across four consecutive three-hour chunks, a procedure referred to in Figure 1 as latent-space rollout. This approach has improved the sharpness and accuracy of the GraphDOP forecasts compared to the original model that decoded all the observations inside the 12-hour target window in a single step. Currently, the training objective is the weighted mean square error.

Once trained, GraphDOP is able to forecast any of the seen observation types, either at their true location or at any point in time and space, producing global forecasts of any observed parameter or satellite channel (Figure 2).

A series of experiments has been conducted to provide evidence that the GraphDOP model develops internal representations of the Earth system state, structure and dynamics, as well as the characteristics of different observing systems (Lean et al., 2025). GraphDOP simultaneously embeds information from diverse observation sources spanning the full Earth system into a shared latent space. This enables predictions that implicitly capture cross-domain interactions in a single model, without the need for any explicit coupling (Boucher et al., 2025).

Growing set of observations

A key advantage of a purely observation-driven system such as GraphDOP is its ability to ingest a wide range of observational data, including many that are not currently used in traditional DA systems. In parallel with model development, we have been steadily expanding our collection of ML-ready observational datasets for training, spanning both in-situ (conventional) measurements and satellite observations (Figure 3).

Not all curated datasets have been used thus far in training, and the performance results presented in the next section reflect only a subset of them. Nevertheless, we are gradually evaluating and incorporating additional datasets into our standard configuration as the system matures. Quality control tools initially developed for 4D-Var have been implemented in GraphDOP to better quantify the importance of the different instruments (Laloyaux et al., 2025).

In addition, work has begun on developing a real-time workflow built on top of existing observation-receiving infrastructure, supporting a potential future operational end-to-end observation-to-forecast system.

Medium-range forecast scores

The model updates combined with an enhanced set of observations have led to large improvements in the medium-range forecast scores. GraphDOP is especially skilful at forecasting surface parameters, and its performance is steadily improving for upper-air parameters, now matching the physics-based Integrated Forecasting System (IFS) at 12-hour lead time. In Figures 4 and 5, we show root-mean-square errors (RMSE) of gridded forecasts against in-situ observations for a summer month (June 2022) and compare those of GraphDOP to the IFS at the same resolution. Similar evaluations for winter months show higher errors, which we believe can be attributed to a harder predictability but also to a lack of observations representing winter conditions, such as snow observations.

Figure 4 illustrates these improvements for upper-air parameters at 850 hPa, 500 hPa, and 200 hPa (forecasted on an O96 grid) in the northern hemisphere extratropics (top row) and in the tropics (bottom row). In the extratropics and for all three levels, GraphDOP and the IFS exhibit very similar skill in the short range, with differences generally within the sampling uncertainty during the first 24–36 hours. Beyond this point, the IFS retains an advantage, particularly for the dynamically sensitive 500 hPa geopotential height and 200 hPa wind speed, where its errors grow more slowly through the medium range. Nevertheless, GraphDOP tracks the IFS closely throughout the ten-day period. To provide additional context, the figure also includes a dashed line indicating the performance of an older version of GraphDOP. Over the tropics, GraphDOP does especially well for 850 hPa temperature and 500 hPa geopotential in the first five days. This comparison highlights that while upper-air forecast skill continues to be led by the physical model, GraphDOP is making steady progress, supported, for example, by the improved training configuration introduced recently.

Figure 5 presents corresponding results for key surface parameters, showing scores for 2 m temperature and 10 m winds forecasted on an N320 grid, over the northern hemisphere extratropics. For both parameters, GraphDOP delivers substantially improved short-range performance. In the case of 2 m temperature, GraphDOP begins with a markedly lower error (1.8 K compared with 2.4 K for the IFS at the same resolution), and preserves this advantage through to day 4, at which point the two systems converge. Ongoing work focuses on improving the rollout strategy to curb error growth into the medium range, with the aim of extending GraphDOP's short-

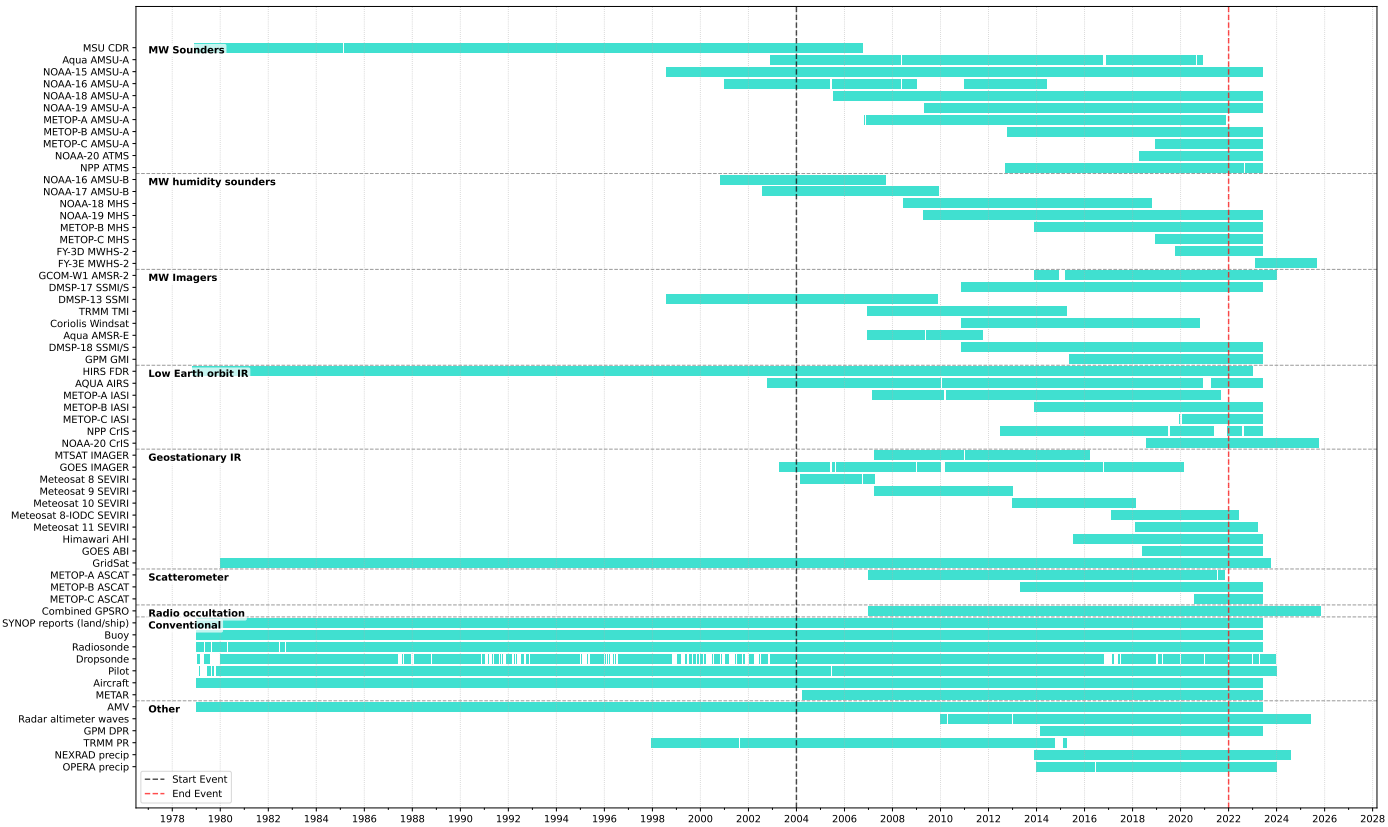

FIGURE 3 Observation types that have so far been converted into ML-friendly training datasets.

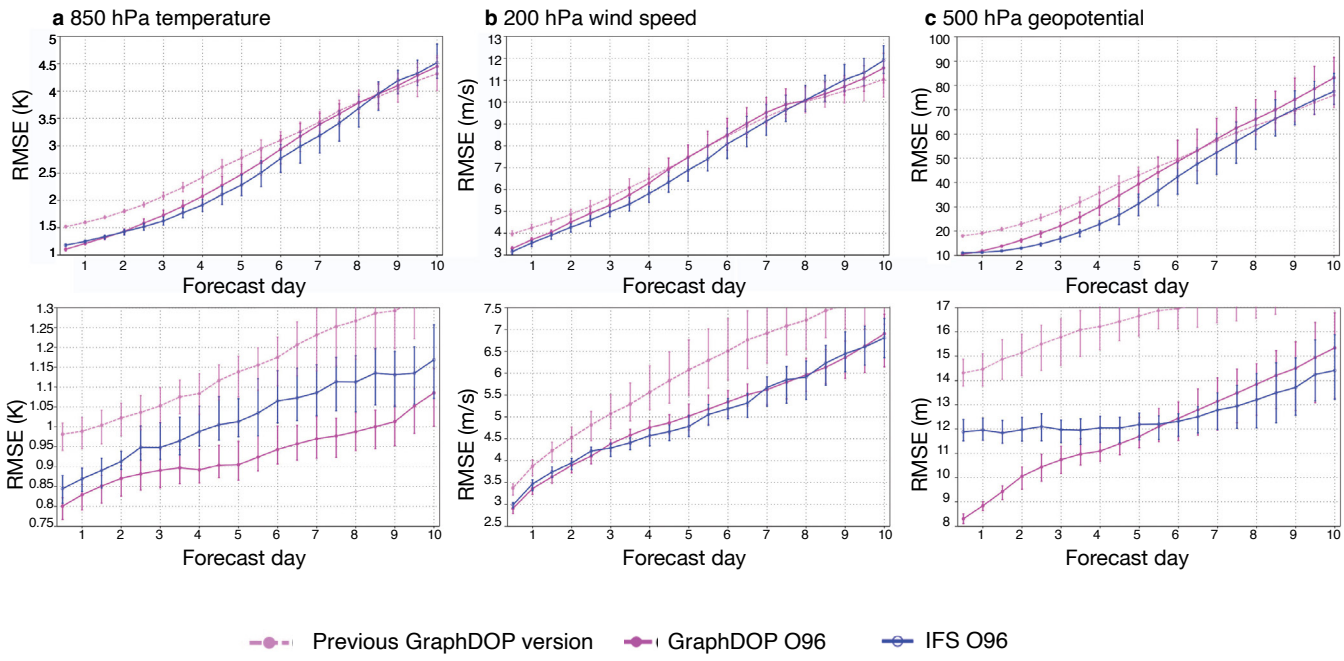

FIGURE 4 Root-mean-square error (RMSE) of (a) 850 hPa temperature, (b) 200 hPa wind speed and (c) 500 hPa geopotential forecasts by GraphDOP (magenta) and the IFS (blue) over the northern hemisphere extratropics (top row) and the tropics (bottom row) for June 2022. Scores are computed against observations.

range advantage. The largest gains appear for 10 m winds, where GraphDOP maintains significantly lower RMSE out to day 10. Together, these results suggest that the implicit, observation-driven, coupling embodied in GraphDOP – where near-surface conditions are constrained directly rather than through parametrized land–atmosphere exchange pathways – yields particularly strong benefits for surface variables.

Conclusions and outlook

The progress achieved over the past year demonstrates that the AI-DOP research track is maturing rapidly, both in terms of model capability and in the breadth of the observational information it can exploit. As we continue to expand the range of ML-ready observations, several key observing systems stand out as high-priority candidates.

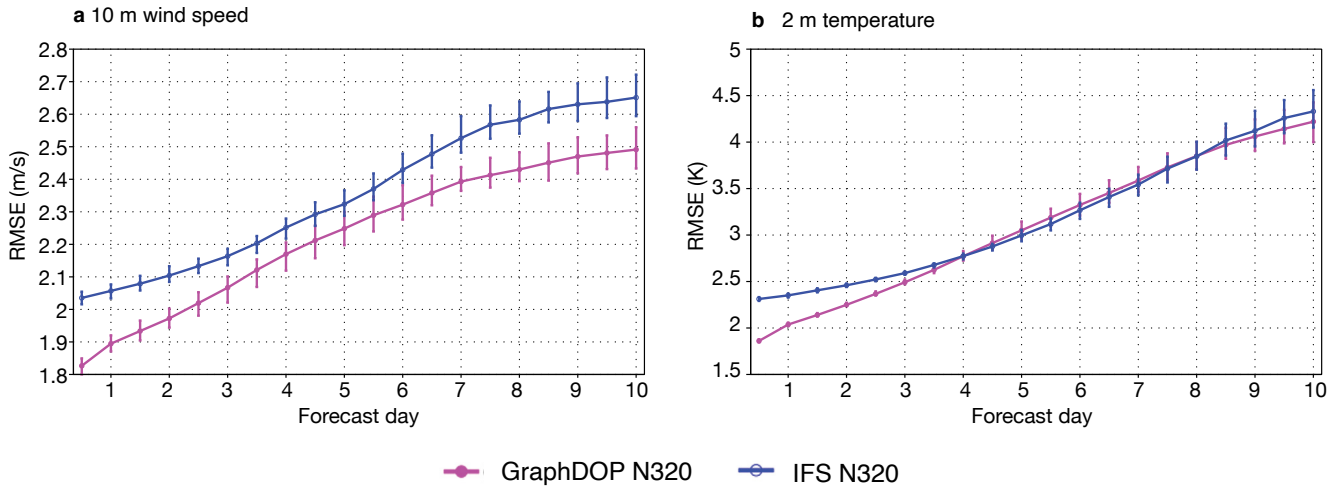

FIGURE 5 Root-mean-square error (RMSE) of (a) 10 m wind speed and (b) 2 m temperature forecasts by GraphDOP (magenta) and the IFS (blue) over the northern hemisphere extratropics for June 2022. Scores are computed against observations.

These include GNSS radio-occultation (GNSS-RO) data, in which the bending angles have been reprocessed for ML use. This includes more sophisticated normalisation and quality control of the bending angles, allowing the model to pick up a useful signal. This is showing to be particularly valuable in current prototype experiments. At the modelling level, several promising directions of research are being explored to allow the network to benefit from longer sequences of observations and to reduce error growth in the rollout.

Overall, the results presented here highlight the growing potential of observation-driven AI models to complement the capabilities of traditional NWP.

Further reading

Alexe, M., **E. Boucher**, **P. Lean**, **E. Pinnington**, **P. Laloyaux, A. McNally et al.**, 2024: GraphDOP: Towards skilful data-driven medium-range weather forecasts learnt and initialised directly from observations. *arXiv* 2412.15687. *https://doi.org/10.48550/arXiv.2412.15687*

Boucher, E., **M. Alexe**, **P. Lean**, **E. Pinnington**, **S. Lang**, **P. Laloyaux et al.**, 2025: Learning Coupled Earth System Dynamics with GraphDOP. *arXiv* 2510.20416. *https://doi.org/10.48550/arXiv.2510.20416*

Laloyaux, P., **M. Alexe**, **E. Boucher**, **P. Lean**, **E. Pinnington**, **S. Lang et al.**, 2025: Using data assimilation tools to dissect GraphDOP. *arXiv* 2510.27388. *https://doi.org/10.48550/arXiv.2510.27388*

Lang, S., **M. Alexe**, **M. Chantry**, **J. Dramsch**, **F. Pinault**, **B. Raoult et al.**, 2024: AIFS – ECMWF's data-driven forecasting system. *arXiv* 2406.01465.*https://doi.org/10.48550/arXiv.2406.01465*

Lean, P., **M. Alexe**, **E. Boucher**, **E. Pinnington**, **S. Lang**, **P. Laloyaux et al.**, 2025: Learning from nature: insights into GraphDOP's representations of the Earth System. *arXiv* 2406.01465. *https://doi.org/10.48550/arXiv.2508.18018*

McNally, T., **C. Lessig**, **P. Lean**, **M. Chantry**, **M. Alexe** & **S. Lang**, 2024: Red Sky at night... producing weather forecasts directly from observations. *ECMWF Newsletter* **No. 178**, 30–34. *https://doi.org/10.21957/tmc81jo4c7*

McNally, A., **C. Lessig**, **P. Lean**, **E. Boucher**, **M. Alexe**, **E. Pinnington et al.**, 2025: An update on AI-DOP: skilful weather forecasts produced directly from observations. *ECMWF Newsletter* **No. 182**, 15–19. *https://doi.org/10.21957/tmi6y913dc*

Assimilating aerosol visible reflectances for improved air quality forecasts

Samuel Quesada-Ruiz, Cristina Lupu, Tobias Necker, Roberto Ribas, Volkan Firat, Leonhard Scheck (DWD), Christina Köpken-Watts (DWD), Angela Benedetti

As a part of a collaboration with the German Meteorological Service (DWD), in addition to the efforts to monitor and assimilate cloud-sensitive visible data (Necker et al., 2025), ECMWF is also advancing the use of aerosol visible reflectance data. This work will enhance the initialisation of aerosol forecasts within the Integrated Forecasting System's (IFS) four-dimensional variational data assimilation (4D-Var) in the atmospheric composition configuration (IFS-COMPO), used operationally in the Copernicus Atmosphere Monitoring Service (CAMS). The aerosol visible reflectance assimilation system (AVRAS) prototype was developed under the CAMS EvOlution project (CAMEO) and followed the approach of Benedetti et al., 2020. This system currently leverages Level-2 cloud-screened aerosol visible reflectance observations from the Moderate-resolution Imaging Spectro-radiometer (MODIS) instrument aboard the Aqua and Terra satellites, which are already operationally received at ECMWF. Future work could extend to VIIRS or similar products developed for other imagers. The first assimilation experiment has been successfully conducted over several cycles, with promising preliminary results. This feature article summarises major findings from this exciting endeavour.

Essential steps that paved the way

Monitoring atmospheric composition is a key objective of Copernicus, the European Union's flagship Earth observation initiative. CAMS provides free and continuous data and information on atmospheric composition, supporting air quality monitoring, progress towards sustainable development goals, and the transition to sustainable energy. The CAMEO project was funded to enhance the quality and efficiency of the CAMS service and to strengthen its ability to respond to policy needs. Within CAMEO, efforts have concentrated on implementing a fast observation operator that ensures the assimilation of replace with aerosol visible reflectances can be performed within the time constraints of operational IFS cycles.

Infrared and microwave satellite radiances have long been assimilated into ECMWF's IFS to help estimate the best possible initial conditions for global numerical weather prediction (NWP). Visible observations, on the other hand, have largely been underexploited for data assimilation. This is due to the heavy computational cost associated with resolving complex interactions between radiation, clouds, aerosol particles and the reflecting surface at visible wavelengths. As a result, direct assimilation of reflectances to estimate aerosols was not feasible until recently. Instead, the aerosol analysis has been constrained by assimilating aerosol optical depth (AOD) products.

The first attempt to assimilate aerosol visible reflectances at ECMWF was made at ECMWF within the Aerosol Radiance Assimilation Study (ARAS) project, a European Space Agency-funded project that ran from 2018 to 2020. As part of ARAS, an observation operator for visible reflectances based on look-up tables was developed by the Oxford-RAL Aerosol and Cloud team and incorporated into the IFS, although it did not reach operational maturity.

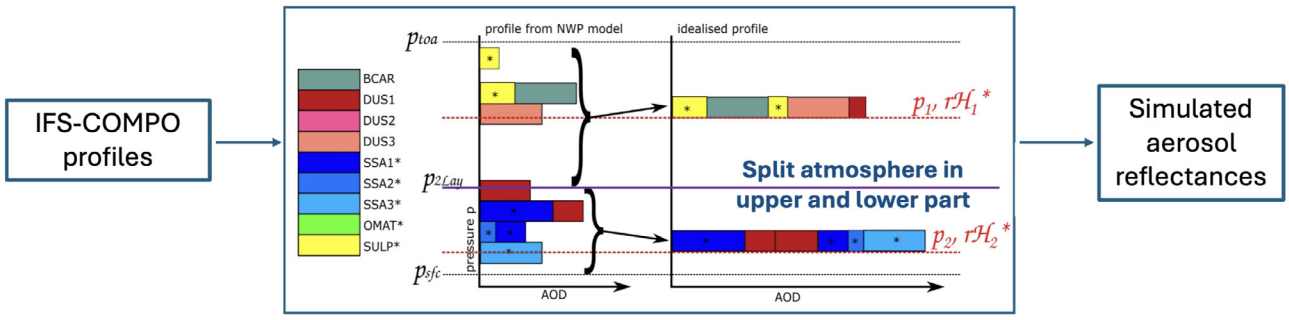

FIGURE 1 Schematic view of an NWP model profile and the corresponding aerosol profile simplification in MFASIS-Aerosol in which only two layers are filled with aerosols. * Refers to hydrophilic.

doi: 10.21957/6tq7-xr22

a MFASIS-Aerosol in the IFS-COMPO

The MFASIS-Aerosol version integrated into the upcoming RTTOV-14.1 release takes as input vertical profiles of atmospheric pressure, temperature, humidity and aerosols, along with information about the surface, and it outputs top-of-atmosphere visible reflectances (see Figure 1). Adjoint and tangent-linear codes are also provided. Aerosol content from the original profile is represented in a simplified profile by concentrating the CAMS aerosol species defined in RTTOV (i.e. sea salt, desert dust, organic matter, black carbon, sulphate) into two distinct layers. This approach captures boundary-layer aerosols and a middle-to-upper tropospheric plume while preserving the total vertical integral of the aerosol loading.

A more recent CAMS global atmospheric composition profile dataset was produced in 2025 (Turner, 2025). It incorporates all previous variables as well as some new ones, specifically three aerosols: ammonium, nitrate and secondary organic matter. Future re-training of MFASIS-Aerosol will make use of the full set of CAMS-species to improve the aerosol reflectance modelling.

Benefits of integrating RTTOV MFASIS-Aerosol in the IFS-COMPO include:

- **Enhanced aerosol representation.** Visible reflectances are sensitive to aerosol size, and composition, adding information beyond the standard AOD products.
- **Better aerosol initial conditions.** These improve both CAMS air-quality analyses and NWP forecast skill through more accurate shortwave radiation and surface energy balance.
- **More consistent assumptions.** The approach is compatible with the IFS-COMPO's 4D-Var, unlike those used in external satellite retrievals.
- **Multi-sensor capability.** This enables the joint assimilation of visible reflectances from multiple sensors, exploiting their combined information content and reducing limitations tied to individual retrieval algorithms, viewing geometries, or missing channels.

ARAS results showed that the assimilation of aerosol visible reflectances increased the aerosol load in the analysis to a level comparable to the MODIS aerosol optical depth data (Benedetti et al., 2020), while also improving other aerosol parameters.

A fast radiative transfer model for visible wavelengths

Realistic simulation of visible reflectances in the presence of aerosols requires a fast radiative transfer operator capable of accurately and efficiently simulating top-of-atmosphere visible reflectances. Traditional radiative transfer solvers for aerosols, such as the Discrete Ordinate Method (DOM), are highly accurate but computationally expensive and impractical for large-scale data assimilation systems that require repeated simulations with the forward, tangent linear, and adjoint observation operators.

A major advance in this area is the Method for Fast Satellite Image Synthesis (MFASIS), developed by DWD and the Ludwig Maximilian University of Munich within the framework of the Hans Ertel Centre for Weather Research. The MFASIS is implemented in the Radiative Transfer for TOVS (RTTOV, currently at version 14.0; Saunders et al., 2018) within the EUMETSAT NWP SAF. It provides a fast and accurate approximation of the one-dimensional radiative transfer model DOM solution for simulating visible reflectances in the presence of multiple scattering due to clouds and aerosols. The MFASIS is under continuous development to further increase accuracy and to extend its capabilities. Examples of this are the introduction of a neural-network-based solver for cloudy visible reflectances (MFASIS-Cloud; Scheck, 2021), and the application of MFASIS for the direct assimilation of aerosol visible reflectances, where aerosol scattering dominates (MFASIS-Aerosol, Box A).

Handling cloud sensitivity in visible-channel aerosol assimilation

Assimilating aerosol visible reflectances requires careful treatment of cloud sensitivity, as clouds strongly influence observed reflectances and can introduce significant biases if not properly accounted for. Three potential approaches for handling cloud sensitivity in visible reflectance assimilation have been identified:

1. **Apply in-house cloud screening to Level-1 reflectances.** This approach relies on a dedicated cloud-detection step to remove visible reflectance observations affected by clouds before assimilation can be performed. However, as Level-1 visible reflectances usually do not include flags to distinguish and separate clouds from aerosols, relying only on model cloud fields can lead to misinterpreting observed cloud signals as aerosol ones and to a degradation in the analysis.
2. **Simultaneous assimilation of clouds and aerosols.** A more advanced option would be to assimilate cloud and aerosol properties together using a single, combined observation operator. While this method

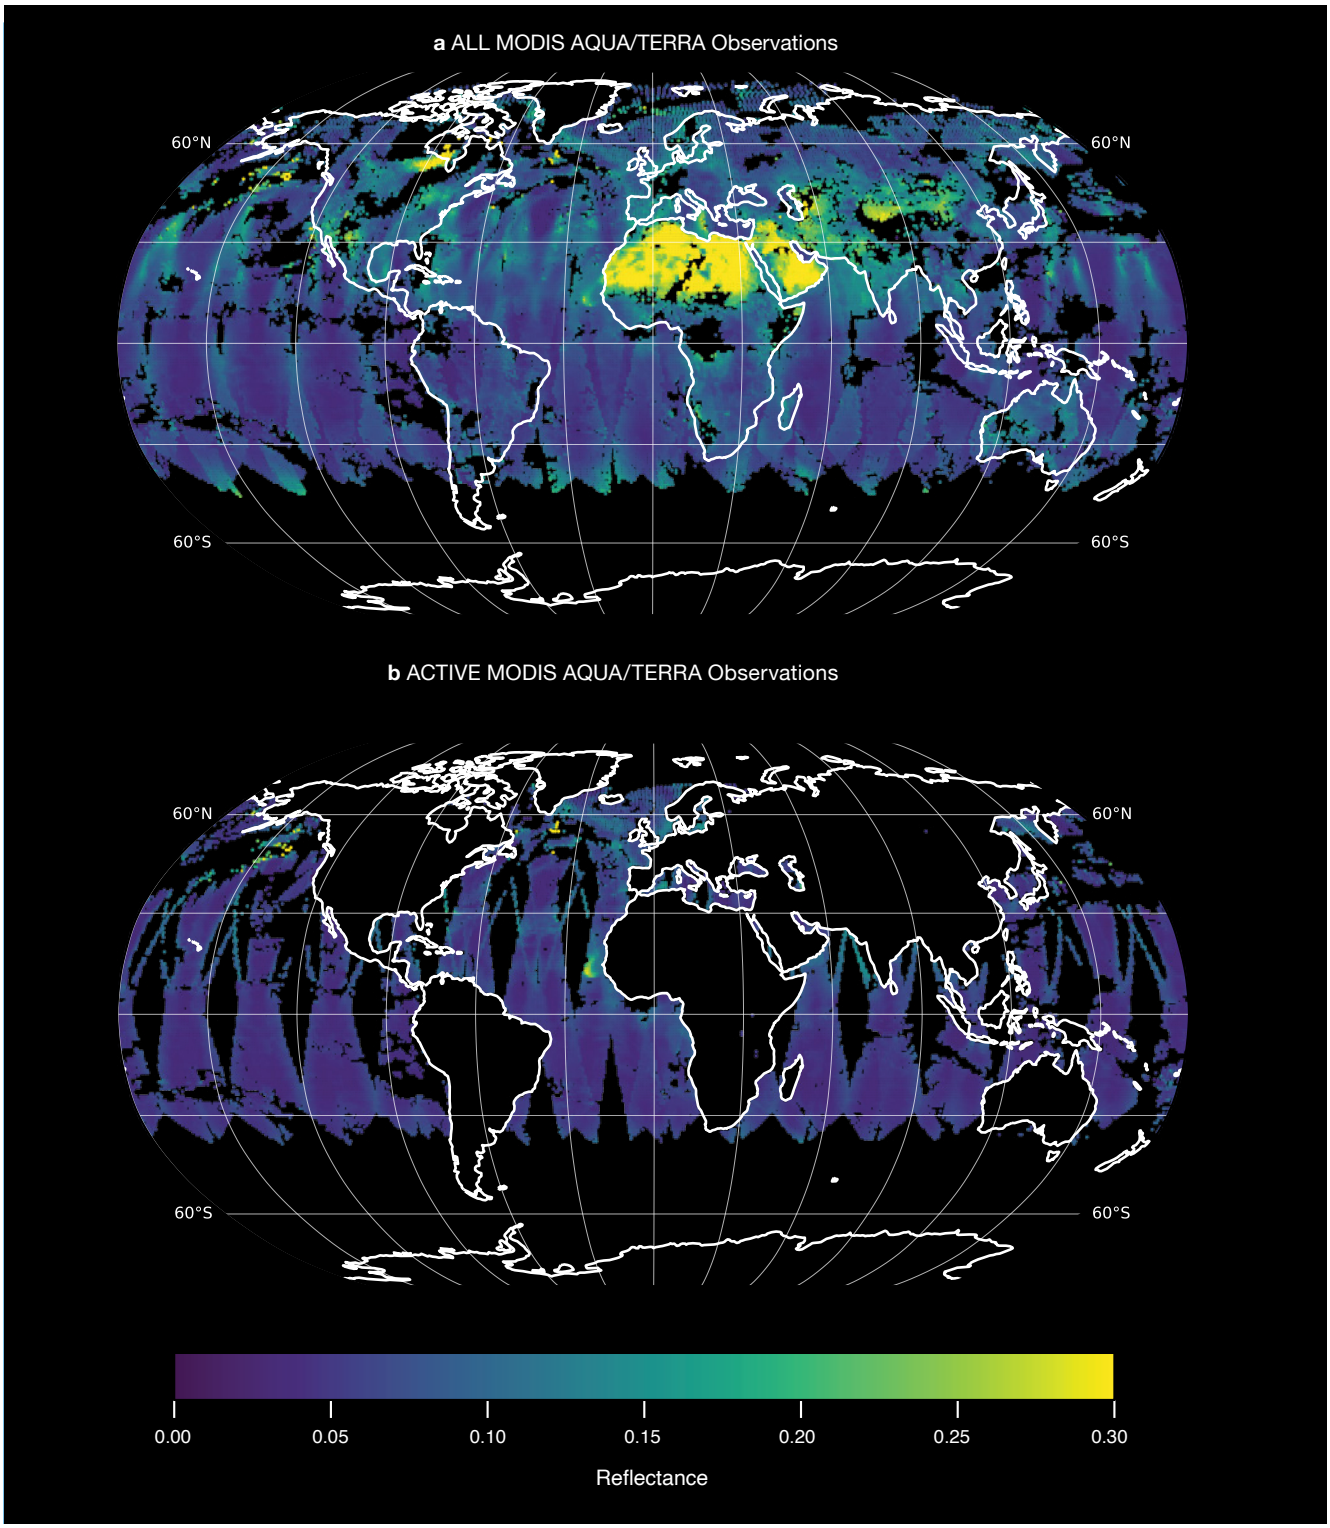

FIGURE 2 MODIS on Aqua and Terra 665 nm visible reflectance observations on 4 June 2025 (a) before and (b) after quality control.

would allow the assimilation system to account for both aerosol and cloud contributions to the observed reflectance, in practice, this is not currently possible. In the future, when a unified MFASIS Cloud and Aerosol operator will be available, this joint-assimilation approach could be revisited.

3. **Use Level-2 cloud-screened observations.** The most practical approach involves the use of pre-processed Level-2 cloud-screened aerosol reflectances from instruments like MODIS (Aqua/Terra) or VIIRS (NOAA) on polar-orbiting satellites. These products apply established cloud-screening algorithms, greatly reducing cloud aliasing risks and supporting operational implementation. This approach was selected for the CAMEO project. MODIS Collection 6.1 Level-2 observations were used since they are operationally available in IFS-COMPO, but future work

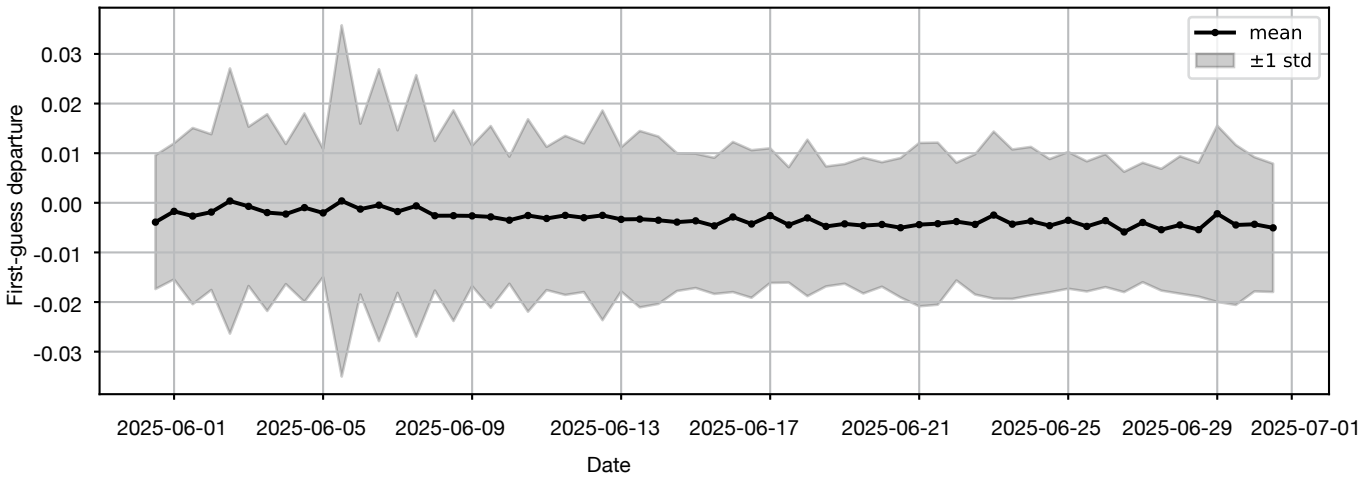

FIGURE 3 Time series of first-guess departures for MODIS 665 nm observations on Aqua/Terra during June 2025.

could extend to VIIRS or similar products developed for other imagers.

Observation processing

A bespoke workflow was established to pre-process aerosol visible reflectance data, facilitating their integration within the IFS and enabling their monitoring and assimilation. MODIS reflectances at 665 nm are extracted from the Level-2 cloud-screened aerosol reflectance sequence dedicated to the AOD processing and converted into a local template, dedicated to the visible reflectance processing. The volume of visible reflectance data available for monitoring is substantial, with datasets arriving at their native high spatial and temporal resolution. A critical issue in using visible data in NWP models is the resolution mismatch between high-resolution satellite observations and coarser model grids used for the analysis. MODIS visible reflectances were experimentally superobbed at a resolution of 80 km to reduce data volume, as well as to ensure consistency with the model at the analysis scale and to minimise the impacts of possible horizontal correlations on the observation error. Screening mechanisms are applied to exclude data that may degrade analysis quality, especially for areas affected by ice, snow, or observations at extreme sun or satellite angles. The final number of observations retained depends strongly on the chosen superobbing strategy, the stringency of screening criteria, and the temporal sampling frequency. In the current setup, following pre-processing and superobbing, a single channel from each MODIS instrument typically yields around 15,000 active observations per day. The data selection and the impact of the quality control are illustrated in Figure 2.

Evaluation of monitoring and assimilation results

Assimilation of AOD in the IFS-COMPO 4D-Var has been operational since 2008. The aerosol variable that is adjusted in the analysis is the total aerosol mixing ratio, representing the aerosol load in the atmospheric column. Other variables that are adjusted in the analysis are temperature, humidity, winds, and surface pressure. The aerosol observations are only indirectly related to those meteorological variables. Therefore, when using AOD or aerosol reflectance, the biggest impact is on the aerosol load. Moreover, the current configuration of the IFS-COMPO 4D-Var aerosol analysis does not allow the extraction of information on winds via the so-called “tracing effect”, which is connected to the transport of a given species by the model winds. This effect is active for humidity and ozone, but not for aerosols. However, because the 4D-Var leverages on the dynamical transport model to produce a short-range forecast of aerosol mixing ratio over the 12-hour assimilation window, the effect of assimilating aerosol observations over one location (i.e. over sea) can also be felt away from that location (i.e. over land). This is a unique feature of assimilation systems which are four-dimensional (4D) as they include a time evolution of the fields which are being assimilated.

The timeseries of first-guess departures (Figure 3) and the first-guess departure mean map (Figure 4a) show a small negative bias over ocean in aerosol-clear scenes. In contrast, the presence of aerosols introduces a positive bias, indicating that the aerosol loading influences the reflectance values. Additional regional biases are observed around 50°S, likely associated with strong surface winds. These conditions can alter surface reflectance and aerosol distribution, contributing to localised discrepancies between observations and model equivalents.

The impact of the visible reflectance assimilation is shown in Figure 4. The analysis departure mean map (Figure 4b) exhibits smaller values than the first-guess departures, especially in areas where higher aerosol concentrations are present. Departure reduction (Figure 4c) is negative across the globe, meaning the analysis is closer to the observation than the first guess. A significant reduction

can be seen in the Gulf of Guinea, north Atlantic and the west-coast of Africa.

Verification against ground-based observations

Figure 6 depicts the temporal evolution of the mean bias and the root mean square error (RMSE) between

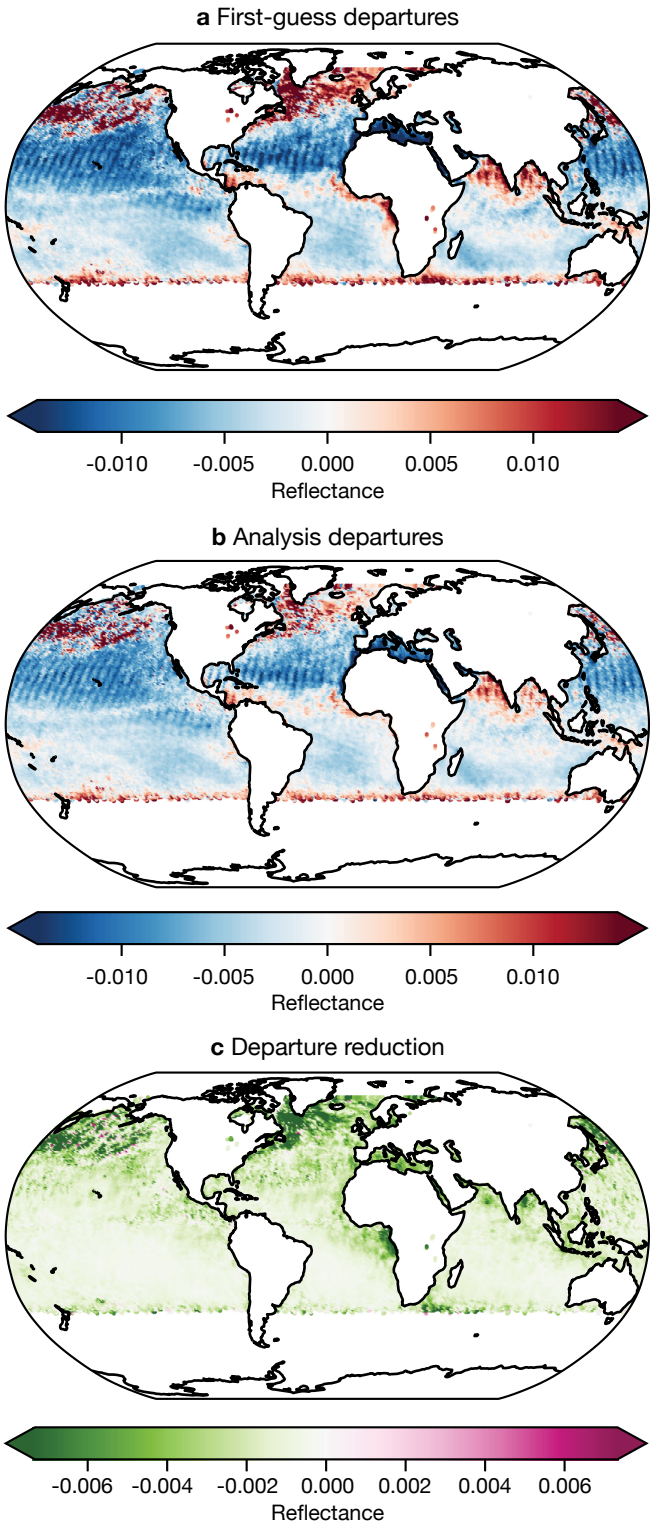

FIGURE 4 Three-panel comparison illustrating (a) the first-guess departures; (b) analysis departures; and (c) departure reduction for June 2025.

case study

Dust event on 4 June 2025

A Saharan Desert dust event on 4 June 2025 provides an illustration of the impact of assimilating visible reflectances (Figure 5). The presence of a large aerosol load, confirming the spatial extent of the dust outbreak, is clearly seen in the image captured by the Flexible Combined Imager (FCI) instrument onboard Meteosat-12. Visible reflectance observations from MODIS confirm the presence of a strong dust plume over the region. The model-simulated reflectances capture the aerosol signal in the same area but underestimate its intensity. Assimilation increases aerosol concentrations in the analysis, enhancing the event representation. The assimilation of the reflectances brings clear benefits in this case, highlighting the ability of the assimilation system to improve aerosol representation and reducing biases when visible reflectances are assimilated.

forecasted and observed AOD at 550 nm from the global AErosol RObotic NETwork (AERONET) network (Holben et al., 1998; *https://aeronet.gsfc.nasa.gov/*). Although the network is spatially sparse in some regions and most stations are located over land, it provides high-quality reference measurements of aerosol optical properties and is widely regarded as the state-of-the-art for aerosol validation. Three experiments are compared: a control experiment without AOD or aerosol visible reflectance assimilation; a aerosol visible reflectances assimilation experiment which only includes MODIS Level-2 cloud-screened visible reflectances from one single wavelength (665 nm); and a baseline experiment which includes the assimilation of AOD from the various sensors assimilated operationally in CAMS (e.g. MODIS, Polar Multi-Sensor Aerosol Product (PMAp), VIIRS). The assimilation of visible reflectances exhibits a smaller bias for most of the period compared to the AOD assimilation. The AOD assimilation experiment shows the smaller RMSE for most of the period. Visible reflectance assimilation shows, in general, lower RMSE than the control experiment, even improving upon the AOD experiment on specific days (e.g. on 6 June 2025).

Outlook and next steps

This article documents substantial progress in integrating aerosol-sensitive visible reflectances within the CAMS analysis framework. Historically limited by complex radiative transfer and aerosol modelling challenges, recent innovations, particularly the introduction of the MFASIS-Aerosol fast radiative transfer operator developed by DWD, have made

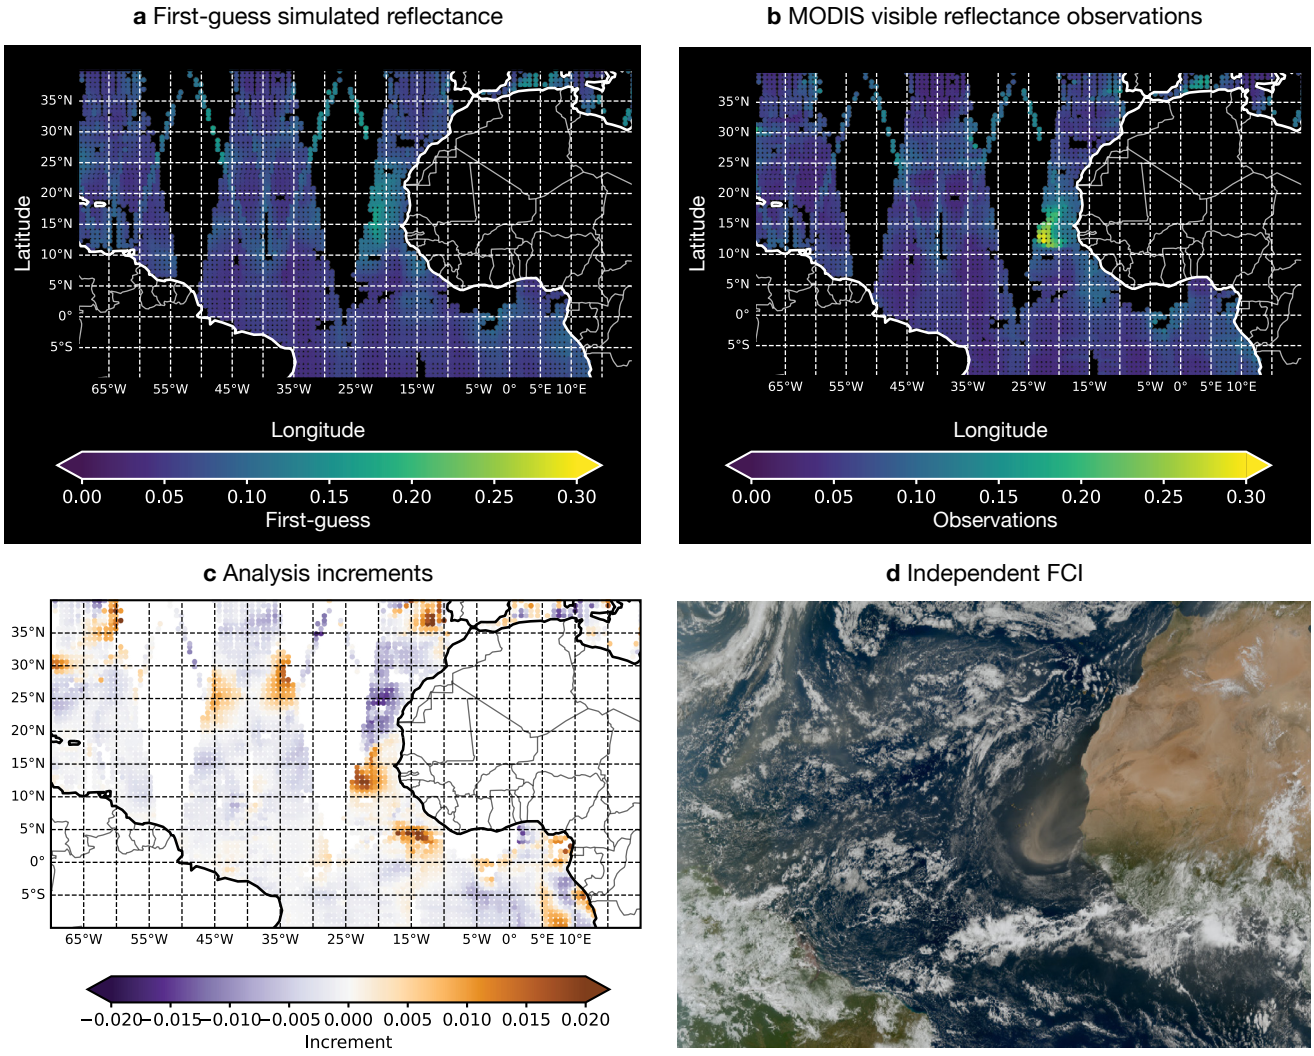

FIGURE 5 Four-panel comparison illustrating the assimilation of MODIS 665 nm visible reflectances: (a) first-guess simulated reflectance; (b) MODIS visible reflectance observations; (c) analysis increments; (d) independent FCI image highlighting a dust outbreak on 4June 2025 (Credit: EUMETSAT; *https://www.eumetsat.int/saharan-dust-over-atlantic-ocean*).

operational assimilation feasible. While previous studies served primarily as demonstrators, this study has introduced the first implementation of aerosol visible observations in the IFS-COMPO. Directly assimilating

b Next steps towards operational application

The research conducted highlights the potential of direct visible reflectance assimilation for improved aerosol analysis and forecasts, while also identifying areas requiring further development:

1. Scientific evaluation and expansion. Perform a comprehensive scientific evaluation using the MODIS MFASIS-Aerosol parameters for relevant periods, focusing on high-impact events such as dust outbreaks. Explore the integration of additional sensors, for example VIIRS or the Ocean and Land Colour Instrument (OLCI) on Sentinel 3A and 3B, subject to the availability of a Level-2 cloud-screened reflectance product.

2. Multi-channel extension. Extend monitoring and assimilation capabilities beyond the 0.6 µm channel, incorporating additional visible wavelengths to exploit the full information content of satellite reflectances.

3. Refinement of the assimilation system. Improve quality control, implement advanced screening techniques to ensure robust observation selection, use a bias correction to address systematic biases identified in the first guess departures, and improve observation and background error characterisation to optimise assimilation impact.

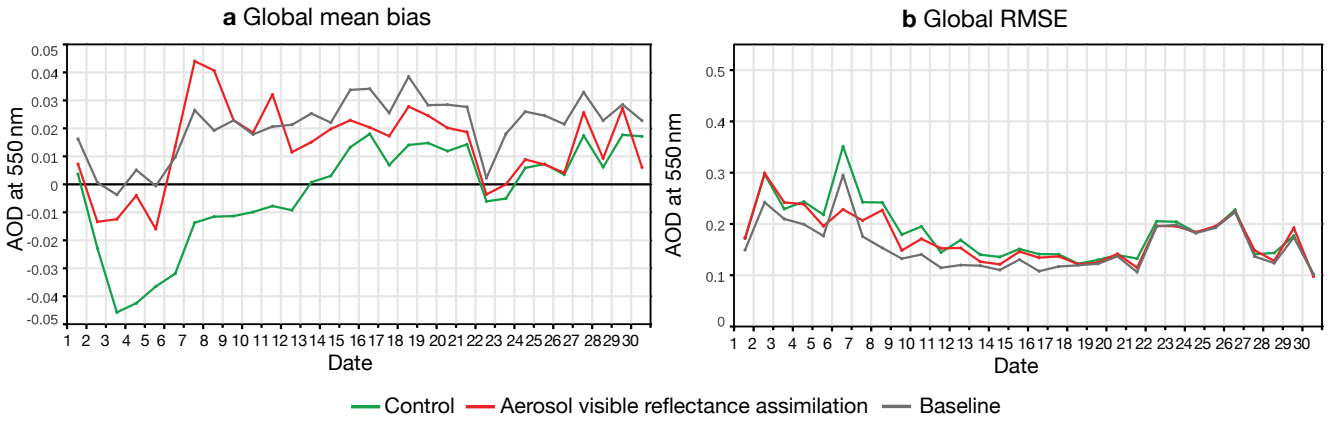

FIGURE 6 Time series of (a) the global mean bias and (b) the global root mean square error between forecasted and observed AOD at 550 nm, calculated over the AERONET network for June 2025. Each line corresponds to a distinct experiment: Control (without aerosol assimilation, green), visible aerosol reflectance assimilation (red) and Baseline (AOD assimilation, grey), enabling intercomparison of forecast skill.

observed aerosol-affected reflectances into ECMWF's assimilation system is anticipated to enhance CAMS aerosol forecasts and atmospheric composition analyses. The methodological advancements constitute a robust basis for sustained innovation and the prospective operational implementation of aerosol-sensitive visible reflectance analyses.

Further reading

Benedetti, A., **S. Quesada Ruiz**, **J. Letertre-Danczak**, **M. Matricardi** & **G. Thomas**, 2020: Progress towards assimilating visible radiances. **ECMWF Newsletter No. 162**. *https://www.ecmwf.int/en/newsletter/162/news/progress-towards-assimilating-visible-radiances*

Holben, B. N., **T. F. Eck**, **I. Slutsker**, **D. Tanré**, **J. P. Buis**, **A. Setzer et al.**, 1998: AERONET - A Federated Instrument Network and Data Archive for Aerosol Characterization. **Remote Sensing of Environment**, Volume **66**, Issue 1, 1–16, ISSN 0034-4257. *https://doi.org/10.1016/S0034-4257(98)00031-5*

Necker, T., **C. Lupu**, **S. Quesada-Ruiz**, **V. Firat**, **L. Scheck** & **A. Benedetti**, 2025: Visible radiances in ECMWF's analysis. **ECMWF Newsletter No. 184**, 17–24. *https://doi.org/10.21957/tn61af97px*

Saunders, R., **J. Hocking**, **E. Turner**, **P. Rayer**, **D. Rundle**, **P. Brunel et al.**, 2018: An update on the RTTOV fast radiative transfer model (currently at version 12). **Geosci. Model Dev.**, **11**, 2717–2737. *https://doi.org/10.5194/gmd-11-2717-2018*

Scheck, L., 2021: A neural network based forward operator for visible satellite images and its adjoint. **J. Quant. Spectrosc. Ra.**, **274**, 107841. *https://doi.org/10.1016/j.jqsrt.2021.107841*

Turner, E., 2025: Diverse profile datasets from the ECMWF CAMS 137-level short range forecasts. **Technical Report NWPSAF-EC-TR-044**. *https://nwp-saf.eumetsat.int/downloads/profiles/nwpsaf_cams137_2025_doc.pdf*

GPU-adaptation of the IFS for EuroHPC machines

Michael Lange, Balthasar Reuter, Ahmad Nawab, Michael Staneker, Olivier Marsden, Patrick Gillies, Zbginiew Piotrowski, Johan Ericsson, Ioan Hadade, Michael Sleigh

The computational performance of the Integrated Forecasting System (IFS) has always been heavily optimised for operational use in ECMWF's data centre. Thanks to the European Commission's Destination Earth initiative (DestinE), in which the IFS forms a key part of two digital twins (DTs), the IFS is routinely run on a number of external EuroHPC systems. Many of these systems contain large numbers of graphics processing units (GPUs) to boost their computing power and support machine learning methods. In recent years, ECMWF, supported by DestinE, has worked in close collaboration with Member States, in particular Météo-France and the ACCORD consortium, with whom we share significant parts of the code, to adapt core components of the IFS forecast model to GPU accelerators. The introduction of modern software-engineering methods, paired with the thorough refactoring of the data structures underpinning the IFS, has enabled the continued optimisation for GPU accelerator architectures without impacting the central processing unit (CPU) code paths and scientific development. This has enabled ECMWF to include a GPU-enabled version of the IFS in its current high-performance computing (HPC) procurement benchmark, and it will allow for better utilisation of EuroHPC computing resources in the future.

Driving the shift to GPU computing

The use of GPUs has become ubiquitous in modern HPC. In addition to machine-learning models, and after the pioneering GPU-adaptation efforts of MeteoSwiss, several ECMWF Member States are using or exploring the use of GPU-based HPC systems for operations (Lapillonne, 2025). The ability to explore and leverage GPUs alongside CPU architectures was a driving force behind ECMWF's decision to pursue GPU support for the IFS, as it provides the ability to flexibly react to fast-changing trends in the hardware market. The desire to increase competition during HPC procurements by enabling the use of different GPU architectures from multiple vendors is a key driver behind ECMWF's GPU-adaptation strategy.

In addition, there is a long history of using external GPU resources for exploratory research runs at ECMWF that have often pushed far beyond the boundaries of operational HPC capacity (Wedi et al., 2020). With a significant increase in hardware diversity across HPC and with access to the EuroHPC ecosystem, support for GPU architectures is a key requirement to fully utilise such resources to their full potential. Such resources allow ECMWF to run the IFS at resolutions beyond the in-house operational capacity, enabling research and exploration towards the next generation of physical forecast models. Recognising this trend early, ECMWF has been involved in various European projects that pioneered the use of novel software technologies on external systems (Bauer et al., 2020; Müller et al., 2019; Segura et al., 2025; Targett et al., 2021).

The EuroHPC ecosystem and Destination Earth

The EuroHPC Joint Undertaking is a major European initiative aimed at developing a world-class supercomputing ecosystem across Europe. To date, EuroHPC has procured eleven state-of-the-art supercomputers, four of which have entered the top ten in the Top500 list of the fastest supercomputers (*https://www.top500.org/*), including the first European exascale supercomputer, JUPITER (*https://jupiter.fz-juelich.de/*).

Several of the supercomputer systems procured by EuroHPC are the target platforms for the DestinE digital twins (Geenen et al., 2024; Wedi et al., 2025). The IFS underpins the Continuous Global Extremes Digital Twin (Extremes DT) and is used in the Climate Change Adaptation Digital Twin (Climate DT; Doblas-Reyes et al., 2025), coupled to the NEMO 4 or FESOM 2 ocean model. Consequently, the IFS must be adapted and tuned to run optimally on these EuroHPC systems. ECMWF develops and operates the Extremes DT and collaborates with the Climate DT consortium on porting and optimising the IFS-based Climate DT configurations for the target platforms.

Many EuroHPC systems feature two main partitions of different node types:

- A traditional CPU-only cluster consisting of nodes with two CPUs each (e.g. LUMI-C, LEONARDO-

doi: 10.21957/f7d8-pv88

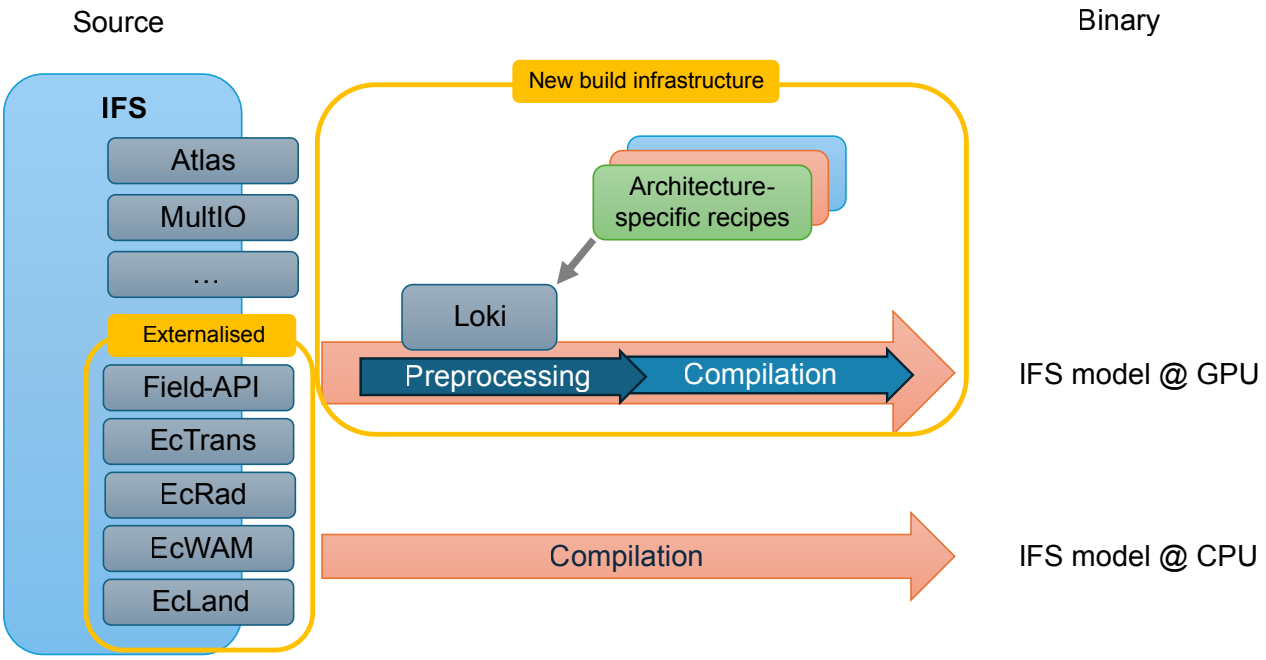

FIGURE 1 Overview of the technical infrastructure changes that enable GPU-adaptation alongside existing CPU compilation. In addition to existing external modules, like Atlas or MultIO, new technical packages and scientific sub-models have been externalised into standalone repositories. Additional build infrastructure, including the Loki source-to-source preprocessing, is then used to automatically generate GPU-optimised code paths during the compilation process.

DCGP, MareNostrum5-GPP), connected with a high-bandwidth network.

- An accelerated partition comprised of nodes with one or two CPUs and four GPUs each (e.g. LUMI-G, LEONARDO-BOOSTER, MareNostrum5-ACC).

The new JUPITER system marks a shift in this architecture: it currently consists solely of an accelerated partition (JUPITER-BOOSTER), with a much smaller CPU-only cluster planned to be installed later. Moreover, the four GPUs per node are complemented by the same number of CPUs, with each pair working in tandem (a so-called "super-chip"). This allows for faster data copies between CPU and GPU memory spaces and even a unified memory space that removes the need for explicit data transfers in code.

Software architecture

Contrary to conventional GPU porting efforts, where technical work is performed on a fixed version of the code, the IFS porting strategy focused on creating a sustainable solution in which scientific development can continue alongside the technical adaptation to GPUs without compromising operational performance on the CPU-based in-house HPC system. In close collaboration with Member States, a set of modern software engineering methods was devised that allowed a "GPU mode" to be embedded in the IFS without impeding the CPU code paths. The capability to run the IFS on GPUs and CPUs from the same code base allows the continual adaptation, testing and maintenance of both code paths, as well as future extensions to support multiple programming models and hardware vendors.

To achieve this separation of concerns, in which technical optimisations are done independently of scientific changes, a combination of library development, data-structure refactoring and source-to-source translation tools was used to create GPU-specific build modes for the Fortran source code. A particular emphasis was put on the extraction of several model sub-components into standalone open-source software libraries. This allowed ECMWF to benefit greatly from ongoing collaborations with HPC hardware vendors that provided significant performance optimisations to core components of the IFS. The provision of freely accessible standalone mini-apps, the IFS "dwarfs" (Müller et al., 2019), enabled not only the integration of hardware-specific optimisations, but also significantly increased the test coverage for technical changes before integration into the full system. Notable standalone packages include the spectral transform library ecTrans (*https://github.com/ecmwf-ifs/ectrans*), the wave model ecWAM (*https://github.com/ecmwf-ifs/ecwam*) and the radiation model ecRad (*https://github.com/ecmwf-ifs/ecrad*).

A pivotal element in GPU porting is the management of program data across two separate memory spaces:

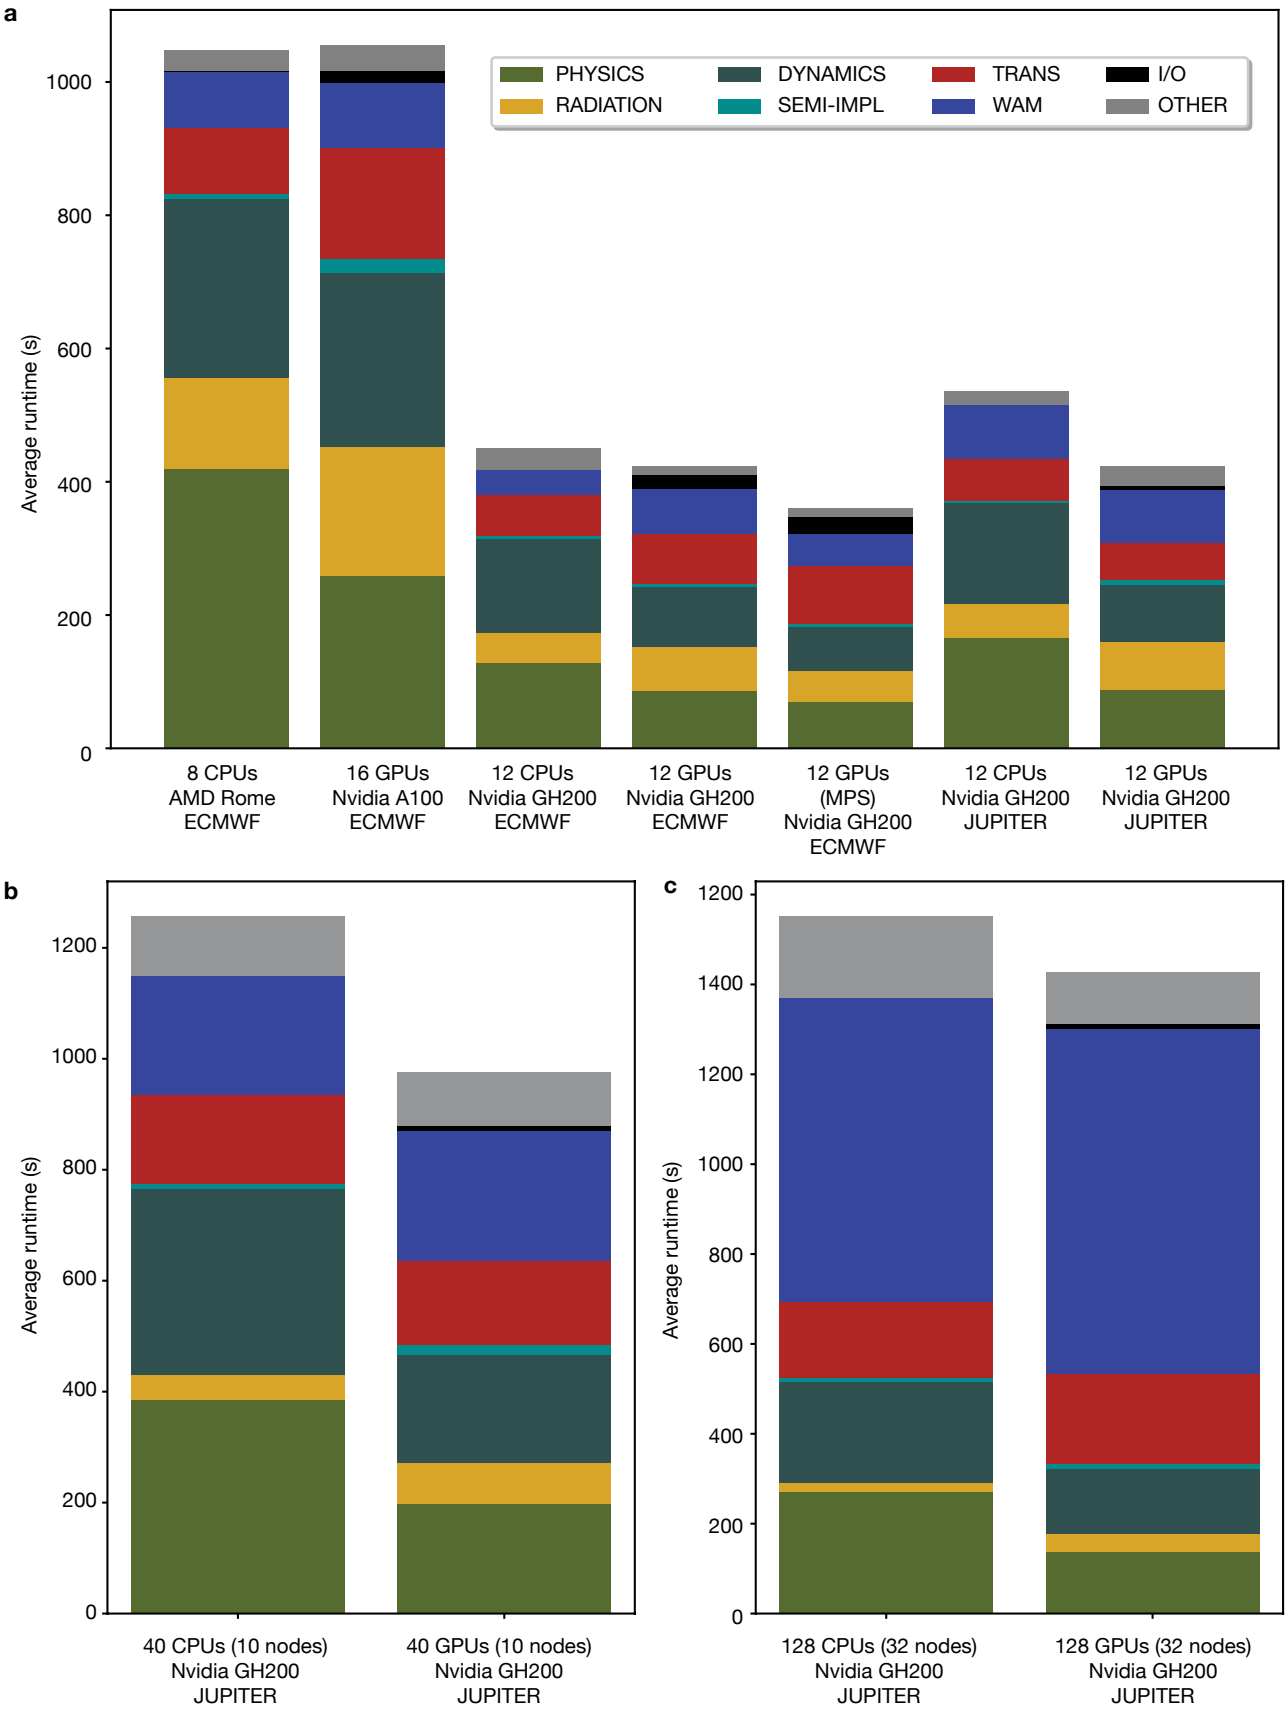

FIGURE 2 CPU–GPU performance comparison across resolutions and architectures. (a) Runtime breakdown of 48-hour forecast at 18 km resolution (TCo639) on GPU partitions of the ECMWF HPC and JUPITER. The initial column shows the runtime breakdown per component on the current operational AMD "Rome" CPU architecture and compares it to discrete GPUs (Nvidia A100). The performance of next-generation CPU-GPU "superchips" (Nvidia GH200, "Grace-Hopper") is then shown for ECMWF's internal machine learning partition and the EuroHPC system JUPITER. The use of GPU-oversubscription via Nvidia's Multi-Process Service (MPS) on the internal system brings further performance gains for GH200 but was not yet available on JUPITER. (b) Comparison of CPU and GPU runtimes for 48-hour forecast at 9 km resolution (TCo1279) on JUPITER, showing per-component breakdowns on minimal node counts (10 nodes, 40 CPUs/40 GPUs). (c) Comparison of CPU and GPU runtimes for 12-hour forecasts at 4.4 km resolution (TCo2559) on JUPITER, using 32 nodes (128 CPUs/128 GPUs).

GPU memory and CPU memory. This challenge is addressed through the "Field-API" (*https://github.com/ecmwf-ifs/field_api*), an array data management abstraction that handles data movement between CPU and GPU while remaining compatible with the complex memory data layout in the IFS. This dedicated Fortran infrastructure library provides vendor-agnostic support for GPU data offload and delivers optimised backends that use low-level programming paradigms, such as CUDA and HIP, to ensure efficient data movement and layout in GPU device memory.

The key to separating the scientific algorithm from the necessary code optimisations is the source-to-source translation tool Loki (*https://github.com/ecmwf-ifs/loki*), a vital component of the Digital Twin Engine (DTE) developed at ECMWF with support from DestinE. Loki is a Python package that allows HPC specialists to encode a set of complex source code transformations as automated "recipes" that are applied during the model compilation process to transform and optimise the CPU code to GPUs. These recipes are highly bespoke and use IFS coding conventions to derive GPU-optimised code from conventional IFS subroutines, using domain-specific knowledge of the underlying algorithms.

Implementing multi-architecture capabilities

The extraction of several key abstractions and scientific sub-components into open-source, standalone modules was key to implementing the IFS's new multi-architecture capabilities. This approach not only allows the development of prototypes of new technical features independently but also enables external experts to optimise performance for specific hardware. This ability was crucial, as the IFS includes a range of diverse computational patterns, often requiring tailored porting strategies for CPU and GPU. For example, the spectral transforms, central to the model's dynamical core, demand highly customised optimisations to scale on large HPC systems. While these optimisations often differ between CPU and GPU, the externalised ecTrans package now offers dedicated code paths for both, unified under a common library interface. Thanks to its open-source release, the GPU paths have also been optimised by experts from Nvidia and AMD for both GPU architectures, resulting in a highly optimised library that supports a variety of CPU and GPU architectures.

As highlighted in Figure 1, for model components where code replication was avoided through automated source-to-source translation, extracting small, representative examples also proved essential for sustainable multi-architecture support. This included both individual routines – like the CLOUDSC cloud microphysics scheme – and entire sub-models such as ecWAM (waves) and ecRad (radiation). These examples enabled the development of tailored solutions combining GPU-enabled data structures, automated translation via Loki, and manual porting to test GPU support. Many of these efforts were highly collaborative, involving Member States, European project partners, the DestinE initiative, and industry experts. This open development model supports ongoing improvements to translation recipes and GPU-specific optimisations, allowing ECMWF to react quickly to emerging hardware trends.

Sustainability and testing

ECMWF has recently added a new GPU partition to its Bologna HPC facility. In addition to enabling and enhancing the machine learning abilities at ECMWF (Pappenberger et al., 2024), this new hardware allows the continuous-integration testing (CI) of the newly developed GPU capabilities. The automated deployment of GPU-enabled test runs, in conjunction with the automated generation of GPU-enabled code paths from standardised CPU code, allows the correctness of the GPU code paths to be routinely checked and verified with only a moderate maintenance overhead.

Thanks to the access provided by DestinE, this capability is even extended to external HPC systems (e.g. LUMI) for individual externalised model components, such as the spectral transform library ecTrans. As a result, the IFS has become a true multi-architecture model that is tested and deployable across a range of HPC systems. The sustainability of this approach – and the effectiveness of separation of concerns – was demonstrated when the GPU-adaptation recipes developed for IFS Cycle 49r1 were migrated successfully to the next version, Cycle 50r1. This was achieved in a relatively short time by a few technical experts, without the need for any intervention from scientific developers.

Status and initial results

With the recent rise of GPU architectures to the forefront of HPC systems, GPU hardware characteristics have become increasingly diverse. In addition to traditional discrete GPU systems, where data transfers between CPU and GPU memory regions are often the main performance bottleneck, shared-memory systems have been developed. These combine CPU and GPU architectures in the same chip or contain dedicated hardware for more efficient CPU–GPU memory transfers. The ECMWF data centre contains two GPU partitions, representing both hardware types.

ECMWF's overall GPU-adaptation strategy aims to stay agnostic to specific GPU types. The recent focus has been on porting most of the atmospheric time step of the IFS to GPU to allow data to stay resident in the GPU memory, reducing the number of data transfers required. This development allows the systematic testing and

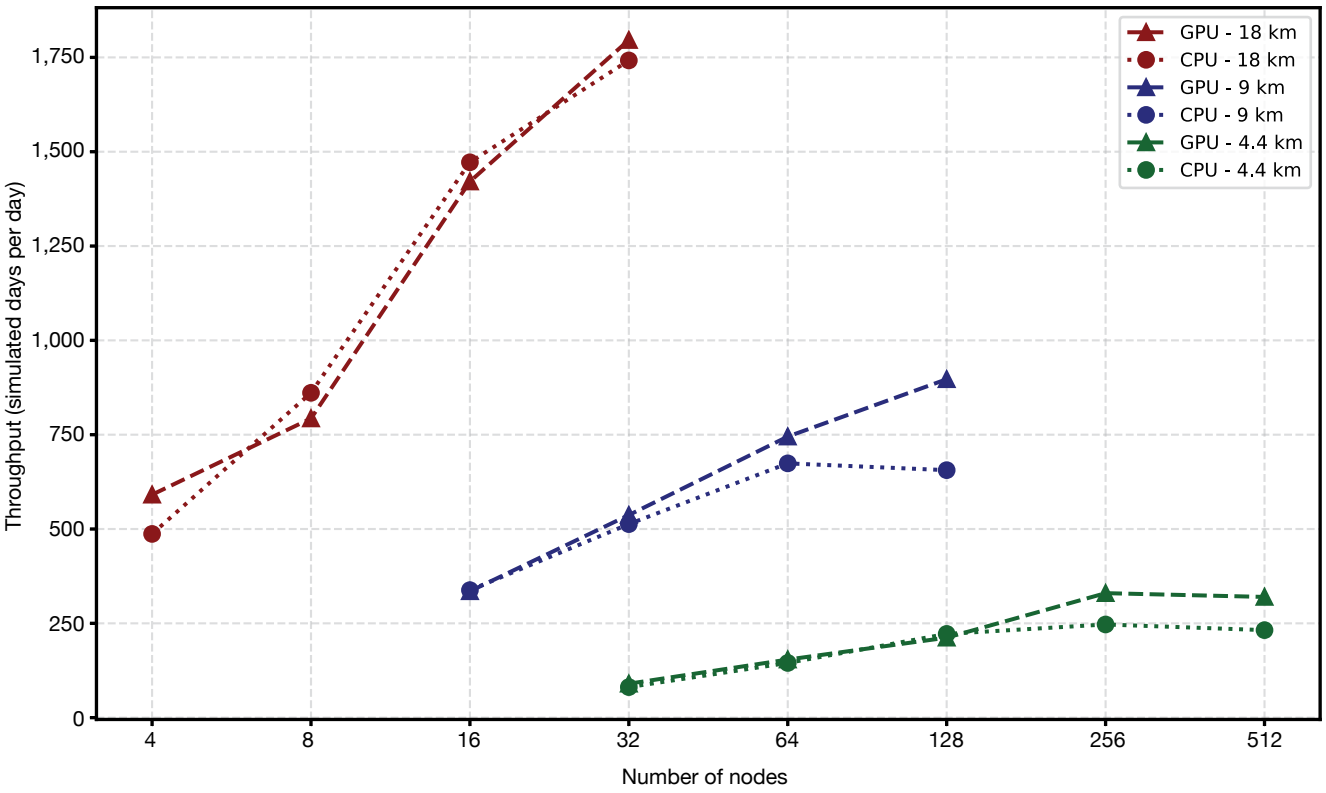

FIGURE 3 Strong scaling of CPU and GPU runs on JUPITER for research and operational resolutions. This highlights that initial GPU optimisations are already sufficient to preserve the excellent strong scaling efficiency of the original IFS implementation for CPUs.

assessment of performance of the GPU port and enables the verification that operational CPU performance has not been affected. The most recent GPU-adaptation efforts have also been aligned with the latest scientific developments in IFS Cycle 50r1 – a step made possible by the high degree of automation involved in the derivation of the GPU code paths. The resulting GPU capabilities are intended to be sustainable across several cycles, allowing further optimisation in ongoing and future developments.

Initial results are already showing promise. A recent performance comparison is shown in Figure 2, where discrete GPUs (Nvidia A100) in the ECMWF HPC centre are compared with next-generation shared-memory GPUs (GH200, "Grace-Hopper"), in both the new in-house GPU partition and the EuroHPC system JUPITER. The initial, *non-optimised* GPU performance is already competitive with the *highly optimised* CPU performance, both at development and operational resolutions. Considering the computational power of the next-generation "Grace" CPUs used in the respective systems and the known optimisation headroom for the GPU implementation, these results provide a strong starting point for further improvements.

Importantly, as shown in Figure 3, the excellent multi-node scalability of the IFS is also preserved on GPUs, suggesting that future explorations of higher resolutions beyond operational constraints will be able to utilise GPU-driven machines just as effectively as they will CPU-based systems. Given the prevalence of GPU systems in the modern HPC ecosystem, this ability to compare and utilise different types of high-end systems with the IFS is essential to objectively assessing and adapting to changes in the HPC landscape over the coming years.

Conclusion and next steps

The initial adaptation of the IFS to GPU accelerators has been achieved and concluded with the inclusion of the novel GPU build modes in the latest HPC procurement benchmark release. Through the involvement in external projects and the DestinE initiative, ECMWF will continue its efforts to maintain and improve upon these capabilities to best utilise any available internal or external computing resources. The high degree of automation and the enhanced automated testing capabilities on internal GPU hardware enable these features to be maintained sustainably. As a result, the novel GPU capabilities are planned to enter the mainline code in IFS Cycle 50r2 and remain available to ECMWF researchers. The capabilities will be retained in subsequent IFS releases, allowing potential inclusion in future releases of OpenIFS.

In addition to novel technical capabilities, one of the key outcomes of the GPU-adaptation efforts has been the restructuring of several components of the forecast model and the improvement in technical testing infrastructure. ECMWF recognises the need to further address existing technical debt in the IFS code base and aims to improve the sustainability of the physical model, ensuring that future technical changes can be implemented with the required agility. The FORGE (Forecast System Regeneration) project (Sleigh et al., 2025) has been established to spearhead the modernisation of the forecast system and associated software infrastructure. This initiative will work closely with Météo-France and ACCORD partners to ensure that ECMWF continues to adapt to the increasing technical demands in a highly diverse environment.

Further reading

Bauer, P., **T. Quintino**, **N. Wedi**, **A. Bonanni**, **M. Chrust**, **W. Deconinck et al.**, 2020: The ECMWF Scalability Programme: progress and plans. *ECMWF Technical Memorandum* **No. 857**. *https://doi.org/10.21957/gdit22ulm*

Doblas-Reyes, F. J. et al., 2025: The Destination Earth digital twin for climate change adaptation. *GMD [preprint]*. *https://doi.org/10.5194/egusphere-2025-2198*

Geenen, T. et al., 2024: Digital twins, the journey of an operational weather prediction system into the heart of Destination Earth. *Procedia Computer Science*, **240**, 99–109. *https://doi.org/10.1016/j.procs.2024.07.013*

Lapillonne, X., et al., 2025: Operational numerical weather prediction with ICON on GPUs (version 2024.10). *Geoscientific Model Development [preprint]*. *https://doi.org/10.5194/egusphere-2025-3585*

Müller, A. et al., 2019: The ESCAPE project: Energy-efficient Scalable Algorithms for Weather Prediction at Exascale. *Geoscientific Model Development*, **12(10)**, 4425–4441. *https://doi.org/10.5194/gmd-12-4425-2019*

Pappenberger, F. et al., 2024: Artificial intelligence and machine learning: revolutionizing weather forecasting. UN United in Science 2024 Report.

Segura, H. et al., 2025: nextGEMS: entering the era of kilometer-scale Earth system modelling. *Geoscientific Model Development*, **18(20)**, 7735–7761. *https://doi org/10.5194/gmd-18-7735-2025*

Sleigh, M. et al., 2025: Modernisation of the Integrated Forecasting System. *ECMWF Newsletter* **No. 182**, 19–23. *https://doi. org/10.21957/m9ad5hv72s*

Targett, **J. et al.**, 2021: Systematically migrating an operational microphysics parameterisation to FPGA technology. *2021 IEEE 29th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pp. 69–77. *https://doi.org/10.1109/FCCM51124.2021.00016*

Wedi, N. et al., 2020: A baseline for global weather and climate simulations at 1 km resolution. *Journal of Advances in Modeling Earth Systems*, **12(11)**. *https://doi org/10.1029/2020MS002192*

Wedi, N. et al., 2025: Implementing digital twin technology of the earth system in Destination Earth. *Journal of the European Meteorological Society*, Volume 3, 100015. *https://doi.org/10.1016/j.jemets.2025.100015*

ECMWF Council and its committees

The following provides some information about the responsibilities of the ECMWF Council and its committees. More details can be found at: *http://www.ecmwf.int/en/about/who-we-are/governance*

Council

The Council adopts measures to implement the ECMWF Convention. Its responsibilities include admission of new members, authorising the Director-General to negotiate and conclude co-operation agreements, and adopting the annual budget, the scale of financial contributions of the Member States, the Financial Regulations and the Staff Regulations, the long-term strategy and the programme of activities of the Centre.

President Dr Roar Skålin (Norway)

Vice President Ms Virginie Schwarz (France)

Policy Advisory Committee (PAC)

The PAC provides the Council with opinions and recommendations on any matters concerning ECMWF policy submitted to it by the Council. The President of the Council can also submit requests to the Committee in the period between Council meetings.

Chair Ms Virginie Schwarz (France)

Vice Chair Prof. Dr Maarten van Aalst (Netherlands)

Finance Committee (FC)

The FC provides the Council with opinions and recommendations on all administrative and financial matters submitted to the Council and exercises the financial powers delegated to it by the Council.

Chair Mr Lukas Schumacher (Switzerland)

Vice Chair Ricardo José Squella de la Torre (Spain)

Scientific Advisory Committee (SAC)

The SAC provides the Council with opinions and recommendations on the draft programme of activities of the Centre drawn up by the Director-General and on any other matters submitted by the Council. The members of the SAC are appointed in their personal capacity and are selected from among the scientists of the Member States.

Chair Dr François Bouyssel (France)

Vice Chair Prof Dr Nedjeljka Žagar (Germany)

Technical Advisory Committee (TAC)

The TAC provides the Council with advice on the technical and operational aspects of the Centre including the communications network, computer system, operational activities directly affecting Member States, and technical aspects of the four-year programme of activities.

Chair Ms Anne-Cecilie Riiser (Norway)

Vice Chair Mr Antonio Vocino (Italy)

Advisory Committee of Co-operating States (ACCS)

The ACCS draws up, for submission to the Council, opinions and recommendations on the programme of activities and the budget of the Centre, on items relevant to Co-operating States, and on any matter submitted to it by the Council.

Chair Mr Ilian Gospodinov (Bulgaria)

ECMWF publications

(see *www.ecmwf.int/en/research/publications*)

Technical Memoranda

935 **Ruparell, K., K. Hunt, H. Cloke, C. Prudhomme, F. Pappenberger & M. Chantry:** A Hydrologist's guide to the CRPS. *December 2025*

934 **van Niekerk, A., B. Suetzl, N. Raoult, M. Janousek & I. Bastak Duran:** Bayesian optimisation of parameters in the ECMWF IFS. *December 2025*

933 **Pillosu, F., T. Hewson, E. Gascón, M. Vuckovic, C. Prudhomme & H. Cloke:** Bridging the scale gap: enhancing point-scale rainfall estimates by post-processing ERA5. *October 2025*

932 **Massart, S. & A. Geer:** Sea-ice concentration analysis in model space in IFS cycle Cy49r1. *October 2025*

ESA or EUMETSAT Contract Reports

El Garroussi, S., F. Di Giuseppe, J. McNorton, P. de Rosnay, S. Garrigues, D. Fairbairn & P. Weston: Satellite-informed fuel estimation using hybrid data assimilation**.** *November 2025*

Kolassa, J., P. Weston, K. Salonen & P. de Rosnay: Quarter 3 2025: Operations Service Report. *October 2025*

EUMETSAT/ECMWF Fellowship Programme Research Reports

Steele, L., N. Bormann, A. Geer, M. Chrust & D. Duncan: Enhancing the exploitation of all-sky microwave sensors at ECMWF using inter-channel error correlations. *December 2025*

Warrick, F.: Quality Assessment and Assimilation of Meteosat-12 FCI AMVs. *November 2025*

Warrick, F.: Quality Assessment of IASI 3D Winds**.** *October 2025*

ECMWF Calendar 2026

Feb 2–5	Training course: Use and interpretation of ECMWF products
Feb 11	WeatherGenerator Dissemination Workshop
Feb 10–12	WeatherGenerator General Assembly
Feb 23–27	Training course: Machine learning and Destination Earth
Mar 2–6	Training course: Parametrization of subgrid physical processes
Mar 16–20	Training course: Data assimilation and Machine Learning
Mar 19	AI Weather Quest: DJF Awards Webinar
Mar 23–27	Training course: EUMETSAT NWP-SAF satellite data assimilation
Apr 13–17	5th ECMWF–ESA Machine Learning Workshop
Apr 13–17	Training course: Numerical methods for weather prediction
Apr 29	Policy Advisory Committee (virtual)
Apr 30	Finance Committee (virtual)
May 18–21	Online Training course: Atos HPC training
Jun 1–5	Using ECMWF's Forecasts (UEF2026)
June 3-5	9th C3S General Assembly
Jun 16–17	Council
Jun 29 – Jul 3	2026 AR Recon workshop and 2nd Observational campaigns workshop for better weather forecasts
Jun 29 – Jul 2	International Workshop on Greenhouse Gas Measurement from Space
Sep 14–18	Annual Seminar 2026
Oct 5–7	Scientific Advisory Committee
Oct 8–9	Technical Advisory Committee
Oct 20–21	Finance Committee
Oct 21	Policy Advisory Committee
Nov 4-6	Visualising Climate Event 2026
Nov 23-27	Workshop on Software Strategies for Sustainable Physical Modelling
Dec 8–9	Council

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/*

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

For any query, issue or feedback, please contact ECMWF's Service Desk at *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.

ECMWF

Newsletter | **No. 186** | Winter 2026

European Centre for Medium-Range Weather Forecasts

www.ecmwf.int