ECMWF
European Centre for Medium-Range Weather Forecasts

Newsletter

No. 187 | Spring 2026

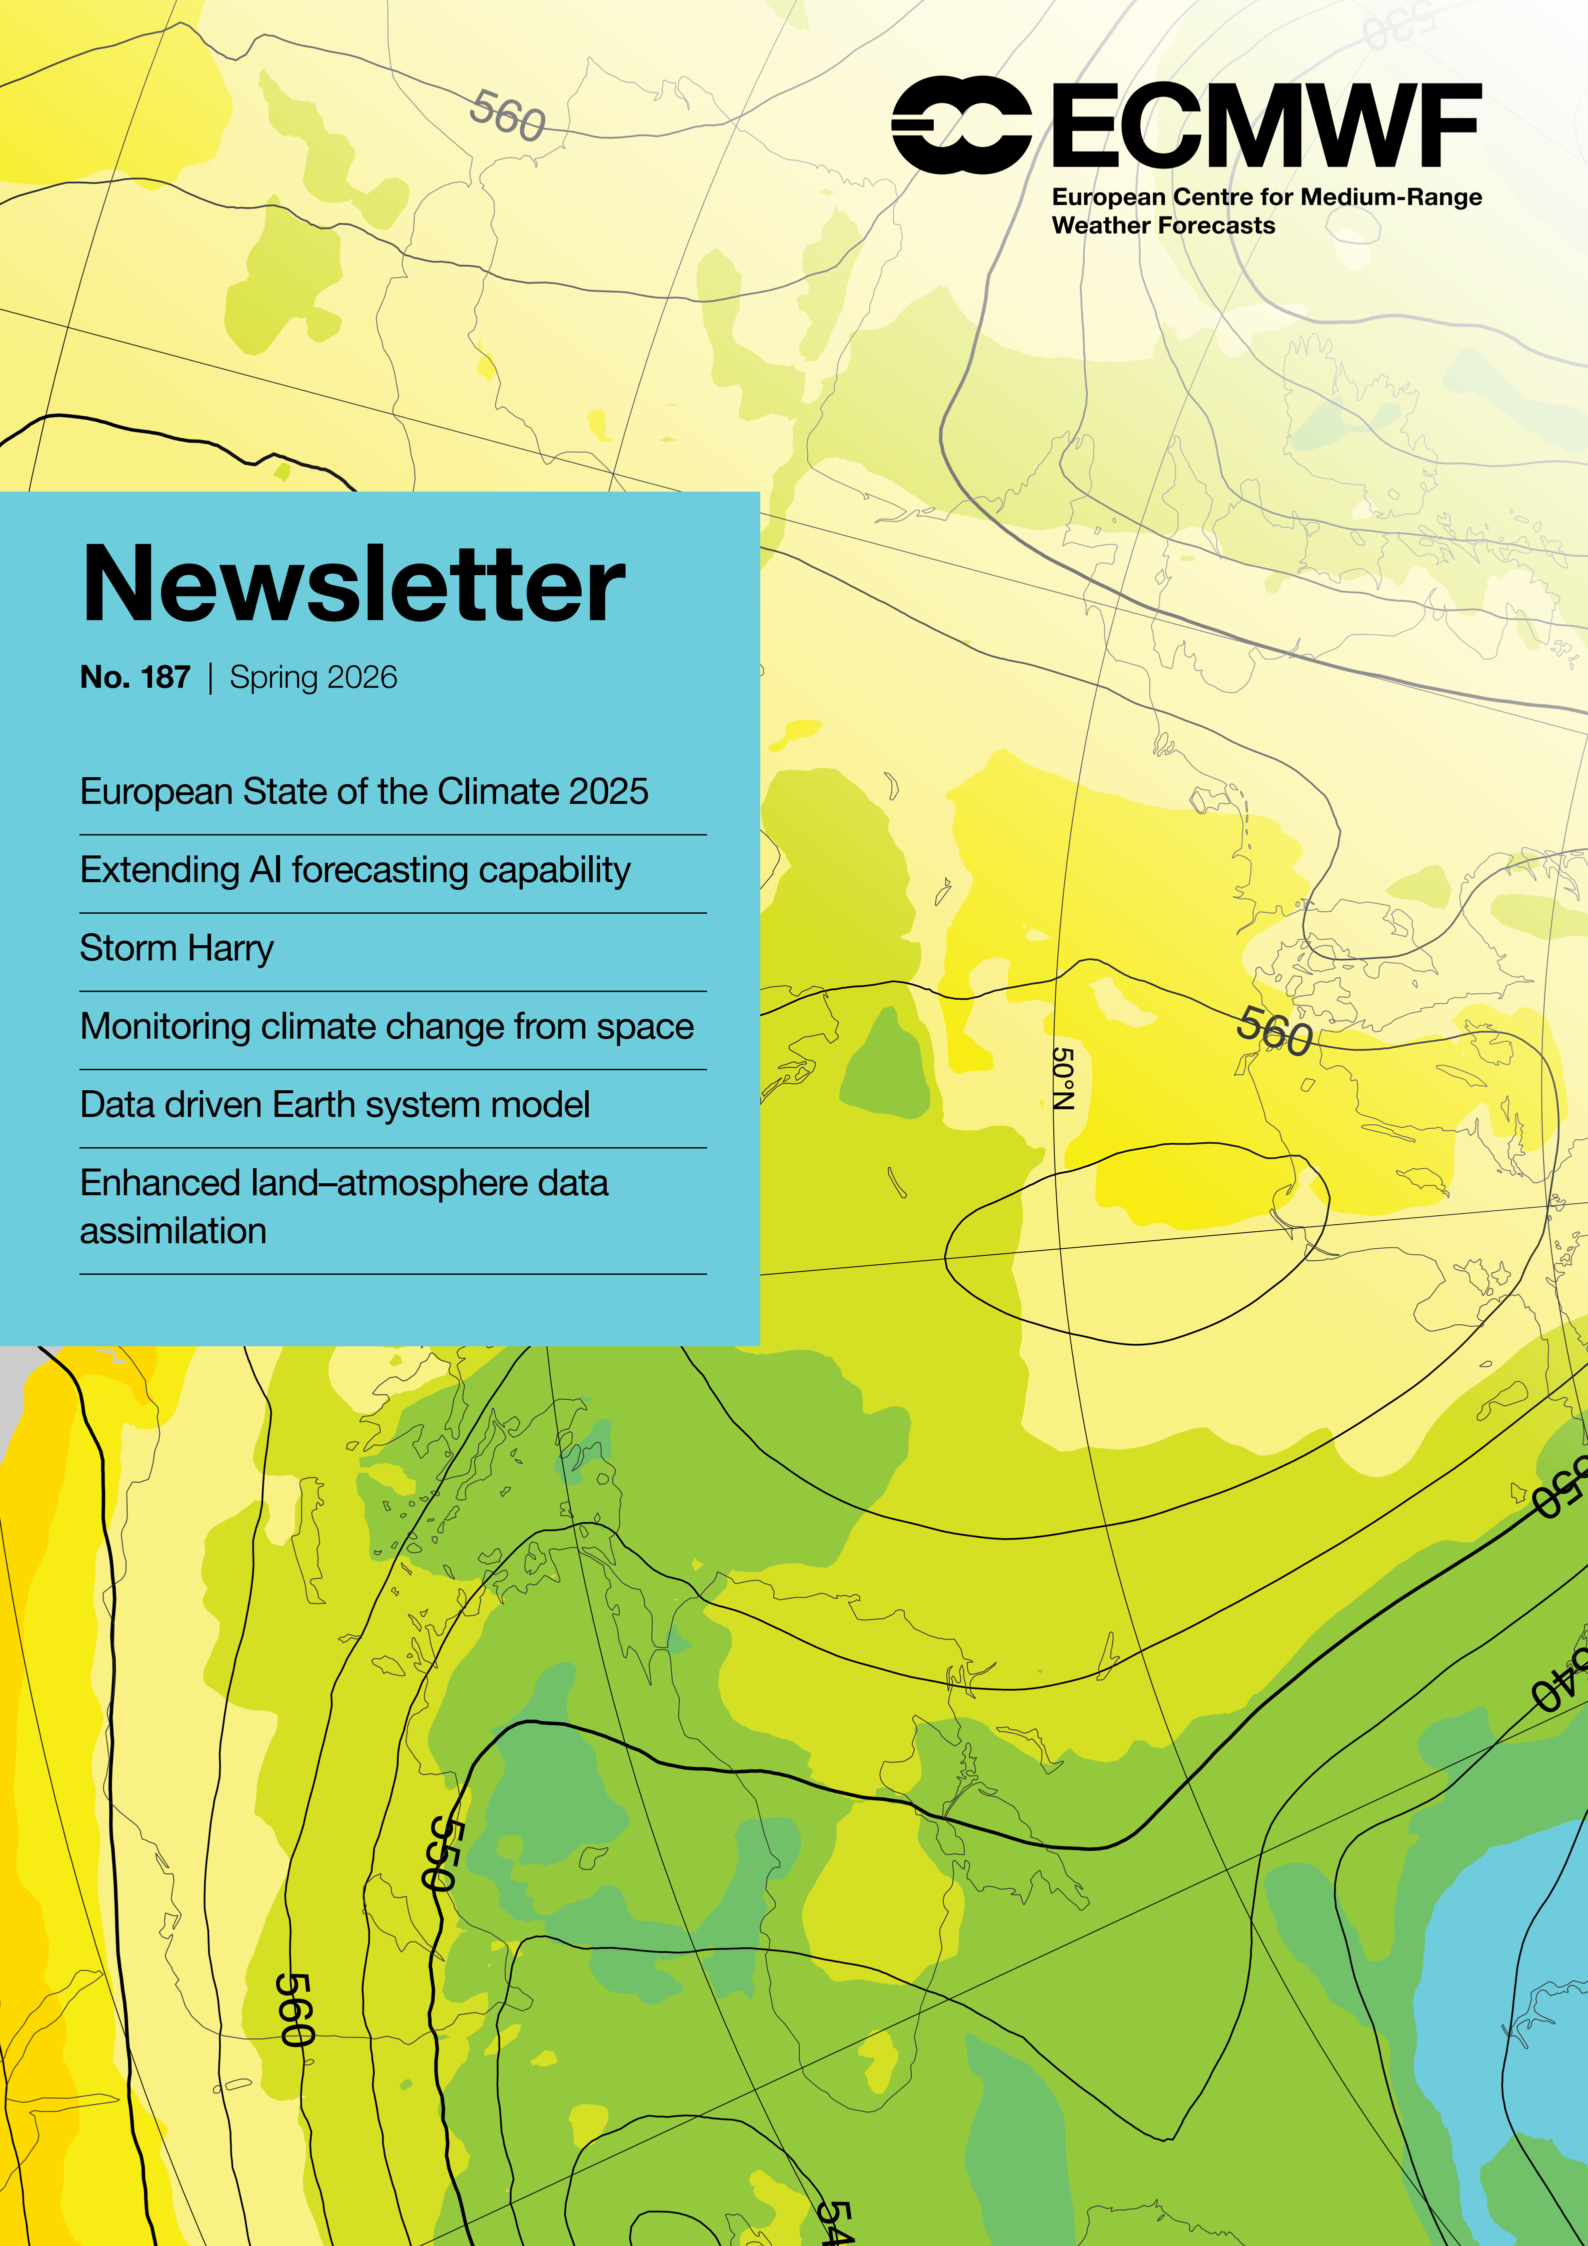

Publication policy

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The ECMWF Newsletter is not peer-reviewed.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*

Strengthening our Earth system approach

As we move further into 2026, ECMWF continues to build on its strong foundations while accelerating progress towards a more integrated, data driven Earth system approach. Innovation across our activities is moving at pace, but our focus remains clear: translating this work into reliable services for our Member and Co operating States and strengthening the capabilities we deliver together with our community.

A key milestone in this journey is the progression of Destination Earth into its third phase. Together with our European partners, we are entering a period of consolidation, operating and further evolving the Climate and Extremes Digital Twins while advancing the Digital Twin Engine.

In parallel, we are advancing the development of a data driven Earth system model that extends beyond the atmosphere to include the ocean, land, waves and sea ice. This represents a significant step forward in how we combine observations, modelling and machine learning at scale, and demonstrates the strength of collaboration across Europe in delivering capabilities that no single organisation could achieve alone.

The transition of OpenIFS to open-source marks another important step in strengthening openness and collaboration across our community. By making a portable version of our global forecasting model openly available, we are reducing barriers to engagement and enabling a wider range of users to work with, contribute to and build on a shared system. This reflects the long standing collaborative nature of the Integrated Forecasting System and reinforces ECMWF's role as a trusted partner, while helping ensure that development, innovation and operational use remain closely aligned and scientifically robust.

A central theme of this edition is the continued evolution of our forecasting system towards a more fully coupled, end to end Earth system. Advances in land–atmosphere coupling and developments in IFS Cycle 50r1 are improving how interactions across components are represented, supporting more consistent and accurate forecasts across timescales and providing a stronger basis for operational services.

Data driven approaches are becoming an increasingly integral part of this landscape. The implementation of AIFS v2, alongside Cycle 50r1, marks an important step in extending AI based forecasting capabilities, including new variables and demonstrated performance gains. Our objective is not to replace physically based modelling, but to combine strengths while maintaining the robustness and integrity that underpin our forecasts.

The newsletter also highlights the breadth of ECMWF's reach. From global engagement through initiatives such as AI Weather Quest, to applications in areas such as renewable energy, our work spans the full pathway from innovation to services that support decision making across a wide range of sectors and users.

Underlying all of this is partnership. ECMWF's success depends on close collaboration with our Member and Co operating States, the European Commission and a wide network of partners. Strengthening these relationships, listening to user needs and co developing solutions remain central to our approach and to ensure that innovation is translated into operational value.

This newsletter provides a snapshot of these ongoing efforts. Together, they reflect an organisation evolving with purpose: integrating new technologies, advancing Earth system approaches, and delivering forecasts and information that help societies prepare for and respond to an increasingly complex and changing world.

Florian Pappenberger
Director-General

doi: 10.21957/gxrf-vv36

Contents

Editor Chelsea Snell • **Typesetting and graphics** Chris Matthews • **Cover** See figure on page 24

European State of the Climate 2025 report published

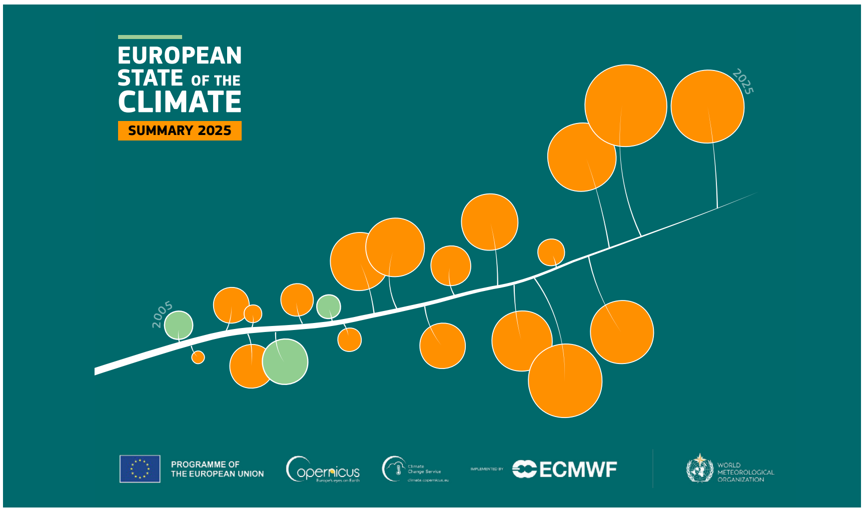

On 29 April 2026, the European State of the Climate (ESOTC) 2025 report was published (https://climate.copernicus.eu/esotc/2025).

The ESOTC report is produced by the EU's Copernicus Climate Change Service (C3S), implemented by ECMWF, and the World Meteorological Organization (WMO), and brings together the work of around 100 scientific contributors.

It provides a comprehensive overview of climate conditions in Europe in 2025 and updates on long-term trends, covering variables from across the Earth system, key events and their impacts, and a discussion of climate policy and action with a focus on biodiversity.

Rapid warming in Europe's colder regions

The report shows a continuing trend of rapid warming across Europe's coldest regions, including in the Arctic and the Alps, where snow and ice play a critical role in slowing climate change by reflecting sunlight back into space.

The area of Europe experiencing winter days with freezing temperatures is shrinking. Europe as a whole also experienced a record low number of cold stress days in 2025, with 90% of the continent experiencing fewer days than average with at least 'strong' cold stress.

Above-average temperatures and below-average precipitation led to a significant reduction in snow and ice cover. In March 2025, the snow-covered area in Europe was around 31% below average – the third lowest snow extent since records began in 1983. Glaciers also experienced a net mass loss, with Iceland recording its second-largest glacier mass loss on record. The Greenland Ice Sheet lost 139 gigatonnes (139 billion tonnes) of ice – contributing 0.4 mm to global sea level rise.

Above-average summer temperatures and heat stress

Summer temperatures were above average across much of the continent in 2025, with several heatwaves.

July saw Europe's second most severe heatwave on record, lasting 25 days (7–31 July) and affecting a large part of the continent

In Türkiye, temperatures reached 50°C for the first time, with 50.5°C recorded at Silopi. In Cyprus, temperatures reached 44.7°C, a new record for July.

Sub-Arctic Fennoscandia saw its longest heatwave on record in July, spanning three weeks. During this period, temperatures close to and within the Arctic Circle exceeded 30°C, peaking at 34.9°C in Frosta, Norway.

In 2025, much of Europe (41%) also experienced more days than average with at least 'strong' heat stress, with southern and eastern Spain recording up to 50 more days than average with 'feels-like' temperatures above 32ºC.

Record marine heat and ocean warming

The annual average sea-surface temperature (SST) for the European region was the highest on record for the fourth consecutive year. Marine heatwaves affected 86% of the region, with 36% experiencing 'severe' or 'extreme' conditions – the highest proportion on record.

The entire Mediterranean Sea experienced at least one day with 'strong' marine heatwave conditions in each of the past three years. In the Norwegian Sea, 'severe' marine heatwave conditions also occurred in 2025.

Drier-than-average conditions across much of Europe in 2025

Much of Europe saw drier-than-average conditions in 2025, but with strong regional contrasts. A large area from northwestern to eastern Europe saw below-average precipitation, and it was one of Europe's three driest years for soil moisture since 1992. In May, around half of Europe (53%) was affected by drought conditions. Annual river flow was below average in 70% of rivers, and while storms and flooding affected some areas, overall extreme precipitation and flooding were less widespread than in recent years.

Wildfires: burnt area in Europe reached a record level in 2025

A record area of around 1,034,000 hectares burnt across Europe in 2025. The largest contributions came from fires across the Iberian Peninsula in August.

Emissions from fires reached their highest levels on record in 2025, with Spain contributing around half of the total. Spain, Cyprus, the United Kingdom, the Netherlands and Germany all saw record annual high emissions.

Destination Earth to move into third phase

The European Commission confirmed the third phase of Destination Earth (DestinE) on 1 February 2026.

DestinE is an initiative of the European Commission to build a highly accurate digital replica of the Earth system, enabling the exploration of past, present and likely future climate conditions and extreme events, including through tailored “what-if” scenarios. This helps European and national institutions better understand, prepare for, and adapt to climate- and weather-related risks.

The initiative is jointly implemented by ECMWF, the European Space Agency (ESA) and the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT), under the leadership of the European Commission’s Directorate-General for Communications Networks, Content and Technology (DG Connect).

The core foundations

During its first two phases, ECMWF worked with more than 100 partner organisations across Europe, including national meteorological and hydrological services, to establish the core components of DestinE.

These include the Climate Change Adaptation Digital Twin, which creates an operational framework for producing global, high-resolution climate projections over multi-decadal timescales, and the Weather-Induced Extremes Digital Twin, which simulates extreme weather events and their impacts at high resolution both globally and across Europe. The Digital Twin Engine enables these systems to run on the supercomputing infrastructure of the EuroHPC Joint Undertaking and manages the production of and access to the high-resolution digital-twin data.

Substantial progress has also been made in artificial intelligence (AI) activities. This included the development of machine-learning components representing different parts of the Earth system, including land, ocean, sea ice, waves and hydrology. These developments contribute to an emerging AI Earth system model designed to complement physics-based simulations, support uncertainty analysis and enable rapid scenario testing. Prototype tools such as a “forecast-in-a-box” and a digital twin assistant have also been created to improve interaction with models and data.

Phase three: consolidation

In phase three, running from June 2026 to June 2028, ECMWF and its partners will focus on operating and further evolving the Climate and Extremes Digital Twins and the Digital Twin Engine. Work will also advance AI models developed in earlier phases and transform digital-twin data into high-quality datasets that can feed into Europe’s AI Factories.

By integrating physical understanding with innovative AI approaches, Europe’s weather and climate prediction capabilities will be further enhanced. These capabilities will support national meteorological and hydrological services and public authorities in preparing for extreme events in a changing climate.

The combination of high-resolution datasets and European AI infrastructure is also expected to enable new AI applications for weather and climate, supporting innovation across industry and the public sector while strengthening Europe’s position in trustworthy AI for environmental intelligence.

A collaborative European effort

DestinE brings together expertise in Earth system modelling, software engineering, high-performance computing and artificial intelligence, drawing on European and national investments across these fields. The initiative is fundamentally collaborative, pooling resources across Europe and Member States to complement existing national and European services and help institutions respond to climate change and extreme weather risks.

With phase three now confirmed, DestinE continues to strengthen Europe’s capacity for preparedness, resilience and innovation in the face of a changing climate.

Plume expands to wind energy plugins: latest updates from the DTWO project

Clara Ducher, Antonino Bonanni, Domokos Sármány, Tiago Quintino

"Federated Digital Twins for Wind-Offshore" (DTWO; **https://dtwo-project.eu/**) is a Horizon Europe-funded project that aims to unlock new possibilities in offshore wind-power production. It does this by developing a digital twin that integrates existing models and simulation assets for greater energy reliability and security. The project started in June 2024 and recently passed its 18-month mid-term review. Involving 12 organisations from 9 countries, DTWO brings together expertise from across the wind-energy value chain, from the world's largest offshore wind turbine manufacturer to research centres, academia, IT, and energy and weather-forecasting institutions.

To achieve its objectives, DTWO is organised around five scientific modules: Earth, Wakes, Siting, Turbine and Grids. These modules combine digital assets from each project partner. They will be integrated and orchestrated within a federated digital twin software architecture, known as the DTWO platform, which is designed to help wind asset owners better manage and optimise their portfolios. Platform users will access the modules via four use cases defined by industrial partners. These are currently being integrated and tested through five test scenarios representing real-world wind farm applications.

ECMWF led the Earth module to completion last November and is now leading the initial scientific integration of the Siting and Turbines modules into the DTWO platform. In addition, ECMWF is facilitating the use of data from the EU's Destination Earth (DestinE) initiative across the various use cases, ensuring that partners can efficiently access and use these resources. ECMWF is also enhancing the usability of global weather prediction data for wind-energy applications by developing dedicated plugins for targeted workflows. These leverage the Plume framework (**https://github.com/ecmwf/plume**) to equip the Integrated Forecasting System (IFS) with additional functionalities tailored to sector-specific needs.

DTWO modules

Earth: Advanced modelling of atmospheric, ocean and wave conditions crucial for offshore wind planning and operating.

Wakes: Modelling inter- and intra-farm wake interactions to optimise turbine placement and maximise wind-farm energy production.

Siting: Assessment of wind resources, environmental conditions and turbine design to identify the most efficient and cost-effective offshore wind-farm locations.

Turbine: Monitoring turbine health and predicting reliability, maintenance needs and component lifetimes to reduce downtime and operational costs.

Grids: Analysis of grid connectivity, energy markets and storage integration to support investment planning and reliable offshore wind power delivery.

Practical impact

From ECMWF's perspective, DTWO fosters the development of the Plume framework with a particular focus on wind energy use cases. Plume is a plugin mechanism for Earth system models and forms part of the Digital Twin Engine developed within DestinE. It is also part of the ECMWF Software Engine, where it is currently at an emerging maturity level. An initial proof-of-concept, co-funded by DestinE, showed promise for applications such as point extraction for the single-column model and machine-learning-based tropical cyclone detection. Plume has since been integrated with DestinE Cycles 48r3 and 49r3. Through DTWO, the framework has been consolidated and the plugin portfolio expanded to include two new plugins (Figure 1):

1. wind farm parametrization, for power estimation and wind field impact assessment; and
2. extreme event capture, for high-impact weather signal identification.

Continuous improvements to Plume are being implemented as user needs are identified. As DTWO's objectives required the creation of the two new plugins described above, a numerical weather prediction emulator was introduced into Plume during the initial months of the project to facilitate the development and testing of plugins. The emulator supports two running modes, with one serving synthetic data and another for historical data from GRIB files to emulate a model run with Plume enabled. This capability makes plugin development accessible to a broader community, enabling users to better shape the future of the framework.

The DTWO project's focus on offshore wind motivated the introduction of support for coupled model runs in Plume, whereas previously it was limited to the atmospheric model. The addition of a Plume entry point in the wave model, ecWAM, now allows plugins to process wave model fields. For example, plugins can detect extreme wave conditions that would prevent turbine maintenance operations.

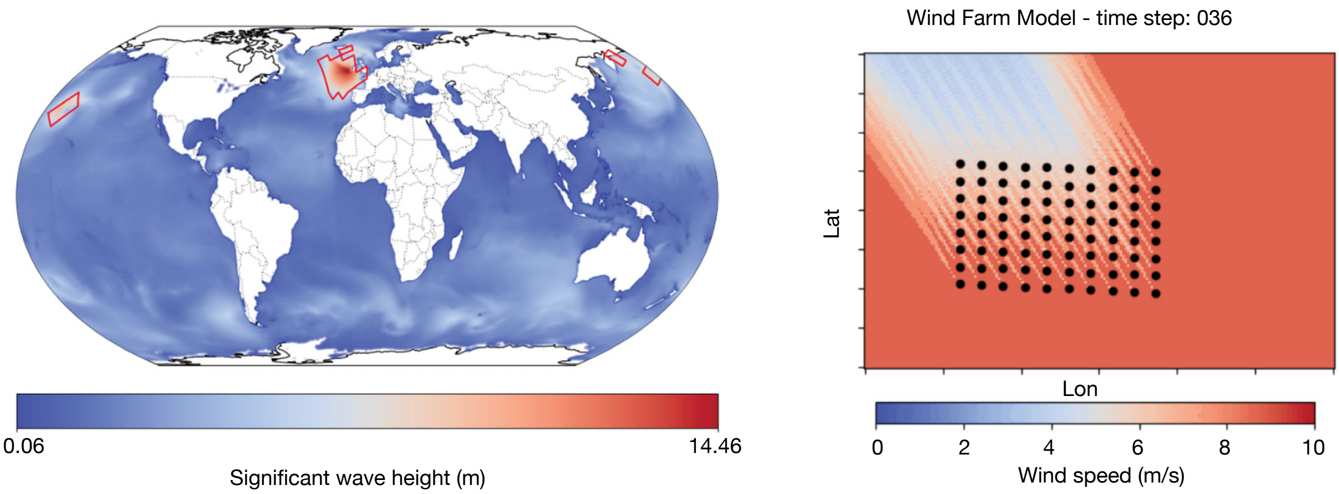

Figure 1 Outputs of the extreme event capture and wind farm parametrization plugins. In the left panel, the red polygons over the significant wave height field are regions where it exceeded 8 metres. The polygon in northwest Europe captures Storm Imogen that hit the UK on 2 February 2016. The plugin performs a coarsening of the signal before producing the output. The right panel shows the effect of wind farm parametrization on the wind field using the Jensen model for wakes. The plot shows the wind deficit behind individual turbines of a hypothetical wind farm. This plugin also computes power estimates which are not visualised here.

These efforts contribute to integrating Plume into research and operational workflows within ECMWF, making it a robust and well-embedded component of ECMWF's digital ecosystem. The framework is evolving to support internal developments, as well as the needs of Member States, Co-operating States and other external partners.

Next steps

The second half of DTWO will focus on the technical integration and evaluation of its use cases into the platform, and on finalising the implementation of the scientific workflows in each use case.

Beyond its coordination role on the technical integration, ECMWF will bring more contributions to the Wakes module. In particular, the architecture of Plume will be challenged to evaluate its suitability for two-way coupling of plugins, allowing them to act as source or sink terms within governing equations. The wind farm parametrization plugin will be used to test and guide the design of this new feature. A two-way coupling between plugins (dynamically loaded applications) and the model would create opportunities for more dynamic interactions within Plume and open new user journeys. For example, Plume could simplify parametrization development by enabling rapid and collaborative scheme iteration and calibration in a separate library, postponing the more complex model integration steps until the scheme is validated.

Funding acknowledgement

The DTWO project (grant agreement No 101146689) is funded by the European Union. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the Commission. Neither the European Union nor the granting authority can be held responsible for them.

Forecast performance 2025

Thomas Haiden, Matthieu Chevallier

ECMWF maintains a comprehensive range of verification statistics to evaluate the accuracy of its forecasts. Each year, a summary of verification results is presented to ECMWF's Technical Advisory Committee (TAC). Their views about the performance of the operational forecasting system in 2025 are given in the box.

The implementation of Cycle 49r1 of the Integrated Forecasting System (IFS) on 12 November 2024 has improved medium-range scores in the extratropics by several percent. It is the largest improvement in upper-air ensemble forecast (ENS) scores from a single upgrade since 2011 and has increased ECMWF's lead over other centres in terms of upper-air scores in the medium range. Cycle 49r1 introduced, among other changes, 2 m temperature assimilation in 4D-Var and the use of the Stochastically Perturbed Parametrizations (SPP) scheme. For surface parameters such as 2 m temperature and 10 m wind speed, the upgrade has reduced the number of large errors in the northern extratropics by about 10% in the ensemble forecast (Figure 1). While precipitation forecast skill has

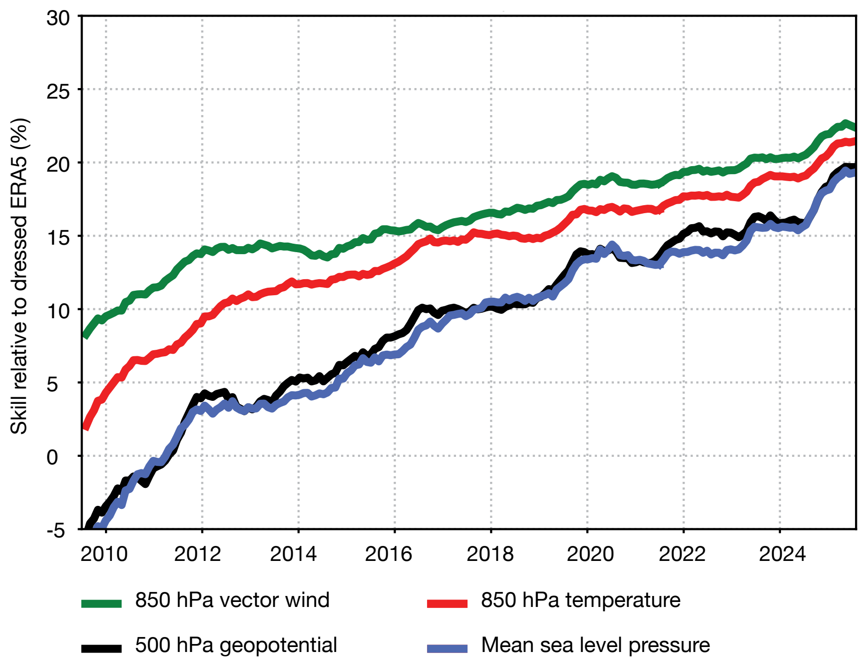

Figure 1 Skill of the IFS ENS relative to a Gaussian-dressed ERA5 forecast at day 5 in the northern extratropics based on the continuous ranked probability score (CRPS). Shown are 12-month running average values. The latest increase reflects model Cycle 49r1.

improved slightly (0.5–1%), scores for 2 m dewpoint and total cloud cover have been largely neutral. IFS Cycle 50r1, scheduled for spring 2026, is expected to deliver substantial improvements in these variables.

On 25 February 2025, the machine-learning (ML) model AIFS Single v1 was implemented, followed by the operational launch of AIFS-ENS on 1 July 2025. Both AIFS Single and AIFS ENS show superior skill for most upper-air and surface parameters compared to the IFS control and ensemble forecasts. Error reductions in the medium range are typically of the order of 5–15%. This includes a reduction in the number of large errors, for example in 2 m temperature, especially in wintertime stable situations. Precipitation is improved as well in the ML vs physics-based forecast by about 10% in the medium range, equivalent to 0.5–1 forecast day. Tropical cyclone (TC) position errors are reduced substantially in the AIFS compared to the IFS, whereas TC intensity errors are generally larger due to a weak-intensity bias.

IFS sub-seasonal forecasts in the extratropics show a continuing trend of increasing skill in week 2, and a marginally significant trend in weeks 3 and 4. On the seasonal timescale, due to the absence of strong forcing from Tropical Pacific sea-surface temperatures in 2025, forecast skill was modest, especially in boreal winter (December-January-February) 2024/25; however, large-scale boreal summer (JJA) temperature anomalies were captured. The strongly negative phase of the Indian Ocean Dipole (IOD) was well predicted.

IFS-COMPO is competitive with other air quality forecasts in predicting NO2, surface ozone, and PM2.5. Scores for column ozone and aerosol optical depth (AOD) indicate a slight improvement in forecast skill due to Cycle 49r1.

The complete set of annual verification results is available in ECMWF Technical Memorandum No. 931 on 'Evaluation of ECMWF forecasts', downloadable from https://www.ecmwf.int/en/publications/technical-memoranda.

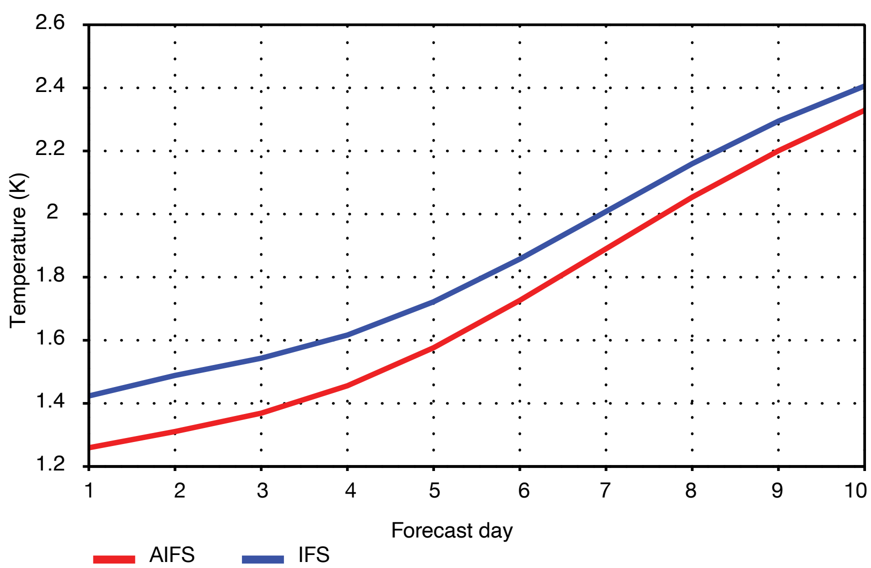

Figure 2 Comparison of physics-based and ML forecasts. Continuous ranked probability score (CRPS) of IFS ENS and AIFS ENS for 2 m temperature in the extratropics for the period August 2025–February 2026.

Additional resources

Verification as part of ECMWF's charts page: https://charts.ecmwf.int

WMO inter-comparison of global model forecast skill: https://wmolcdnv.ecmwf.int

WMO ocean wave model intercomparison results: https://confluence.ecmwf.int/display/WLW/WMO+Lead+Centre+for+Wave+Forecast+Verification+LC-WFV

List of 'Known IFS Forecasting Issues': https://confluence.ecmwf.int/display/FCST/Known+IFS+forecasting+issues

IFS cycle changes since 1985: http://www.ecmwf.int/en/forecasts/documentation-and-support/changes-ecmwf-model

Assessment of ECMWF's Technical Advisory Committee, 20–21 October 2025

With regard to its overall view of the performance of ECMWF's operational forecasting system, the Committee:

a) congratulated ECMWF on its 50th anniversary and is looking forward to many more years of successful and fruitful collaboration between ECMWF and Member and Co-operating States;

b) welcomed the positive benefits of 49r1. The revised assimilation of 2 m temperature, the switch from SPPT to SPP and the increase in resolution of EDA have, relative to the ERA5 benchmark, brought a large improvement in skill. Cycle 49r1 shows the largest improvement in skill from a single model cycle since the introduction of 46r1 several years ago;

c) congratulated ECMWF on maintaining and increasing their lead over other centres in upper air scores for the extratropics;

d) noted that ECMWF does not lead over all centres for all surface parameters, more so in the shorter range especially for precipitation. The difference between ECMWF and those centres who are leading in this respect has narrowed for 2 m temperature due to the introduction of 49r1;

e) welcomed the benefits 49r1 has brought for high-impact weather. For example, EFI skill has improved, especially for precipitation. Heavy summertime precipitation skill scores over Europe also seem to be starting to improve having been relatively stagnant for some time whilst some Member States noted an improvement in 10 m wind details, especially during winter. Cloudiness is also improving;

f) noted that skill scores for DestinE continuous Extremes digital twin (IFS at 4.4 km horizontal resolution) remain higher than for IFS. Since the introduction of 49r1, DestinE upper air scores have shown more improvement, implying 49r1 has an even greater benefit at higher resolutions;

g) welcomed tropical cyclone verification showing the smallest D+3 position error so far and a reduced central pressure bias whilst noting the ENS location forecast has moved from being under- to over-dispersive;

h) welcomed the benefits of 49r1 in IFS-COMPO (CAMS) for prediction of atmospheric composition;

i) congratulated ECMWF on the operational introduction of AIFS Single in late February 2025 and AIFS ENS in early July 2025. These are world-leading achievements and initial skill improvements from AIFS ENS are particularly encouraging. Information ECMWF has provided on verification and model characteristics to help users build trust has been welcomed;

j) noted that, compared to the physics-based forecasts, machine learning forecasts lead on skill for many parameters and are less jumpy;

k) recognised that the performance of AIFS is similar to several other machine learning forecasts, including a small decrease in skill during the last 12 months;

l) noted that, compared to IFS, AIFS leads for 2 m temperature and 10 m wind speed, especially over Europe. In particular, AIFS is better with 2 m temperature in stable wintertime synoptic situations and benefits AIFS brings to 10 m wind speed are larger over flatter, less complex terrain.

m) noted that AIFS shows a smaller position and track error for tropical cyclones than IFS, but has a greater intensity bias;

n) congratulated ECMWF on maintaining its long-term lead over other centres for significant wave height and its lead for peak period scores;

o) welcomed the sub-seasonal model showing some improvements in week 2. Noted that when considering warm summer and cold winter anomalies there is yet to be any statistically significant robust increase in skill for weeks 3 and 4;

p) noted that, although the transition from El Niño to La Niña and thence to neutral conditions was well-captured by the seasonal model, the amplitude of cycle was more muted in reality;

q) recognised that the lack of strong forcing from Pacific SSTs would impact seasonal forecast skill, skill being lower-than-average in such situations. Nevertheless, seasonal forecasts for winter 2024–25 and summer 2025 showed some skill over the northern hemisphere, but missed the cold winter anomaly over North America and were generally less skillful over the southern hemisphere, especially over Australia;

r) appreciated ECMWF's continued development of new diagnostics and products, and its very good support provided to Member and Co-operating States over the last year, with engagement via many mechanisms including online support, the annual UEF, online seminars, site visits and meteorological representatives at Member States.

AI Weather Quest at mid-year: global participation and first insights

Olga Loegel, Joshua Talib, Jörn Hoffmann, Frédéric Vitart

The AI Weather Quest ("the Quest") is an ECMWF-led international competition focused on sub-seasonal forecasting using machine learning (ML). It provides a transparent and operational-style benchmarking and knowledge sharing framework for evaluating ML approaches to forecasts at lead times of 3–4 weeks.

The initiative is endorsed by the World Meteorological Organization (WMO) under the WMO Integrated Processing and Prediction System (WIPPS), reinforcing its role as a coordinated global effort to explore the responsible integration of ML into forecasting systems.

The scientific design and evaluation framework have previously been described in a paper (Loegel et al., 2025; https://doi.org/10.1088/3049-4753/adf649).

Now midway through its first competitive year, the Quest has completed two 13-week forecasting periods, September to November 2025 (SON) and December 2025 to February 2026 (DJF), and is currently in its third period, March to May 2026 (MAM). Experience so far demonstrates strong international uptake, scientific relevance, and sustained participation.

Participation has remained consistently high (Table 1), with more than 25 teams submitting forecasts each week. Across the two completed periods, 42 teams, 194 participants and 74 distinct models have contributed.

	Participating teams	Participating team members	Participating models	Participating models eligible for full period-aggregated scoring
SON period	35	166	60	13
DJF period	37	177	66	18

Table 1 Participation metrics across the two completed competitive periods. The table summarises the number of participating teams, team members and submitted models during the September–November 2025 (SON) and December 2025–February 2026 (DJF) forecasting periods, including models eligible for full period-aggregated scoring.

Broad engagement across the weather and ML ecosystem

Participating teams represent a broad range of institutional backgrounds across the weather and ML ecosystem, spanning research organisations, universities, meteorological institutions, and industry (Figure 1).

Geographically, teams span Europe, China and the United States, with additional participation from Kenya, Niger, Morocco, Peru, South Korea, Canada and Brazil (Figure 2). This distribution demonstrates the increasingly global interest in applying ML to sub-seasonal prediction.

Beyond participation, several teams achieved full period-aggregated scoring by submitting forecasts consistently for all variables across a full forecasting period. Organisations whose teams met this eligibility criterion during the SON or DJF periods include meteorological institutions, universities and companies across multiple regions (Table 2).

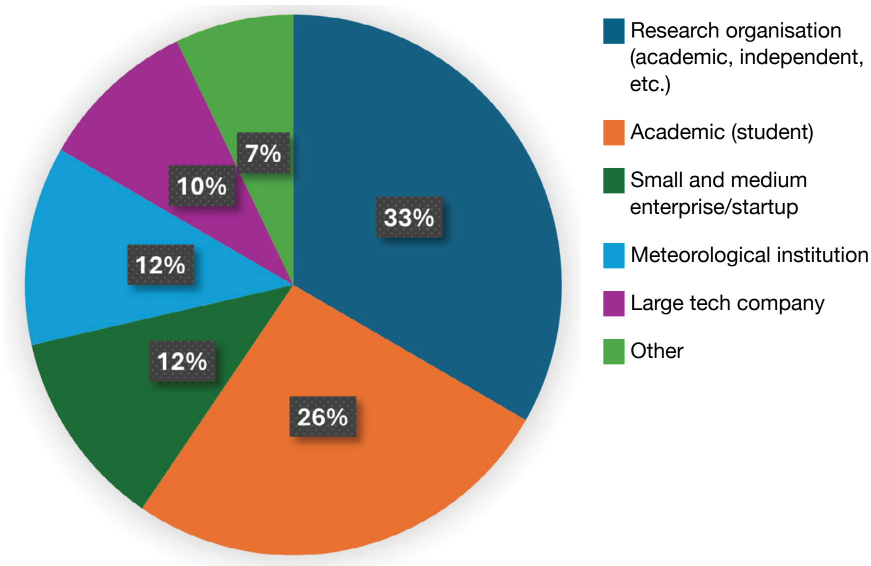

Figure 1 Distribution of participating teams by organisation type across the two completed competitive periods. Classification is based on the institutional affiliation of the team leader.

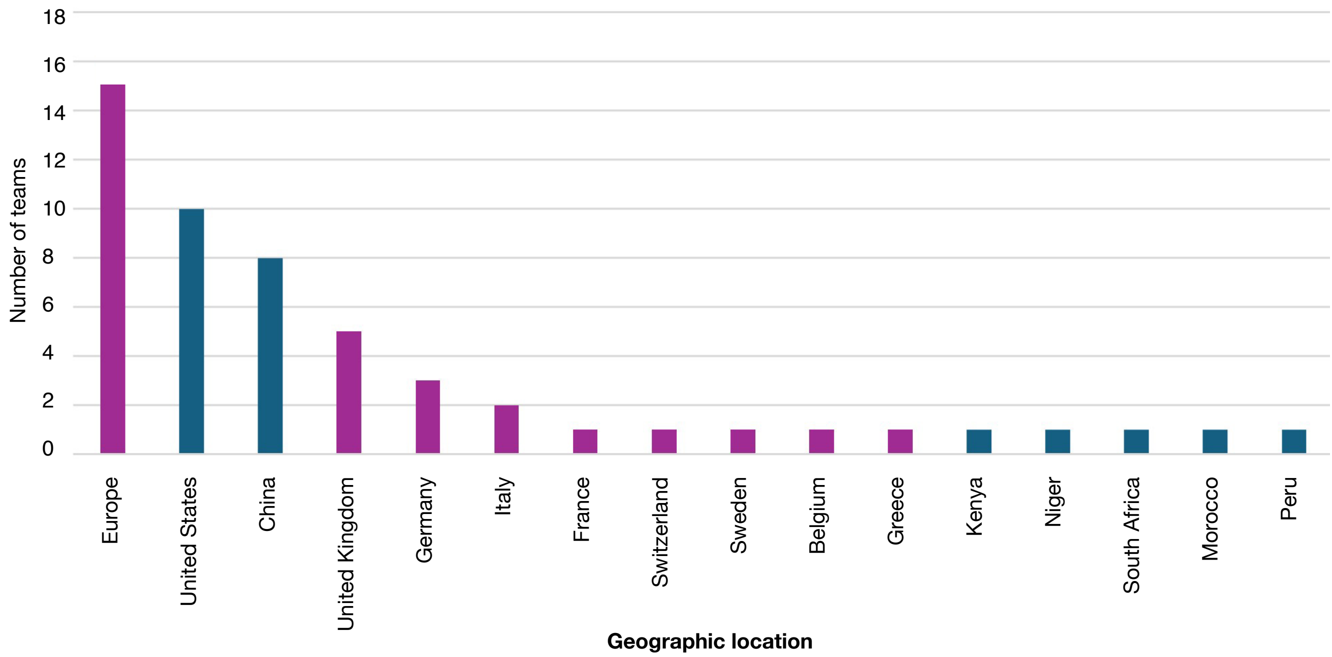

Figure 2 Distribution of participating teams by geographic region across the two completed competitive periods. Classification is based on the team leader's institutional affiliation. Purple bars represent European countries, aggregated in the first bar ("Europe") and also shown individually.

Organisation type	Full period-aggregated scoring eligible organisations (country)
Meteorological institution	• China Meteorological Administration (China) • European Centre for Medium-Range Weather Forecasts (international) • Intergovernmental Authority on Development Climate Prediction and Applications Centre (Kenya/Eastern Africa) • Jiangsu Climate Center (China)
Research organisation (academic/independent)	• Climate Modeling Alliance – CliMA (United States) • ETH Zurich (Switzerland) • Fudan University (China) • Instituto de Matemática Pura e Aplicada (Brazil) • Karlsruhe Institute of Technology (Germany) • Massachusetts Institute of Technology (United States) • Nanjing University of Information Science and Technology (China) • National Science Foundation National Center for Atmospheric Research (United States) • University of Toronto (Canada)
Academic (student)	• Gwangju Institute of Science and Technology (South Korea) • Harvard University (United States) • Imperial College London (United Kingdom) • Linköping University (Sweden) • Nanjing University (China) • University of Washington (United States)
Company (large tech/SME/ Startup)	• Microsoft (United States) • NVIDIA (United States) • Rhiza Research (United States) • WindBorne Systems (United States)

Table 2 Organisations whose non-anonymous teams met the eligibility criterion for full period-aggregated scoring during the SON or DJF periods. Teams listed here submitted forecasts consistently for all required variables across an entire 13-week forecasting period, making them eligible for full evaluation and leaderboard ranking.

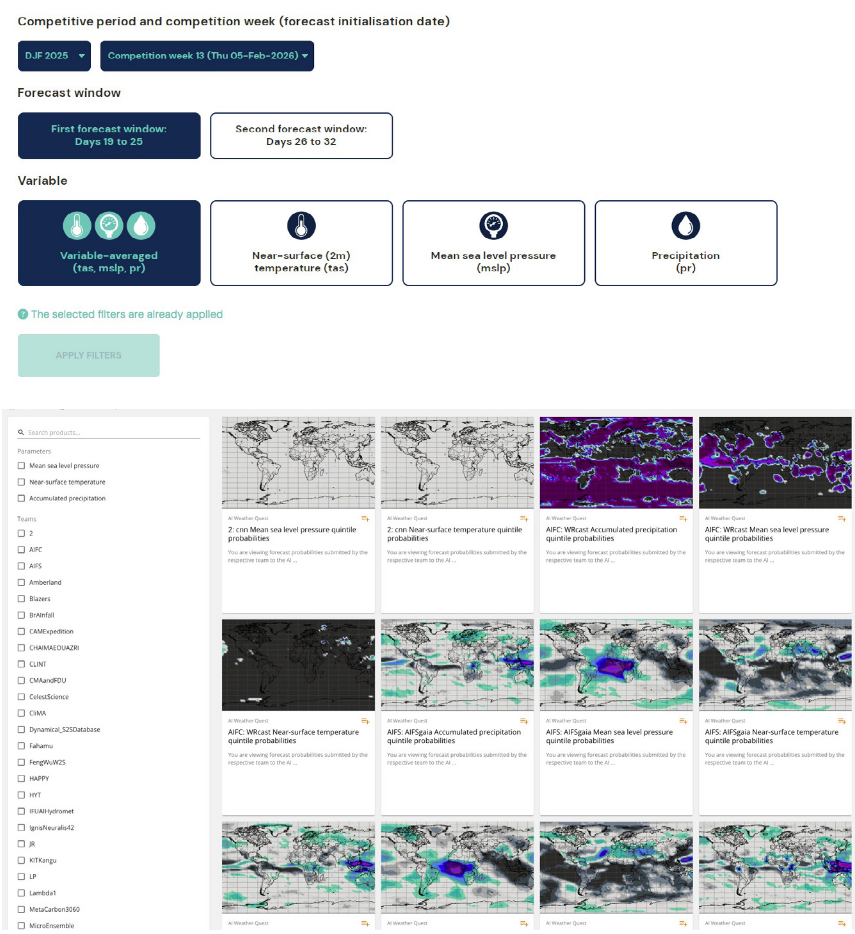

Figure 3 AI Weather Quest forecast visualisation portal. The portal provides weekly visualisations of submitted forecasts.

Variable-averaged RPSSs aggregated over the **DJF 2025** period up to and including **week 13**, for the **first** forecast window:

Team Rank	Model Rank	Team name / Model name	Period-aggregated RPSSs			
			tas	mslp	pr	averaged
1		MicroEnsemble				
	1	MicroDuet	0.086	0.079	0.083	0.083
	2	StillLearning	0.071	0.058	0.083	0.071
	6	Huracan	0.081	0.051	0.042	0.058
N/A		Dynamical_S2SDatabase				
	N/A	multimodel_mean	0.086	0.074	0.043	0.067
2		LP				
	3	LPM	0.089	0.047	0.065	0.067
3		AIFS				
	4	AIFShera	0.076	0.085	0.028	0.063
	7	AIFSgaia	0.062	0.054	0.014	0.043
	9	AIFSthalassa	0.060	0.048	0.014	0.041
4		CMAandFDU				
	5	FengshunHybrid	0.068	0.046	0.066	0.060
	8	FengshunAdjust	0.033	0.045	0.047	0.042
	12	Fengshun	−0.007	0.015	0.020	0.009
N/A		Dynamical_S2SDatabase				
	N/A	ECMWF	0.062	0.056	0.013	0.044

Show More ↓

Evolution of teams' rankings based on period-aggregated RPSS scores in the **DJF 2025** period, for **Variable-averaged RPSSs aggregated** and for the **first**

Figure 4 AI Weather Quest leaderboard showing model performance rankings. The leaderboard displays weekly model scores based on the ranked probability skill score (RPSS) for near-surface (2 m) temperature (tas), mean sea-level pressure (mslp), precipitation (pr) and averaged across the forecast variables.

Transparent benchmarking and structured exchange

The Quest operates with full transparency. Visualisations of submitted forecasts are published each week (Figure 3), while leaderboards provide continuously updated performance metrics (Figure 4). After each completed 13-week forecasting period, all submitted forecasts and team methodology summaries are made publicly available, enabling approaches to be openly compared over time.

The initiative also fosters structured dialogue between participants, for example through its end-of-period awards webinars. During the SON and DJF period awards webinars, five teams presented their methodologies and results to an international audience:

- MicroEnsemble (Microsoft, Harvard University, University of Toronto, IMPA, MIT and Rhiza Research)
- CMAandFDU (China Meteorological Administration and Fudan University)
- LP (Jiangsu Climate Center)
- AIFS (ECMWF)
- Fahamu (IGAD Climate Prediction and Applications Centre)

Leaderboards, submitted forecast visualisations and data, methodology summaries and webinar recordings are accessible on the AI Weather Quest website (https://aiweatherquest.ecmwf.int/).

First scientific insights

The first two competitive periods provide initial evidence on the performance of ML approaches at sub-seasonal timescales. Results suggest that improvements over dynamical models are most evident for precipitation, while added value for temperature and pressure remains limited to post-processing techniques. Many of the strongest-performing systems combine ML with dynamical forecast information, and surpassing climatology at lead time of 3–4 weeks remains challenging.

A detailed scientific analysis is available in the accompanying science blog on the ECMWF website (Talib et al., 2026; https://doi.org/10.21957/0fddea4898).

Looking ahead

The third competitive period is under way, with the fourth scheduled to begin on 14 May 2026. New teams are still welcome to join. Participation is designed to be straightforward. Forecasts are submitted using the AI-WQ package, a dedicated Python package that ensures full compatibility with the evaluation system and visualisation portal. Detailed installation instructions and user guidance are available in the official documentation (https://ecmwf-ai-weather-quest.readthedocs.io/en/latest/).

In parallel, the AI Weather Quest organisers are planning a joint special issue bringing together peer-reviewed manuscripts from participating teams. The issue will be published in the Royal Meteorological Society journal *International Journal of Climatology* and *Meteorological Applications*, titled Advances in Machine Learning for Weather and Climate: Modelling, Forecasting, and Applications (https://rmets.onlinelibrary.wiley.com/hub/journal/14698080/call-for-papers/si-2026-000201). The call for papers is currently open.

Acknowledgements

All authors and organisation of the AI Weather Quest are supported with funding from the European Union, provided to ECMWF under the Contribution Agreement between the European Union, represented by the European Commission, and ECMWF on the implementation of the Destination Earth (DestinE) initiative.

We gratefully acknowledge the hard work and enthusiasm of all AI Weather Quest participants. We also extend our thanks to the Advisory Board, whose interdisciplinary expertise in AI and meteorology has helped shape the structure and direction of this challenge.

The IFS AI Assistant

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

The Integrated Forecasting System (IFS) is at the heart of ECMWF's forecasting capability and is one of its most sophisticated software systems. Having evolved over several decades, it underpins a wide range of weather forecasting and Earth system applications. As a result, the associated codebase, documentation and workflows are extensive, distributed across multiple repositories and file types. The scale and diversity make it challenging for IFS users to navigate the system efficiently, especially when working with unfamiliar components or getting up to speed in a new area. To address this, a new IFS AI Assistant has been developed. The tool provides a local, retrieval-based assistant tailored to the IFS codebase and its documentation, helping users to explore and understand the system more effectively.

Why an AI assistant for the IFS?

The IFS code environment is unusual in several respects. While much of the system is written in Fortran, it also includes shell scripts, Python, notebooks, build files, README files, PDFs and other technical documentation. Knowledge is therefore spread across code and documents that need to be understood together rather than in isolation. At the same time, users need answers that are reliable, explainable and grounded in actual source material. In this context, a generic conversational AI tool is not sufficient: the system must operate on-premises to meet security requirements for closed-source code repositories, while also minimising hallucinations by managing the usage of Large Language Model (LLM) context and by grounding its responses on indexed and verifiable content. The IFS AI Assistant has been designed with these constraints in mind, offering reliable outputs.

Building on foundations laid by EU-funded Horizon Europe projects (namely EO4EU (https://www.eo4eu.eu/), MediTwin (https://meditwin-project.eu/) and DeployAI (https://www.deployaiproject.eu/), the assistant is intended to support IFS users in their daily work. Typical use cases include navigating large repositories, understanding how specific routines or modules are connected, locating relevant documentation, and supporting onboarding for new developers. In this way, the tool complements wider efforts to make the IFS software ecosystem easier to understand, maintain and evolve.

How the AI Assistant works

The IFS AI Assistant uses a hybrid retrieval-augmented generation (RAG) approach. At a high level, the architecture consists of four main parts (Figure 1):

- A frontend where users ask questions.
- A Representational State Transfer (REST) application programming interface (API) that handles requests.
- A search-and-retrieval engine that identifies relevant information.
- An indexing pipeline that parses repositories and builds the underlying search databases.

The indexing pipeline is a key component as it determines the quality and traceability of the information later supplied to the language model.

During indexing, the system scans selected repositories and classifies

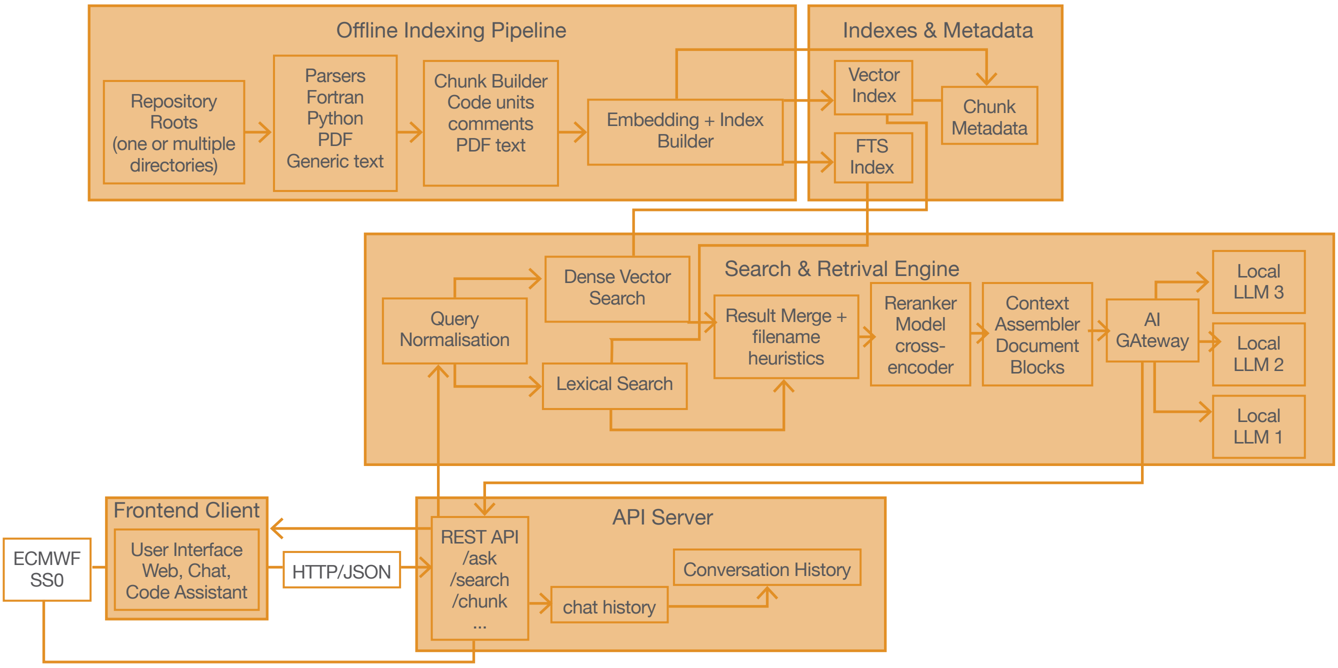

Figure 1 Main components of the IFS AI Assistant. The diagram shows the frontend, REST API, search-and-retrieval engine, and indexing pipeline, together with the lexical and vector databases used to ground responses in IFS code and documentation.

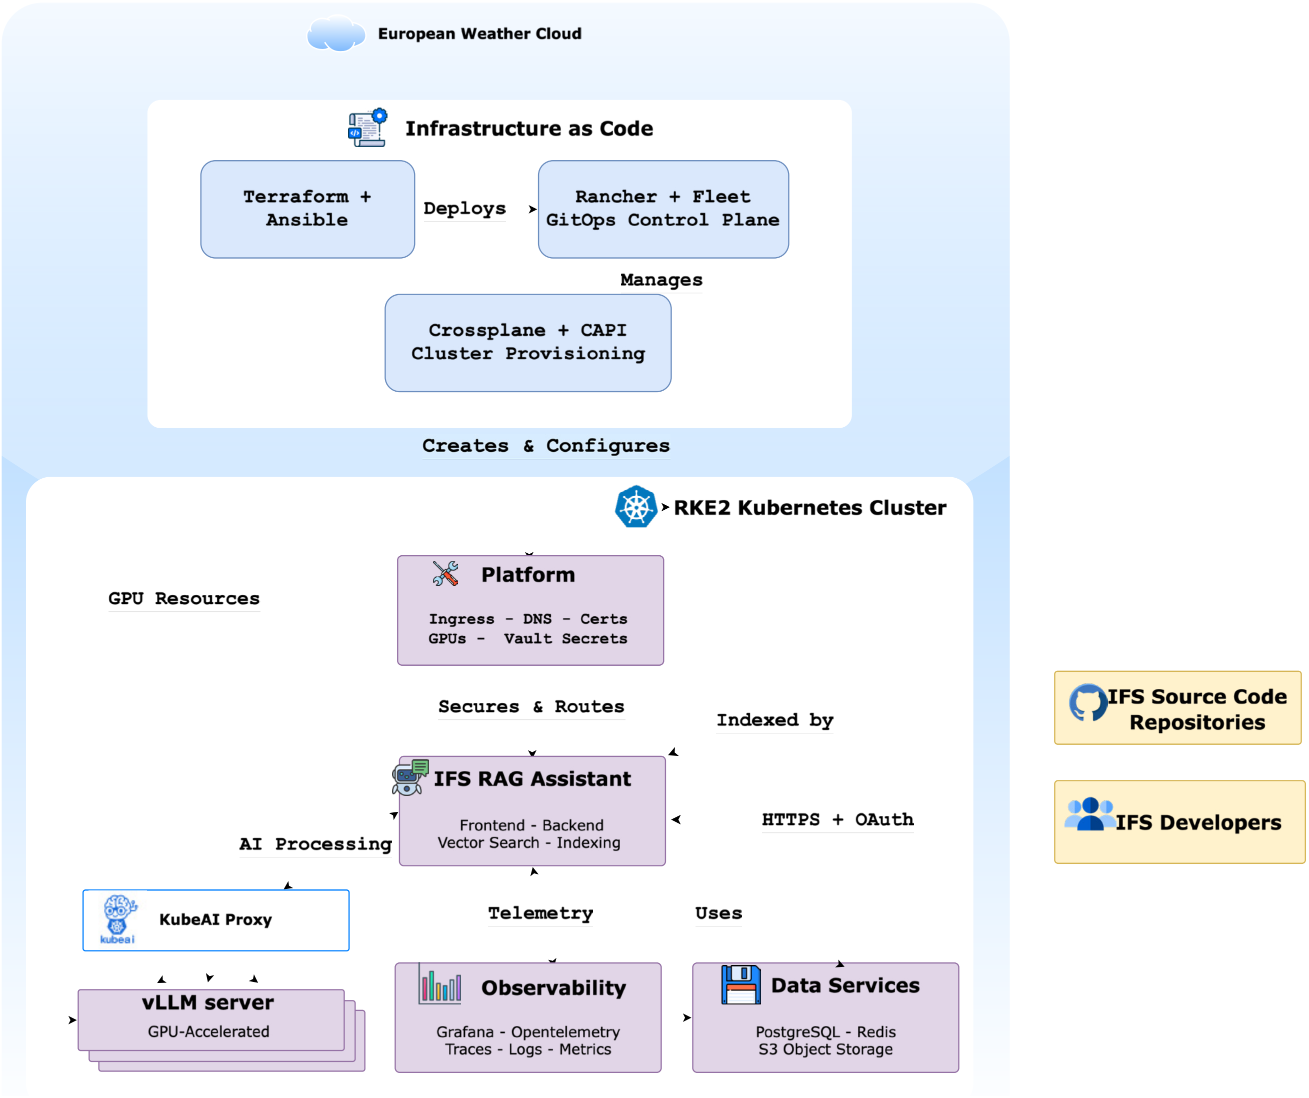

Figure 2 IFS AI Assistant infrastructure architecture. The diagram illustrates the GitOps pipeline and Kubernetes-based deployment, along with the components that make up the innovative tool. All components are deployed in the European Weather Cloud (EWC).

files into types such as Fortran, Python, PDFs, notebooks and generic text. Different parsing strategies are then applied depending on file type. For example, Fortran source code is parsed structurally, leveraging ECMWF's Loki (https://sites.ecmwf.int/docs/loki/main/index.html), to identify modules, subroutines, functions and programs, while Python files are analysed through their syntax tree. The system also captures contextual information such as comments, top-level file descriptions and synthetic file indexes. This helps to preserve relationships between different units in the same source file.

The indexed content is then stored in forms that support different retrieval strategies. Text is stored in a lexical search database, while embeddings are generated and stored in a vector database for semantic search. When a user submits a question, the assistant combines lexical and semantic retrieval, merges and re-ranks the results, and then assembles a context window for the generator model. This hybrid approach improves robustness: lexical search helps with exact terminology and identifiers, while dense retrieval helps capture conceptual similarity across code and documentation. This design supports both data governance and the requirement for explainable, repository-grounded answers.

The system is implemented using Python-based backend components, with LLM runtime support provided through KubeAI (https://www.kubeai.org/), which deploys, scales and routes traffic to AI inference backends. The assistant includes conversation memory and a serving layer exposed through a REST API and frontend. While the main user interface is a web-based chat, users can also interact with the system through command line tools.

Reliability, deployment and current use

The assistant is deployed on an RKE2 Kubernetes cluster (Figure 2) hosted on the European Weather Cloud (EWC; https://www.europeanweather.cloud/), reachable at https://ifs-assistant.apps.ifs-llm.ewcloud.host/ and integrated with the ECMWF Single Sign On (SSO) infrastructure. It is currently in an experimentation phase, piloted internally, expanding to a growing number of users. Early experience suggests that it can reduce the time needed to locate relevant code or documentation and can make it easier to explore unfamiliar parts of the IFS ecosystem (including both closed and open-source code). For large scientific repositories, the proposed approach offers a practical way to access dispersed technical knowledge through a user-friendly development aid.

Outlook

The IFS AI Assistant is still evolving, but it already illustrates how AI methods can support scientific software engineering. Rather than replacing developer expertise, it is intended to amplify it – helping users move more efficiently through a complex codebase while keeping answers anchored in trusted sources. As the system matures and incorporates new features, such as support for user-defined repositories, role-based access control for more fine-grained security and support for code assistants. This positions the assistant to become an increasingly valuable part of the wider IFS development environment, fostering both productivity and knowledge sharing across teams.

Monitoring climate change from space: the ECV programme

Joaquín Muñoz Sabater, Joao Martins, Eleanor Hansford, Erik Loebel

Climate change is no longer a distant threat. It is already reshaping ecosystems, economies and living standards across the globe. Understanding these changes and responding effectively requires high-quality, consistent, long-term climate data. The European Union's Copernicus Climate Change Service (C3S), implemented by ECMWF, contributes decisively to this effort. By enabling the production of quality controlled, validated climate data, C3S provides information that empowers scientists, policymakers and businesses to monitor trends and develop effective adaptation and mitigation strategies.

Why Essential Climate Variables matter

Essential Climate Variables (ECVs) are a critical component of C3S activities. These are variables considered fundamental for understanding the climate system and its changes over time. Examples of ECVs include sea-surface temperature, greenhouse gas concentrations, ocean colour and soil moisture. By tracking these variables over time, we can detect trends, identify anomalies, and evaluate the impacts of climate change on natural systems and human society. ECV quantities form the foundation of observation-based Climate Data Records (CDRs): long-term datasets derived from in situ networks and satellite observations that follow rigorous processing and validation to ensure consistency over decades. These records make it possible to assess climate change and variability, enabling evidence-based research, policies, and operational services.

The ECV programme

C3S operates an ECV programme organised around six thematic domains: Atmospheric Physics,

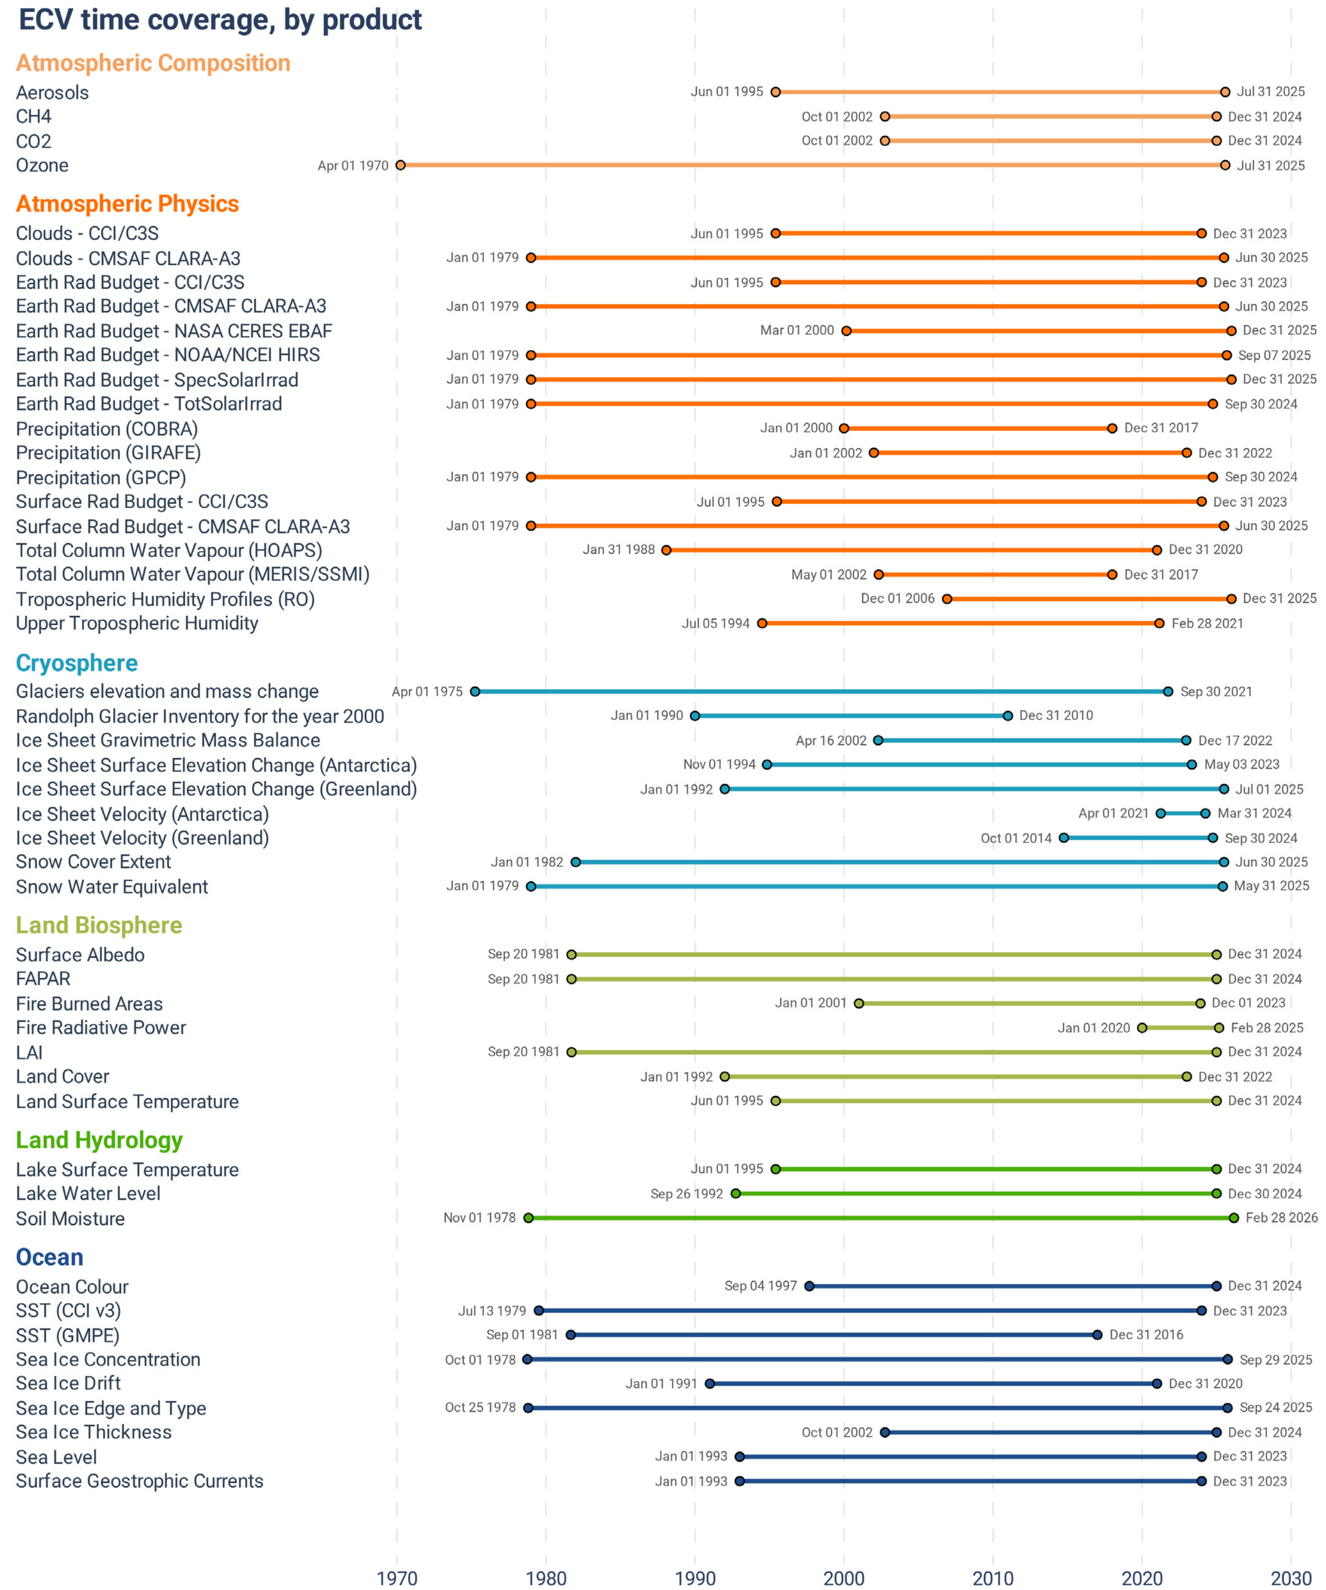

Figure 1 Temporal coverage of Essential Climate Variable (ECV) quantities published in the Climate Data Store in March 2026, organised by thematic domain. Each bar represents the period covered by a dataset.

Atmospheric Composition, Oceans and Sea Ice, Land Cryosphere, Land Hydrology, and Land Biosphere. By applying robust data processing to satellite observations, it provides a global perspective on climate that ground-based networks alone cannot deliver. Through the Climate Data Store, the C3S ECV programme provides access to 30 of the 55 ECVs defined by the Global Climate Observing System (GCOS), along with

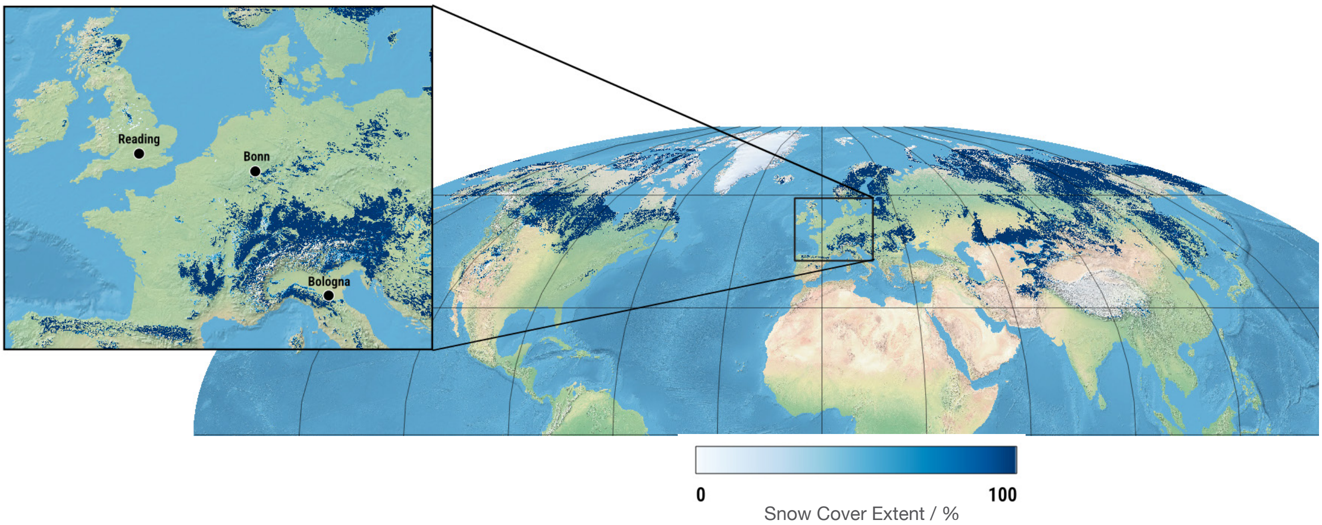

Figure 2 Fraction of pixel covered by snow on 1 March 2005, based on MODIS observations. In forested areas, a correction is applied to account for the masking effect of the forest canopy.

their associated quantities (see Figure1). In addition, the programme delivers value-added services that enhance the usability of its datasets. These include an independent Evaluation and Quality Control (EQC) activity, which ensures that datasets are reliable and suitable for scientific and policy applications. It is this blend of rigorous data and strong service support that makes the ECV programme unique.

From satellite measurements to climate information: the snow case study

An example of an ECV product is the Snow Cover Extent CDR. Snow cover plays an important role in the climate system by influencing the Earth's energy balance and freshwater availability. By reflecting a large fraction of incoming solar radiation, snow helps regulate surface temperatures, while seasonal snowpack acts as a natural water reservoir in many regions. The C3S dataset provides daily maps of snow-covered areas across the northern hemisphere derived from multiple satellites since the early 1980s (see example in Figure 2). These data can be used to monitor changes in the timing and duration of seasonal snow cover. For example, long-term analyses show that in many regions, spring snow melt occurs earlier than in previous decades. Such information helps assess the impacts of climate change on water availability and mountain ecosystems.

The importance of collaboration

Collaboration is central to the success of the C3S ECV programme. By working closely with agencies such as the European Space Agency (ESA) and the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT), C3S ensures that its datasets remain state of the art and seamlessly integrated within the broader Earth observation infrastructure. These partnerships allow the programme to incorporate data from new-generation satellite sensors, including most recent missions from EUMETSAT and the Copernicus Sentinel satellites, into the production chain of CDRs. Future Copernicus missions will further expand climate monitoring capabilities. With their focus on high-latitude observations, satellites like the Copernicus Imaging Microwave Radiometer (CIMR) and the Copernicus Polar Ice and Snow Topography Altimeter (CRISTAL) will enhance our monitoring of the Arctic and Southern Oceans, along with the Greenland and Antarctic ice sheets, significantly improving our capability to observe environmental changes in these remote regions. Close coordination with partner agencies remains essential to integrate these new data streams into operational CDR production and to further develop services that meet the diverse needs of C3S's broad audience.

Expanding the ECV programme

The ambition is to eventually cover all 55 GCOS ECVs to give a complete view of the climate system. Efforts will also focus on strengthening consistency across datasets through cross-ECV initiatives and improved methodologies. Reanalysis products will complement direct observations, particularly for variables that are difficult to measure from space. Where both observations and reanalysis are available, ensemble approaches will help increase confidence and reduce uncertainty. By linking ECV datasets to climate and policy-relevant indicators, the programme will provide timely and actionable information for adaptation and mitigation efforts. Enhanced data access, robust evaluation services and continued collaboration will ensure that this remains an authoritative source of climate knowledge, supporting Europe and the world in navigating the challenges of a changing climate for years to come.

The conclusion of the BUILDSPACE project

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

The BUILDSPACE ("Enabling Innovative Space-driven Services for Energy Efficient Buildings and Climate Resilient Cities") project concluded on 31 January 2026, completing a 36-month Horizon Europe Innovation Action, involving 14 organisations from eight countries. The project set out to leverage Earth observation (EO), Internet of Things (IoT) and drone survey data, together with digital twins and augmented reality (AR) and virtual reality (VR) technologies, to improve urban development and building construction practices, ultimately strengthening the resilience of cities to climate change.

A central challenge in urban resilience is bringing together heterogeneous information – such as satellite products, positioning services, local repositories and sensor observations – into workflows that decision-makers can use. BUILDSPACE addressed this challenge through an underlying core platform that integrates European Global Navigation Satellite System (EGNSS) services, Copernicus core services, and external sources, including repositories and IoT data.

On top of this platform, BUILDSPACE delivered five thematic services spanning building-scale digital twin generation and enrichment, and city-scale climate risk and energy-demand analysis (Figure 1).

The services include:

1. Digital Twin Generation, creating building-level digital twins linked with EGNSS and supporting interactive visualisation via AR/VR.
2. Digital Twin Enrichment, adding Simultaneous Location and Mapping (SLAM) techniques, Unmanned Aerial Vehicle (UAV) imagery and thermal analysis to assess building energy performance.
3. Building Environment Climate Scenarios, supporting stakeholders in understanding how building stock energy demand may evolve under climate change.
4. Urban Heat Analysis and Resilience, combining high-resolution urban heat information with sociodemographic layers to assess heat risk and social vulnerability.
5. Urban Flood Resilience, integrating Copernicus and local data with climate measurements to assess flood risk and support renovation guidelines and blue-green infrastructure planning.

The services allowed for the identification of environmental hotspots, assessment of risks, and generation of evidence to support interventions such as targeted building renovation or urban resilience planning.

Practical impact

ECMWF's contribution to BUILDSPACE was delivered largely through the European Weather Cloud (EWC), a community cloud computing platform jointly operated with EUMETSAT, producing several broad-range tangible and transferable results:

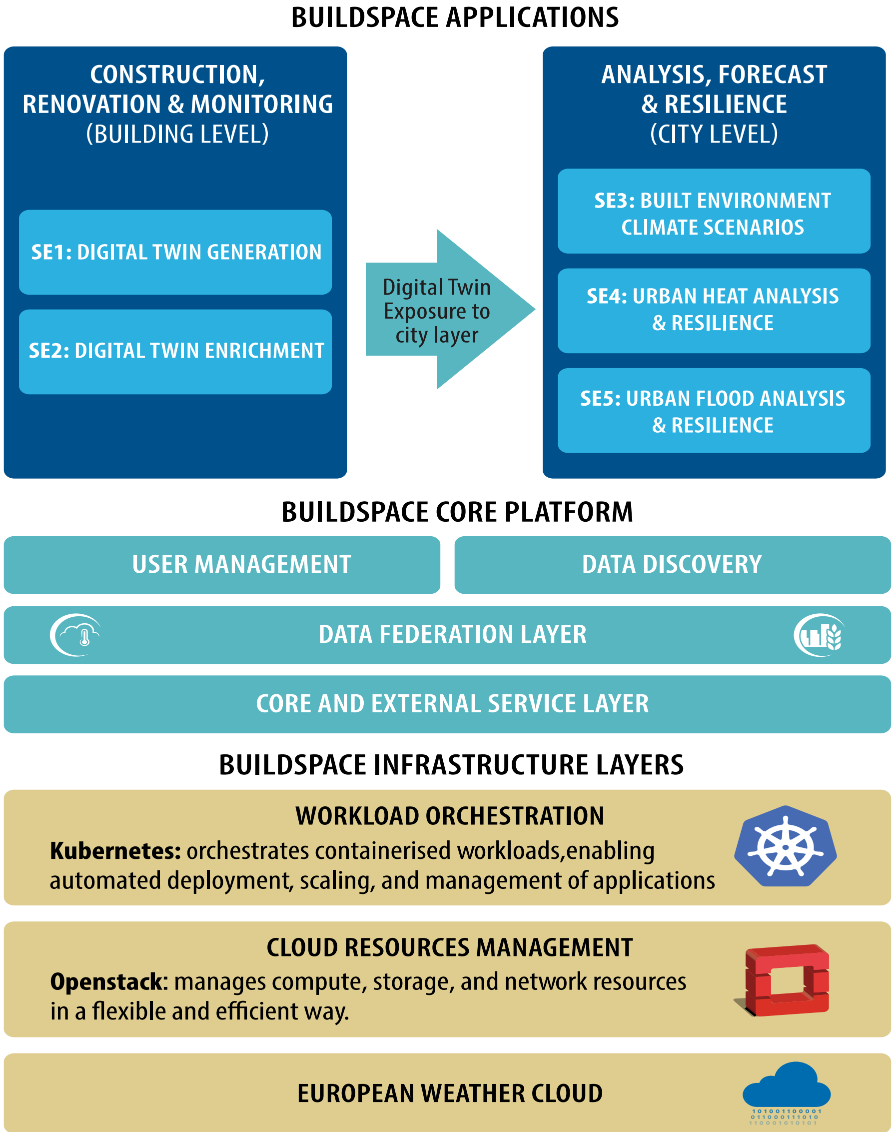

Figure 1 **BUILDSPACE services overview.**

- Infrastructure-as-Code (IaC) blueprints for rapid, highly available Kubernetes on the EWC. Building on earlier experience from the Horizon Europe-funded EO4EU project (https://www.ecmwf.int/en/newsletter/186/news/conclusion-eo4eu-project), ECMWF implemented a fully automated IaC approach that enables the deployment of highly available Kubernetes clusters on the EWC in a matter of minutes.
- Continuous integration and continuous delivery (CI/CD) patterns for secure, scalable geospatial web applications. ECMWF contributed to the scalable design and automated deployment of BUILDSPACE's Urban Heat Analysis and Resilience service on the EWC, establishing CI/CD workflows that are being reused and refined in ongoing projects. ECMWF also designed and managed a security-hardened GeoServer deployment and automated publication workflows using a GitOps approach.
- Copernicus and Destination Earth (DestinE) engagement to support interoperability and uptake. BUILDSPACE strengthened the use of ECMWF-contributed Copernicus datasets across services and helped align project outcomes with DestinE's evolution through engagement at community events and initiatives. These activities are continuing in the context of other ongoing projects (especially CLIMRES; https://cordis.europa.eu/project/id/101147777).

Demonstrations and urban resilience outcomes

BUILDSPACE services were demonstrated in four pilot locations across Europe with different climate profiles (Figure 2): Poland (construction support for facilities at the University of Warsaw and the Poznań University of Technology), Latvia (Riga – energy demand forecasting and climate scenario analysis), Greece (Piraeus – building and urban heat risk assessment), and Slovenia (Ljubljana – urban flood resilience).

Figure 2 Map of the four BUILDSPACE pilot sites across Europe.

Together, these pilots showcased how digital twins and EO-driven analytics can translate into practical outputs such as hotspot maps for heat risk, evidence for renovation planning, and flood-risk assessments supporting resilience measures and blue-green infrastructure planning.

BUILDSPACE also contributed to COSPACE, a clustering initiative under the Horizon Results Booster programme. Building on the technical work, the initiative produced a joint policy brief recommending that decision-makers enhance data-sharing frameworks and interoperability, invest in artificial intelligence (AI)-driven climate services, promote nature-based solutions for urban resilience, strengthen climate-resilient agriculture, and expand climate monitoring and early warning capabilities.

From project to operational deployment

Looking forward, one practical objective is to bundle and maintain the BUILDSPACE cloud-native patterns (IaC, GitOps, CI/CD and observability) into a standardised "project deployment kit" for EWC-based services, accelerating delivery while improving reliability and repeatability.

Funding acknowledgement

BUILDSPACE (grant agreement 101082575) was funded by the European Union. Views and opinions expressed are those of the authors and do not necessarily reflect those of the European Union or the granting authority. Neither the European Union nor the granting authority can be held responsible for them.

Implementation of AIFS v2

Mariana Clare, Sara Hahner, Harrison Cook, Gert Meertes, Josh Kousal, Annelize van Niekerk

ECMWF implemented its machine learning weather forecasting system, the Artificial Intelligence Forecasting System (AIFS), operationally in 2025. It will now upgrade the system with the implementation of AIFS v2, introducing simultaneous improvements to both the single and ensemble configurations. The new version expands the range of predicted variables, adding waves, additional land-surface fields and new pressure levels in the stratosphere. These developments, along with technical changes, allow the model to capture features that were not previously represented directly in AIFS forecasts, while also supporting improvements in forecast skill and physical consistency. Additionally, AIFS v2 adopts MultIO as the encoding pipeline, ensuring consistency with the upcoming migration to GRIB2 in the Integrated Forecasting System (IFS) Cycle 50r2.

AIFS v2 will be implemented alongside IFS 50r1, ensuring compatibility with the latest operational analysis and allowing it to benefit from improvements introduced into the IFS cycle. As in previous cycles, both AIFS models have been developed using the open-source Anemoi framework and have been released as a model on the ML model-sharing platform Hugging Face.

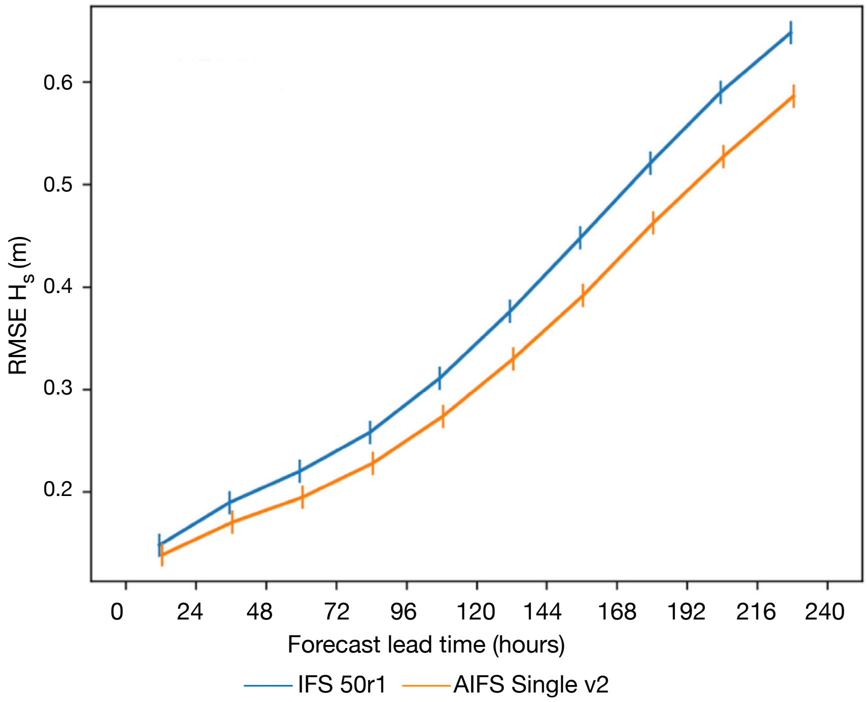

Figure 1 Root mean square error (RMSE) in significant wave height (H_S) over forecast lead time for the IFS 50r1 and AIFS Single v2 models. Error is with respect to satellite altimeter observations.

Expanded Earth-system representation

AIFS v2 continues moving toward a more complete Earth-system representation, extending the representation of land and atmospheric variables and introducing a representation of waves. With this new capability, AIFS forecasts provide key wave parameters such as significant wave height, mean wave direction and mean wave period. Verification shows a substantial improvement in medium-range wave forecast skill compared with IFS 50r1, with significant wave height errors reduced by around 10% (see Figure 1). AIFS v2 also introduces snow cover forecasts. Verification indicates that AIFS forecasts show improved performance compared to IFS 50r1, with the predicted snow cover generally closer to observations. In the stratosphere, the addition of variables at 10 hPa along with other small modifications has improved representation at higher levels. Forecast skill for variables at 50 hPa and 100 hPa has increased substantially in both the northern and southern hemispheres, with a 3-day improvement in skill between AIFS Single v2 and AIFS Single v1. The addition of the pressure level also allows AIFS v2 to represent phenomena such as sudden stratospheric warmings.

Improvements in forecast realism in the AIFS Single

Forecasts from AIFS Single v2 appear partially less smooth than those from v1, reflecting increased dynamical variability. The representation of vertical velocities has also improved, with the Hadley circulation and associated convection cells now correctly formed and their strength consistent with physical expectations (see Figure 2). These improvements did not translate directly into better verification scores, emphasising the need to look at diagnostics when seeking to improve physical consistency in the AIFS.

Improvements in forecast realism in the AIFS ENS

Whilst the core architecture of the AIFS Single remains unchanged as part of the upgrade, there have been several methodological improvements to the AIFS ENS. A new multi-scale loss function (see Lang et al. 2025 DOI., https://doi.org/10.48550/arXiv.2506.10868) replaces the reference-state truncation approach used previously. Changes have also been made to the neural network connectivity within the ensemble model where additional edges in the decoder and new edge features have been introduced. Finally, the same physically motivated bounds (e.g. non-negativity and fractional limits) used by the AIFS Single model are now imposed in the AIFS ENS, improving physical consistency and reducing the spurious small precipitation values sometimes observed in AIFS ENS v1. These changes reduce artefacts in AIFS ENS forecasts and improve the representation of atmospheric structures across different spatial scales.

Adaptation to IFS Cycle 50r1

Because AIFS forecasts are initialised from the operational IFS analysis, changes to the operational IFS cycle can impact AIFS performance. Tests show that the AIFS forecast skill is improved when more recent operational analysis data is included in the training. However, some IFS cycle changes can have larger impacts. For example, substantial updates to the analysis in IFS Cycle 50r1 caused noticeable degradations in AIFS v1 forecast skill when it was initialised with 50r1 analysis, with errors in 2-metre temperature increasing by up to 20%.

To address this, AIFS v2 has been fine-tuned on both current operational analysis and available experimental 50r1 data (mid-May–September 2024 and December 2024). This fine-tuning allows the AIFS to adapt to changes in the structural and statistical characteristics of the IFS 50r1 analysis. As a result, there is no longer a degradation from using 50r1 analysis as initial conditions, although a small negative impact remains for 2-metre temperature in the Arctic. This likely reflects limited training data for winter sea-ice coupling and is expected to diminish as more 50r1 data become available.

Improvements in encoding and preparation for GRIB2

AIFS v2 uses MultIO as the encoding and data processing layer for writing directly to the Fields Database (FDB). This ensures identical GRIB headers for the initial state and forecast fields and simplifies the operational pipeline, reducing encoding inconsistencies between the IFS and AIFS. Additionally, all fields have been brought into compliance with WMO units and parameters. While AIFS v2 is not yet written using the latest version of the GRIB2 standard, this first step will help ensure a seamless transition in the upcoming migration and implementation of IFS 50r2.

Looking ahead

AIFS v2 marks another step in the operational development of AI-based forecasting at ECMWF. Continued evaluation and development will further refine the system and help to strengthen the role of AI within ECMWF's evolving forecasting framework.

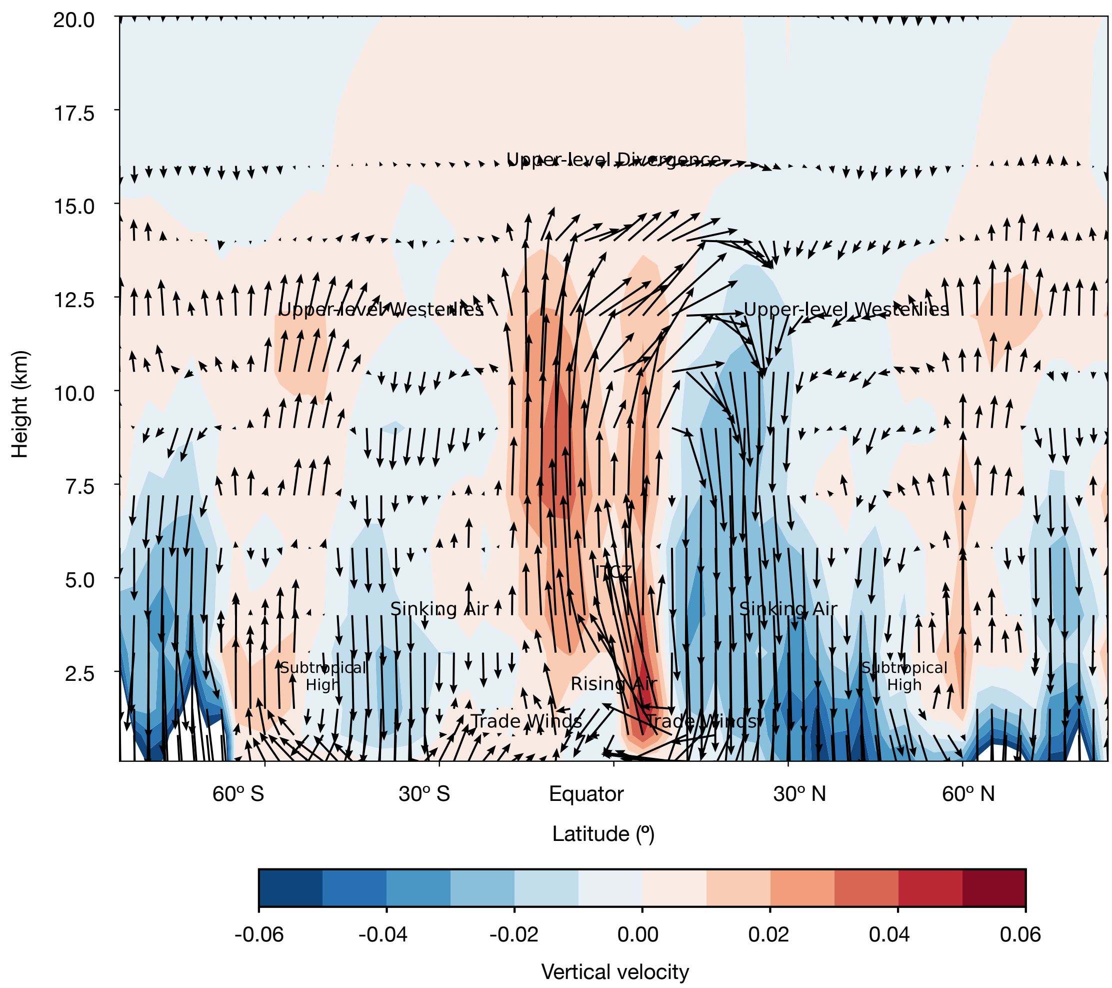

Figure 2 Diagnostic of the Hadley Circulation for AIFS Single v2, showing clearly formed convection cells with air rising and sinking uniformly across the vertical profile, unlike in AIFS 1.1. Vertical velocities are of expected strength whereas previously they were too strong at the top of the atmosphere.

OpenIFS now open source

ECMWF has made its portable global forecasting model, OpenIFS, fully open source for the first time, marking a significant step in its commitment to open science and collaboration.

OpenIFS is a research, education and training version of the Integrated Forecasting System (IFS). Developed over decades, the IFS is a shared achievement of ECMWF's Member and Co-operating States, built through sustained investment, scientific excellence and operational expertise. Its evolution has been shaped by close collaboration with national meteorological services, notably Météo-France, with which ECMWF has co-developed the IFS/ARPEGE system for nearly 40 years.

Until now, OpenIFS has been available under licence, typically through institutions. By removing this requirement, ECMWF is opening access to a much wider community of users.

The move simplifies access and removes administrative barriers, enabling researchers, educators and developers to work more freely with the model. Users can now download, modify and run the code without restriction, including on local systems such as laptops. This flexibility supports a broader range of applications, from academic research and teaching to exploratory development and industry collaboration.

Making OpenIFS open source also strengthens scientific transparency and reproducibility. Researchers can cite the exact version of the code used in their work and share it alongside publications, including in open-access journals. At the same time, all users will have access to a consistent and up-to-date version of the model, supporting more robust comparisons and collaboration across institutions.

The change will help bring OpenIFS closer to the operational IFS. While the operational system is updated annually, OpenIFS releases have historically lagged behind. Open sourcing the model will enable more frequent updates and reduce the gap between research and operations, ensuring that users can work with code that more closely reflects the current forecasting system.

Since its launch in 2011, OpenIFS has supported a growing international community of users. It has become widely used in universities for teaching numerical weather prediction and atmospheric modelling, helping to bridge the gap between theory and practical application. The model has also contributed to a range of research activities, including climate and atmospheric modelling and atmospheric composition research.

OpenIFS has already demonstrated value beyond education and research. It has supported developments that improve computational efficiency and has contributed to innovations that feed back into operational forecasting. By making the model openly available, ECMWF aims to accelerate this cycle of innovation and strengthen links between research and operations.

The transition to open source also aligns ECMWF with broader practices across the global weather and climate community, where open models are increasingly used to foster collaboration and innovation.

Access OpenIFS

The OpenIFS source code is now available on GitHub, alongside documentation and supporting materials, providing immediate access for users worldwide: https://github.com/ecmwf-ifs/openifs

New observations October 2025 – March 2026

The following new observations have been activated in the operational ECMWF assimilation systems during October 2025 – March 2026.

Observations	Main impact	Activation date
Radio occultation bending angles from PlanetIQ YAM-8	Temperature and winds in upper troposphere/lower stratosphere	13 Nov 2025
All-Sky Radiances (ASRs) from Meteosat-12 (replacing Meteosat-10)	Tropospheric humidity and wind	20 Nov 2025
Radio occultation bending angles from Metop-D	Temperature and winds in upper troposphere/lower stratosphere	11 March 2026

Coupled data assimilation in 50r1

Phil Browne, Patricia de Rosnay, Tony McNally, Hao Zuo, Alan Geer, Tracy Scanlon, Sarah Keeley, Kristian Mogensen

Outer-loop ocean–atmosphere coupled data assimilation will be fully integrated into ECMWF's operational systems with the Integrated Forecasting System (IFS) Cycle 50r1. The upgrade introduces interactive 3D ocean and sea-ice components coupled to the existing atmosphere, wave and land components within the numerical weather prediction (NWP) data assimilation system. This milestone reflects extensive development across ECMWF to operationalise an approach first demonstrated in the coupled reanalysis (CERA) systems. For the first time, we will create a unified Earth system analysis to initialise our forecasting system.

As well as aligning the analysis model with that of the forecasts, we will soon actively assimilate oceanic observations across all analysis and forecast systems. These include in situ subsurface profiles of temperature and salinity from Argo floats, moored buoys, ships, air-dropped instruments, and elephant seals with sensors (Figure 1). The system also actively assimilates sea-level anomaly data from altimeters.

Making better use of interface observations

One of the major benefits of the coupled data assimilation is that, for the first time, we can fully utilise so-called interface observations that are sensitive to both the ocean surface and atmosphere. Going well beyond

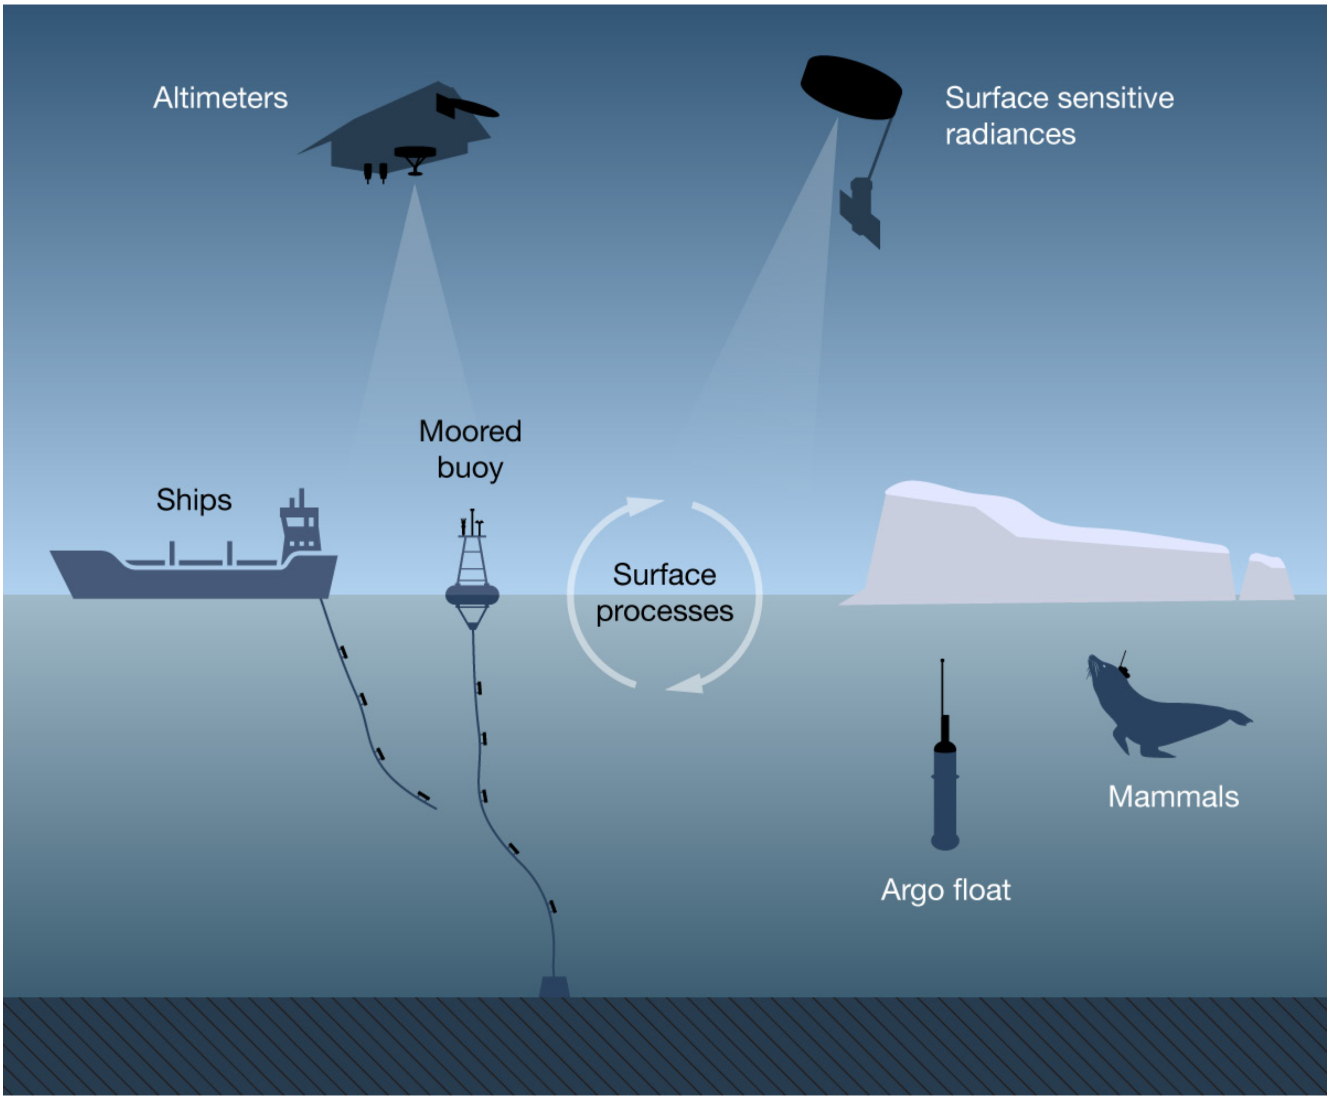

Figure 1 An extended network of observations feeds the coupled system, linking ocean and atmosphere. Satellites, ships, buoys, Argo floats and instrumented marine mammals provide complementary data to constrain the Earth system and allow improved representation of surface processes.

the CERA reanalyses – ECMWF's earlier coupled reanalysis system (Schepers et al. 2018; https://doi.org/10.21957/sp619ds74g), we are now extracting additional information from satellite radiance observations to impact not only atmospheric analyses but also the 3D ocean and sea-ice fields. The system uses observations from both microwave and infrared instruments on low Earth orbit and geostationary satellites

Previously, ECMWF ran a fixed reanalysis and analysis ocean system, last updated in 2016 with the ocean reanalysis ORAS5. With Cycle 50r1, the coupled analysis and the daily forced ocean and sea-ice analysis are part of the operational cycle, enabling regular updates of the ocean component with each IFS release. Although the system remains anchored to the new ocean reanalysis ORAS6, it will be possible to add new observation types, modify perturbation methods, and generally respond to user requirements. Anchoring to ORAS6 also ensures consistency with re-forecasts, as well as drastically reducing testing requirements and therefore speeding up the research-to-operations process for cycle development.

Delivering better forecasts

These changes bring benefits for forecasts. They improve consistency in the boundary layer and allow better use of surface processes within 4D-Var. For example, when using scatterometers in the data assimilation system we are now able to use surface stress relative to the ocean currents to compare with satellite measurements. Similarly, ocean currents are present in the wave data assimilation, allowing waves to diffract as they do in the real world. The system reduces reliance on external data products and provides the ability to extract more timely information. In the case of sea-surface temperatures, information could be up to 69 hours more timely.

The coupled system also strengthens cross-component consistency. Ocean subsurface observations of temperature and salinity can detect and verify changes in the atmosphere, while satellites and traditional atmospheric observations can reveal and verify changes in the ocean. In this way, each component now helps constrain the quality of the other. In addition, the system now fully utilises information from interface observations that was previously discarded.

The lower boundary condition for the analysis is no longer the sea or ice surface; it is now the bathymetry and geothermal heating at the ocean floor. This shift increases the number of prognostic variables available to be used within the assimilation, such as sea-surface salinity, sea-ice albedo, and more. As a result, this opens new avenues of research to extract more data from existing observations and to allow more data to be used.

With the long journey from research to operations nearing completion, this new platform will allow exciting new scientific developments. For example, work will exploit data from future satellite missions such as the Copernicus Imaging Microwave Radiometer (CIMR) and CRISTAL, with impacts across the Earth system. Research has started exploring the direct assimilation of satellite microwave radiances to initialise the ocean salinity, supplementing the existing sparse conventional observations. Future coupling also extends to land–atmosphere (see article by Herbert et al. in this Newsletter).

In the new age of data-driven forecasting, we are continuing to expand and develop the physics-based system. However, the benefits of this will be felt in the data-driven systems, too, by providing better initial conditions and more consistent coupled analysis datasets for training.

Forecasting for Storm Harry in the Mediterranean

Linus Magnusson, Jean Bidlot, David Lavers, Nigel Roberts, Josh Kousal, Sara Hahner

The beginning of 2026 was characterised by cold weather in the Nordic countries, while south-western Europe was extremely wet and hit by a number of storms. In January, Storm Harry impacted several Mediterranean countries with extreme precipitation, strong winds and destructive waves. The hardest-hit regions included southern Italy, where persistent gale-force winds brought significant wave heights exceeding 8 metres, and some parts of Sicily observed nearly 500 mm of rainfall in 72 hours (19–21 January). Greece and Tunisia also experienced extreme rainfall. In this short report, we focus on the prediction of the large-scale flow and ocean waves.

On 16 January, an upper-level trough amplified over the eastern Atlantic and formed a cut-off low on 18 January over the Iberian Peninsula. The cut-off low then propagated towards the south-east, settling over Tunisia on 20–21 January (Figure 1a). At the surface, a strong pressure gradient developed over the Ionian Sea and the central Mediterranean between a surface low associated with the upper-level cut-off low and a high-pressure system over the Balkans.

To investigate the predictability of the upper-level trough, we applied a curvature diagnostic, CURV (Curvature Using Radial Variation), recently developed at ECMWF. The diagnostic was computed for 500 hPa geopotential height, with values at each grid point ranging between -16 (closed cyclone) to 0 (no curvature) to +16 (closed anticyclone). For both the Integrated Forecasting System (IFS) and Artificial Intelligence Forecasting System (AIFS) ensembles, the signal for cyclonic conditions (negative CURV) started to appear around 10 January, and both confidently predicted a very cyclonic pattern in forecasts from 12 January onwards (Figure 1b). This indicates that the large-scale flow was well predicted for this case.

The Extreme Forecast Index (EFI) for maximum significant wave height from the IFS ensemble indicates an early signal of high waves across a large part of the central Mediterranean on 20 January (Figure 2), already evident six days before. In a corresponding forecast evolution plot for waves south of Sicily (Figure 1c), the signal starts to appear around 12 January. The gain of the signal in the waves corresponds well with the appearance of the signal for the cut-off low as seen in the forecast evolution for z500 CURV (Figure 1b).

The upcoming AIFS version 2 will introduce several wave parameters (see article by Clare et al. in this Newsletter). The forecast maps in Figure 3 show significant wave height on 20 January at 12 UTC in the IFS analysis 36 h forecasts from the IFS control and AIFS Single v2, with the AIFS model predicting waves reaching above 8 metres south of Sicily. The structure of the forecast field from AIFS agrees well with the IFS forecast and the analysis.

This is supported by the altimeter observations, with the altimeter to the west of the storm reporting waves of 7.5 metres. Measurements from a buoy south of Sicily, operated by the Italian National Institute for Environmental Protection and Research (ISPRA), also reported significant wave height of around 8 metres at this time. Looking at the forecast evolution plot for significant wave height for AIFS Single v2, we also see a consistent prediction of high waves (7–8 metres) from 12 January, in line with the prediction of the cyclonic conditions.

In summary, the high predictability of the upper-level trough in the medium-range forecasts translated well into the ocean wave forecast. This was probably due to the large-scale nature

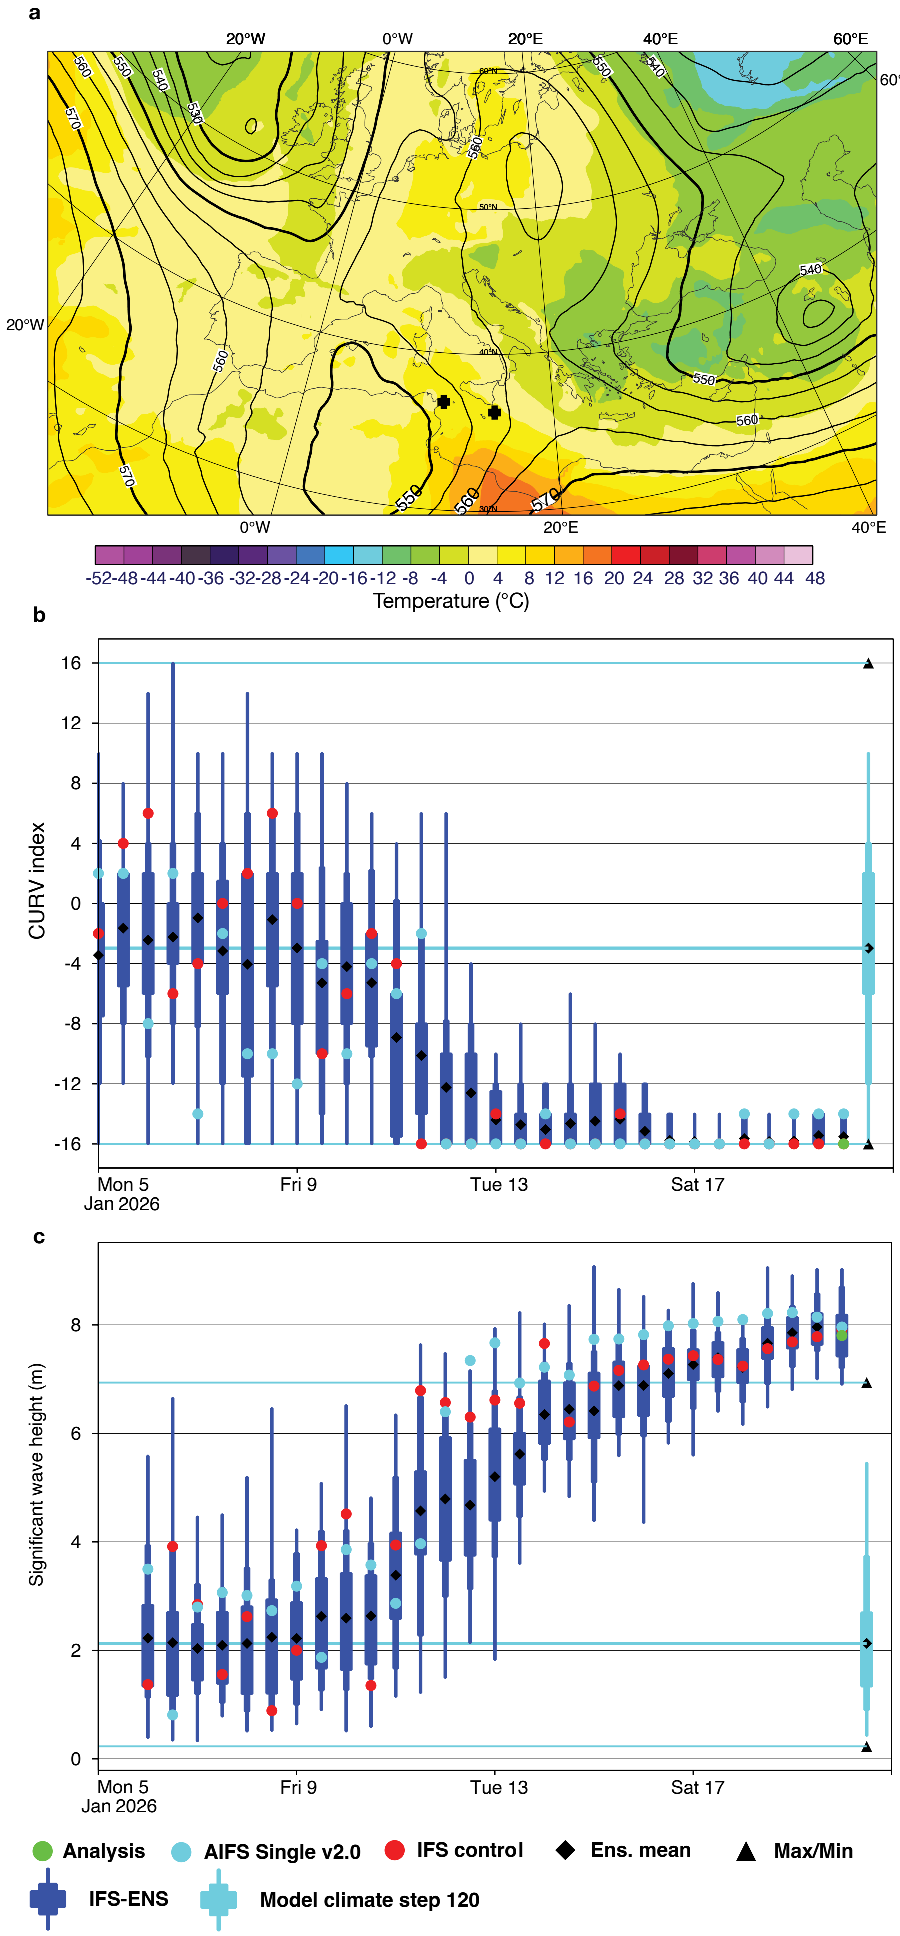

Figure 1 (a) Analysis of z500 (contours) and T850 (shading) from 20 January 00 UTC (b) forecast evolution for z500 CURV (-16 = cyclone) over northern Tunisia on 20 January and (c) forecast evolution of maximum significant wave height south of Sicily on 20 January.

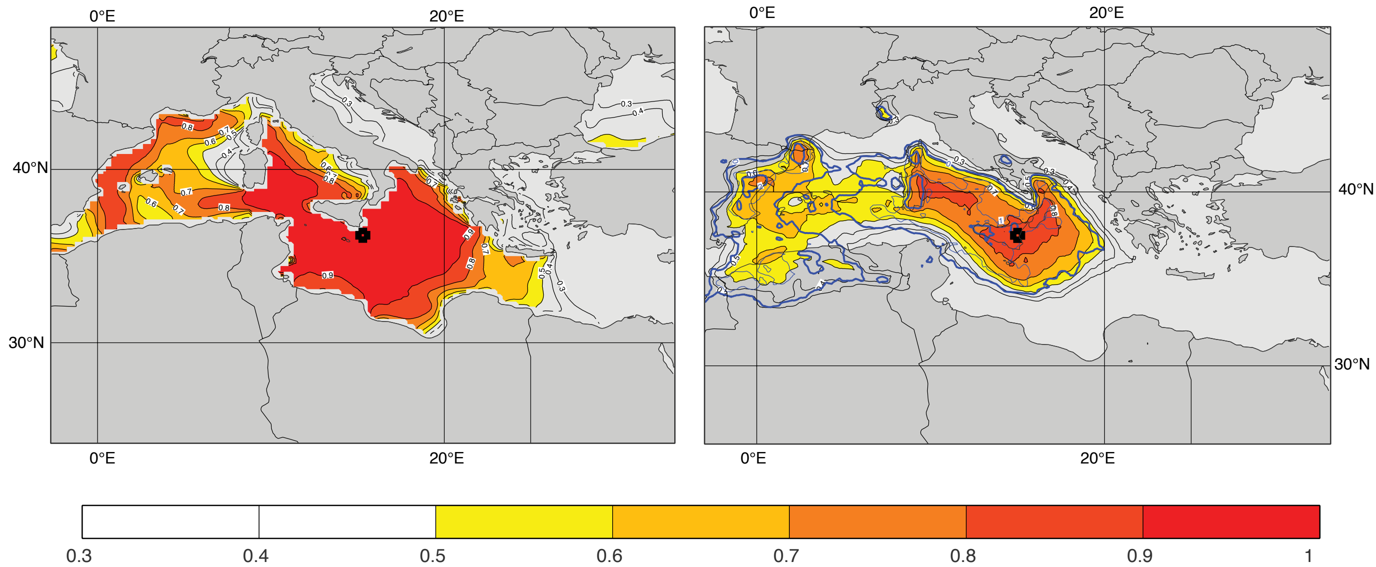

Figure 2 Extreme Forecast Index for 24-hour maximum significant wave height valid on 20 January from the IFS ensemble, initialised 20 January (left) and 14 January (right).

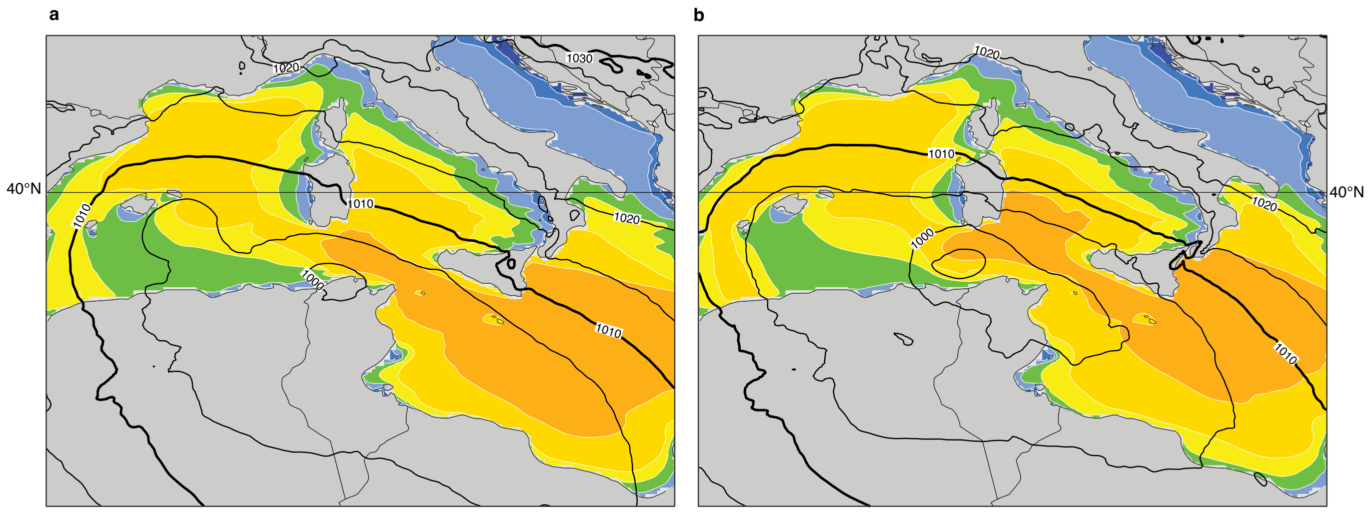

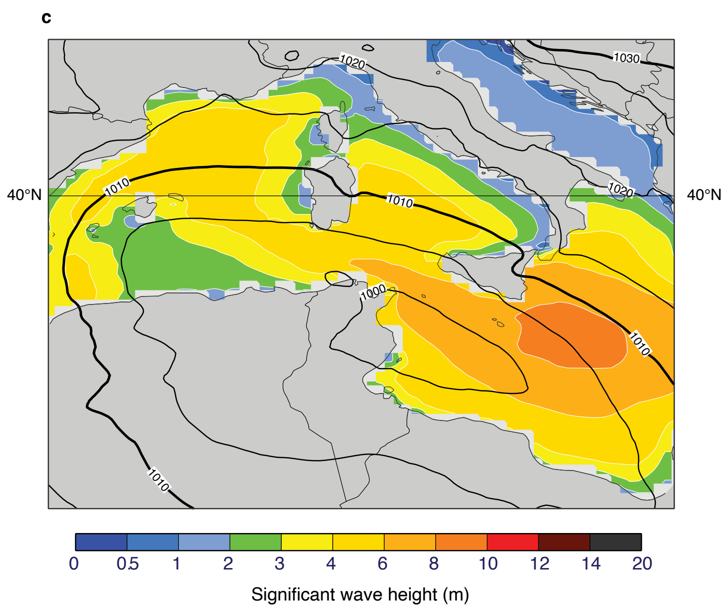

and slow movement of the low-pressure system. While not covered here, the predictability of extreme precipitation varied more, with relatively high predictability over southern Italy, but lower predictability, even in the short range, over northern Tunisia.

Figure 3 Significant wave height (shading) and mean sea level pressure (contours) valid on 20 January 12 UTC in (a). ECMWF analysis, (b).IFS control forecast and (c). AIFS Single v2 forecast from 19 January 00 UTC.

Turning historical operational forecast data into impact

Lauren Rootham, Alba Gomez Segura, Xiaobo Yang

Users of ECMWF data are addressing some of the most pressing challenges in the meteorological and climatological communities.

From improving early warning systems to strengthening the resilience of energy infrastructure, ECMWF products underpin innovation and deliver real-world impact.

Here we present three recent examples that illustrate how ECMWF products are being used in practice and how access to consistent, high-quality datasets delivers impact.

Managing renewable energy risk

At CentraleSupélec, researchers developed an innovative approach to managing financial risk for solar energy producers facing increasingly unpredictable weather. Their method combined a weather-based insurance model with a peer-to-peer risk-sharing system, which demonstrated that solar farms can dramatically reduce financial uncertainty – cutting risk variability by more than half.

The work relied on access to ECMWF products – specifically irradiance and temperature forecasts from the ECMWF operational archive – which was fundamental in resolving the lack of publicly produced data, allowing researchers to compute forecast errors and translate them into monetary values. The datasets provided a high temporal resolution, enabling a realistic simulation of a solar farm producer's operational workflow. The researchers modelled decision making at 12 p.m. daily, synchronised with electricity spot market timings, to recompute actual electricity production.

Without access to ECMWF data, such empirical validation would have been impossible, leaving the research paper without the necessary practical grounding to demonstrate the algorithm's utility.

A contrastive learning model for predicting tropical cyclone rapid intensification

Researchers at the Institute of Oceanography, Chinese Academy of Sciences, developed a novel contrastive learning model that significantly improves the forecasting of tropical cyclone rapid intensification utilising ECMWF forecast and analysis data. The ability to predict rapid intensification – when a tropical cyclone strengthens dramatically over a short period more accurately – has critical implications for emergency management and coastal community preparedness, potentially saving lives and reducing economic losses.

By training advanced machine learning algorithms with ECMWF datasets, the team was able to better capture the complex atmospheric conditions leading to sudden storm strengthening.

Such research was made successful through ECMWF datasets being broadly consistent both spatially and temporally, as well as the breadth of variables. This enabled the efficient alignment of environmental fields with tropical cyclone samples, reducing the additional effort required in relation to missing data or differences in formatting, providing well-aligned input features and improving comparability across samples.

In this work, ERA5 reanalysis data, produced by the Copernicus Climate Change Service (C3S) at ECMWF, provided comprehensive three-dimensional representations of atmospheric structure, critical for understanding vertical dynamics driving rapid intensification.

Whilst ERA5 reanalysis fields were used for model training, and as part of the environmental inputs, in the operational setting, ERA5 was replaced with forecast fields from the Integrated Forecasting System (IFS). The IFS offers higher spatial resolutions and stronger representation of wind-field intensity, leading to an added improvement in the forecasting of rapid intensification performance. This combination of reanalysis and forecast data reflects the growing role of ECMWF products as both training data and real-time inputs for data-driven applications.

Without ECMWF data, the work would have required substantial harmonisation of variables, spatial grids, and time axes, as well as the reengineering of the entire data processing and model training pipeline. This would likely have led to reduced sample availability, loss of consistency across datasets, and lower model performance and reproducibility. In short, ECMWF data were not just helpful, they were instrumental to the success and reliability of the research.

Understanding wildfire impacts on the stratosphere

Researchers at the U.S. Naval Research Laboratory used ECMWF's operational archive alongside ERA5 to investigate how wildfire smoke influences the stratosphere (https://doi.org/10.1029/2023JD040289). Their work focused on two major Australian wildfire events in 2009 and 2019–2020, which generated intense pyrocumulonimbus clouds capable of forming anticyclonic circulation known as Smoke With Induced Rotation and Lofting (SWIRLs).

Access to ECMWF's operational archive proved invaluable to the team's work. It enabled the analysis of forecasts without aerosol heating, allowing them to test whether the SWIRLs could be sustained in the absence of smoke-driven radiative effects.

The results showed that ECMWF's operational model could not maintain

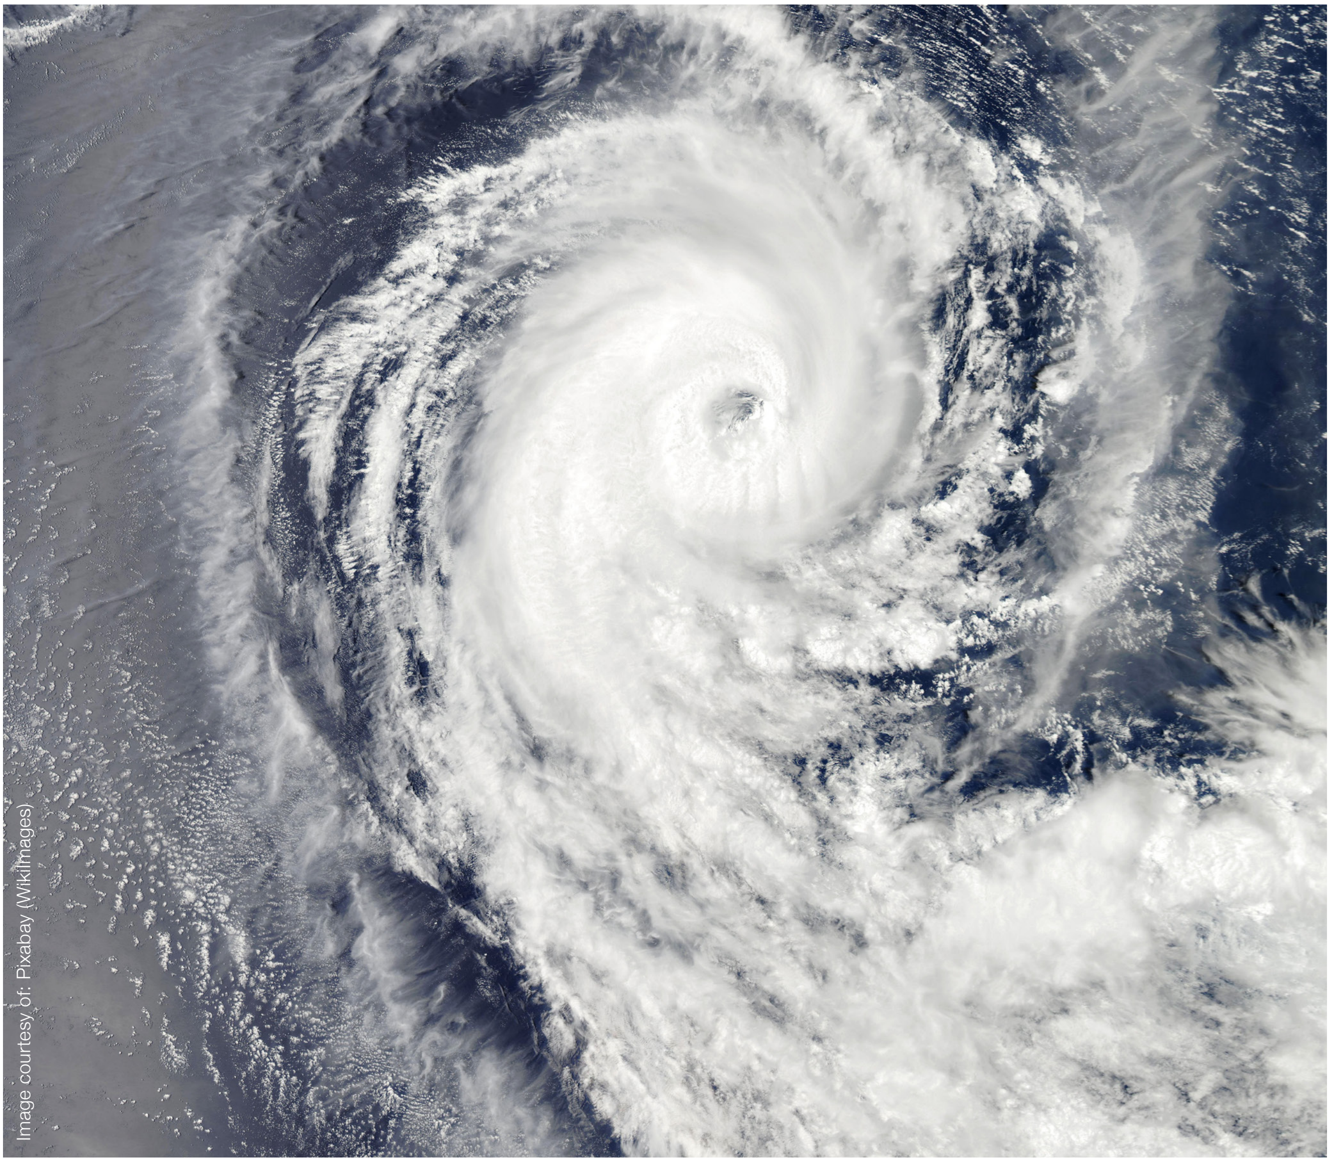

the observed SWIRLs when aerosol heating was removed. This finding provided clear evidence that accurately representing aerosol heating in forecast model physics is essential for capturing smoke-induced stratospheric circulations.

The ability to download historical forecast data, without needing to run high-cost forecast models, made this analysis possible and helped accelerate the research.

This study underscores how ECMWF's high-quality, readily accessible data enables cutting-edge atmospheric research, directly supporting new insights into extreme wildfire impacts and advancing our understanding of how smoke influences the upper atmosphere.

Supporting global users

These examples, drawn from Research Service Agreements for access to the Archive Catalogue and from across multiple sectors, demonstrate how ECMWF products support a wide and evolving range of user needs.

As the demand for highquality data grows, ECMWF remains committed to providing reliable datasets, expertise, tools and support that help users turn complex challenges into measurable impact.

Access ECMWF data

Explore available datasets and access options via the ECMWF data portal: https://www.ecmwf.int/en/forecasts/datasets

Did you know? ECMWF Fee Waiver Scheme

If you are a research institution, humanitarian agency, NGO, or organisation working for societal benefit, you may be eligible to have ECMWF service charges **waived**. Waivers are subject to management approval and permitted only for Official Duty or Research Purposes. To find out more about this: ecmwf.int/en/forecasts/accessing-forecasts/service-agreements.

Enhanced land–atmosphere data assimilation

Christoph Herbert, Peter Weston, Patricia de Rosnay, Hans Hersbach

The coupling methodology in current state-of-the-art numerical weather prediction systems is primarily based on weakly coupled data assimilation. This involves running separate assimilation approaches for each Earth system component, with coupling occurring only through cycling of the coupled forecast model. As a result, observations from one component cannot influence others until the next data assimilation window. We developed an outer loop coupling approach for increased consistency and more balanced initial conditions between land and atmosphere.

In line with the ECMWF Earth system strategy, coupled data assimilation is central to providing consistent initial conditions across the atmosphere, land, and ocean – allowing observations to simultaneously influence the analysis of multiple components (de Rosnay et al., 2022). Figure 1 illustrates the ECMWF reanalysis coupled data assimilation configurations used in ERA5 (Herbsach et al., 2020) and plans for future reanalysis generations.

The Copernicus Climate Change Service Evolution (CERISE) project aims to advance coupled surface–atmosphere assimilation in preparation for the next generations of seasonal prediction and reanalysis systems beyond ERA6. CERISE builds upon quasi-strongly coupled assimilation, as described by Penny et al., 2017. In this method, the system is not treated as a single fully integrated state vector, as it is in a strongly coupled data assimilation. However, cross-component state exchange between the different assimilation systems occurs within the same data assimilation window. This is implemented through outer loop coupling within the incremental formulation of 4D-Var used for atmospheric analysis (see Box A). ECMWF's Integrated Forecasting System (IFS) Cycle 50r1 will already feature outer loop coupled assimilation between the ocean/sea-ice and the atmosphere (see article by Browne et al., 2026 in this Newsletter). For land, ECMWF has been developing a unified Land Data Assimilation System (LDAS) built upon

a 4D-Var outer loops and the LDAS in detail

Incremental 4D-Var relies on outer loop cycling. In each iteration of what is referred to as an outer loop, a nonlinear model trajectory is produced. We linearise the nonlinear model trajectory around the most recent state estimate, resulting in a quadratic cost. This is then minimised by an iterative solver in what are called inner loops. The minimisation yields an increment being the update to the forecast state using observations based on their respective error statistics. The increment is valid at the start of the assimilation window, and we use it to update the nonlinear state as the improved guess for the next outer loop iteration. In the final outer loop, we propagate the increment obtained from the minimisation associated with the penultimate trajectory by the tangent-linear model from the start of the assimilation window to the required times to form the analysis. The analysis is then used to initialise the short forecast to provide the initial condition for the subsequent assimilation window.

The LDAS includes screen-level analysis, snow analysis (including single-layer snow temperature), as well as soil moisture and soil temperature analyses at multiple soil depth layers (0–7 cm, 7–28 cm, and 28–100 cm). The screen-level analysis is an intermediate step required to produce gridded analysis fields from SYNOP relative humidity (RH2m) and temperature (T2m) station observations at screen level (2 metres). These observations are typically reported at standard synoptic hours (e.g. 00, 06, 12, 18 UTC) and rely on a 2-dimensional Optimal Interpolation (2D-OI). The analysed RH2m and T2m fields are assimilated as pseudo-observations, i.e. not measured directly by instruments but treated as observations. We use an SEKF to analyse soil moisture, soil temperature and snow temperature. The sensitivities between the observations and these prognostic variables are derived from their spread in the 50 members of the Ensemble of Data Assimilations currently produced. For the soil moisture analysis, in addition to using the screen-level observations, we assimilate satellite soil moisture retrievals from the Soil Moisture and Ocean Salinity (SMOS) and the Advanced SCATterometer (ASCAT) missions.

doi: 10.21957/gxrf-vv36

TABLE 1 Configurations of numerical experiments run at horizontal resolution TCo399 (~28 km) covering boreal summer (JJA) 2022 and winter (DJF) 2022/23 periods. The ERA-CERISE prototypes use outer loop coupling and are compared to the weakly coupled approach implemented in current operations (IFS Cycle 49r1).

	Description	LDAS in outer loop(s)	Period
1	Weakly coupled (Control)	1 (no land feedback)	1 Jun to 31 Aug 2022
2	ERA-CERISE prototype summer	1, 2, 3	1 Jun to 31 Aug 2022
3	Weakly coupled (Control)	1 (no land feedback)	1 Dec 2022 to 28 Feb 2023
4	ERA-CERISE prototype winter	1, 2, 3	1 Dec 2022 to 28 Feb 2023

the Simplified Extended Kalman Filter (SEKF). This system already includes a multi-layer soil moisture analysis in operations and has been extended to other variables, making it suitable for enhanced coupling (Herbert et al., 2024).

As part of CERISE, we developed capabilities in the IFS to enable outer loop land–atmosphere coupling using a similar infrastructure approach as for the ocean. In this setup, we run the LDAS in multiple outer loops within each assimilation window, initialising each Earth system component at every outer loop with updated land and atmospheric fields. We investigated different configurations to identify the optimal degree of land–atmosphere coupling in terms of number of coupled loops.

The coupling configuration with the LDAS in the first three outer loops allows for the most balanced initial conditions across the land and the atmosphere, showing an overall positive impact on near-surface atmospheric forecasts and improved fit to independent atmospheric observations. However, enhanced coupling can also lead to degraded skin temperature at regional scale when evaluated against land surface temperature data, mainly due to deficiencies in the representation of the diurnal cycle.

Outer loop coupling developments

We conducted major infrastructure developments to support outer loop land–atmosphere coupled data assimilation. In the current weakly coupled system, the SEKF runs its own nonlinear trajectory to provide the necessary forecast fields used for quality control, screening, first-guess departure calculations and as the background for the SEKF analysis. The trajectory represents most of the SEKF computational cost, and effectively re-runs the first nonlinear trajectory in the atmospheric 4D-Var. As a cost-effective alternative, we developed an externalised SEKF configuration which does not run a trajectory. In this configuration, the required trajectory fields are written out to disk in the first 4D-Var trajectory and are simply read back in the SEKF without the need to run additional trajectories. In combination with several additional optimisations, this

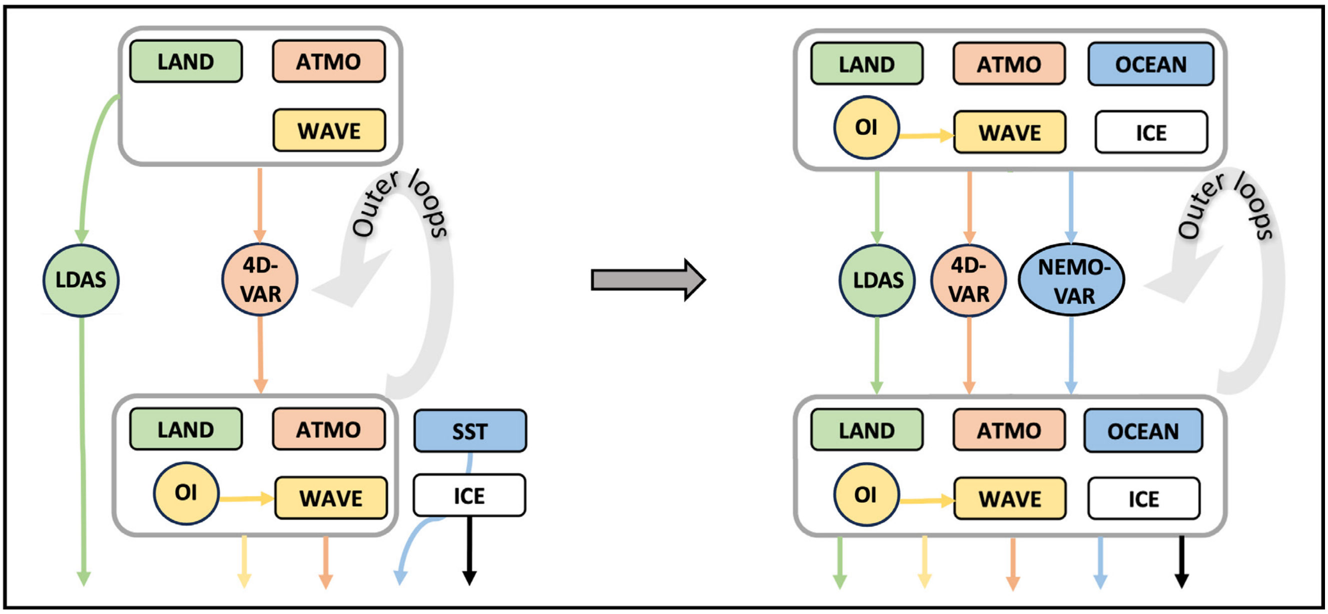

FIGURE 1 Coupled data assimilation schematics showing the ERA5 system (based on IFS Cycle 41r2) on the left compared to the next generation of coupled reanalysis systems beyond ERA6 on the right. Rectangular boxes represent Earth system component models, while circular shapes denote the assimilation approaches used in the respective analyses. Note that the OI wave analysis is performed only once within the penultimate outer loop.

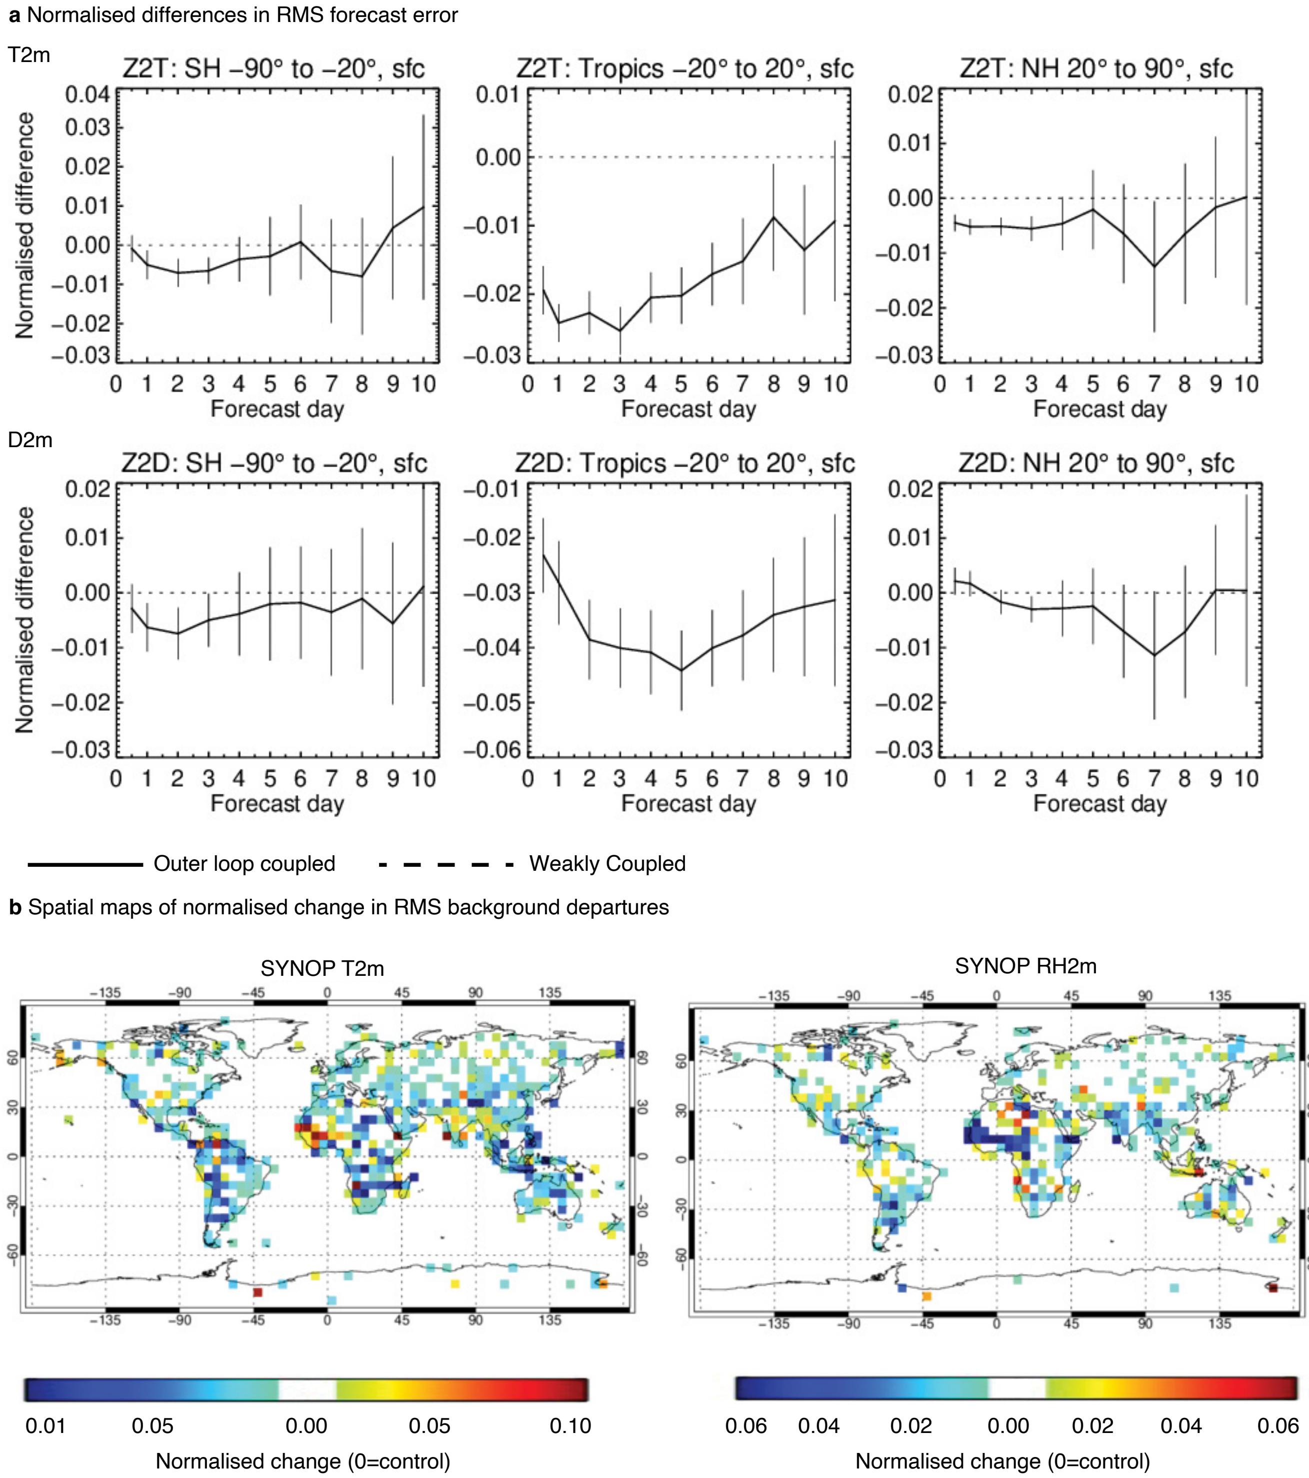

FIGURE 2 Differences in forecast skill for near-surface variables comparing the outer loop coupled system with the weakly coupled system over the boreal winter period 2022/23. (a) Zonal averages of normalised differences in root mean square (RMS) forecast errors for 2-metre temperature (T2m) and 2-metre dewpoint temperature (D2m) against operational analyses (IFS Cycle 47r3) as a function of forecast day; (b) Spatial maps of normalised change in RMS background departures (12-hour forecast in observation space) against T2m and 2-metre relative humidity (RH2m) SYNOP station observations.

reduces the SEKF computational cost by a factor of 80 at horizontal resolution Tco1279 (~9 km), while maintaining equivalent meteorological performance.

In the IFS, we are integrating the LDAS into the 4D-Var workflow as a separate family. We follow a structure similar to that of ocean–atmosphere coupled assimilation, where we run the surface tasks in parallel to the atmospheric minimisation. The LDAS can be flexibly activated up to the penultimate outer loop, incorporating both the 2D-OI and the SEKF in each loop, while the snow analysis runs only in the first activated instance.

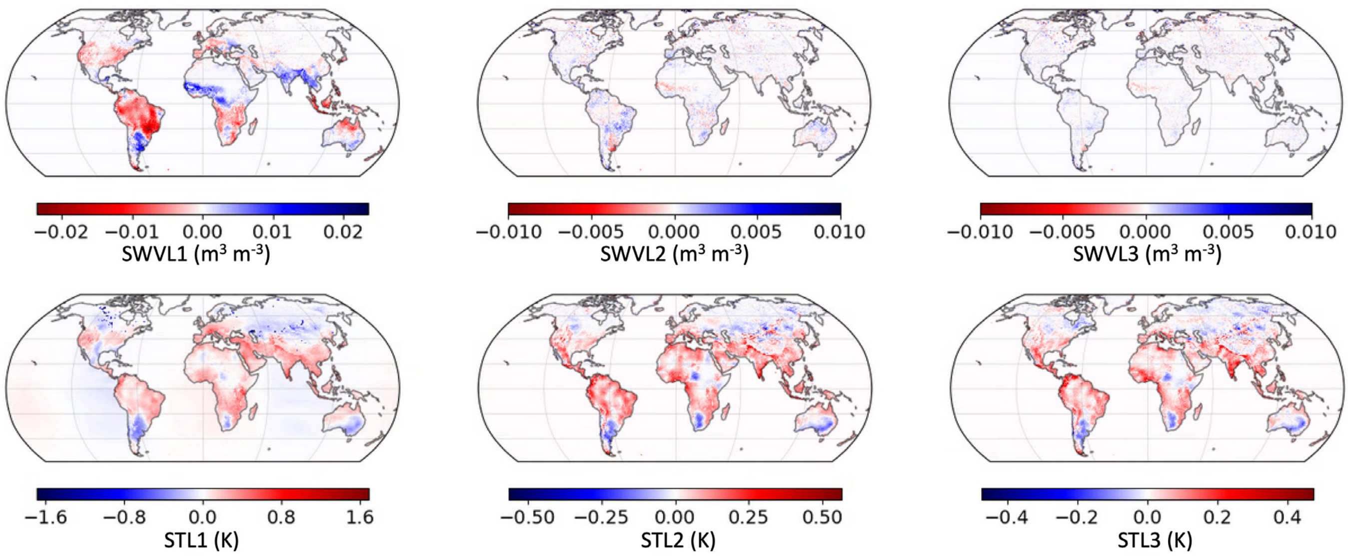

FIGURE 3 Temporal mean of multi-layer soil moisture (SWVL1–3) and soil temperature (STL1–3) analysis increments over the boreal winter period 2022/23, based on the outer loop coupled land–atmosphere data assimilation system.

We can independently initialise each outer loop trajectory and LDAS from one of the recent land-surface states obtained in previous loops. The SEKF produces a land-surface analysis for all the prognostic variables in the top three soil layers, and the respective analysis increments are added at the start of the assimilation window to produce updated land-surface fields for the atmospheric and land analyses in the subsequent outer loop(s).

Coupling methodology

To enhance coupling between the LDAS and the atmospheric 4D-Var, we activate the LDAS within the first, and in several subsequent outer loops. In this framework, a 2D-OI and SEKF analysis can be produced within any outer loop, using the most updated state from the corresponding nonlinear trajectory. Performing the coupling at the outer loop level enables a more consistent exchange of information between components than achieved by weakly coupled configurations, particularly for variables that exhibit variability on short time scales. After running a land analysis, we update the land-surface state and feed it back to the atmosphere to initialise subsequent outer loops. Initial testing showed that more frequent updates between outer loops lead to improved forecast skill compared with simply reducing observation errors in the weakly coupled system.

In the outer loop coupled data assimilation system, the observation error covariance matrix (R) is scaled across successive outer loops, with progressively larger values prescribed in later iterations. The deliberate inflation of observation errors reduces the influence of the same observations when the land analysis is repeatedly updated.

This approach is necessary because the SEKF provides an exact minimisation of the land-surface cost function in a single application and has no knowledge that a given observation has already been assimilated in previous outer loops. Without scaling R, repeated applications of the SEKF would repeatedly “pull” the land state towards the same observations, potentially leading to excessive increments and drifts, particularly in subsurface variables under stronger coupling configurations. Increasing R across outer loops therefore limits the impact of reused observations while still allowing the land analysis to benefit from updated atmospheric conditions arising from nonlinear trajectory effects and the progressively refined 4D-Var atmospheric analysis.

Numerical experiments

Before assessing which coupling strategy performs best, we first examined the benefit of outer loop coupling relative to a re-tuning of the weakly coupled system. We made a direct comparison between short-term IFS experiments of two setups:

1. Running the LDAS once in the weakly coupled setup with reduced observation errors

2. Running the LDAS in several outer loops with updated conditions

In the first setup, the LDAS was activated only in the first outer loop, but with halved observation errors for the SEKF – an approach roughly equivalent to performing the analysis twice. In the second setup, the LDAS was activated in the first two outer loops using the same observations. This assessment revealed that halving the observation error in the weakly coupled system has a

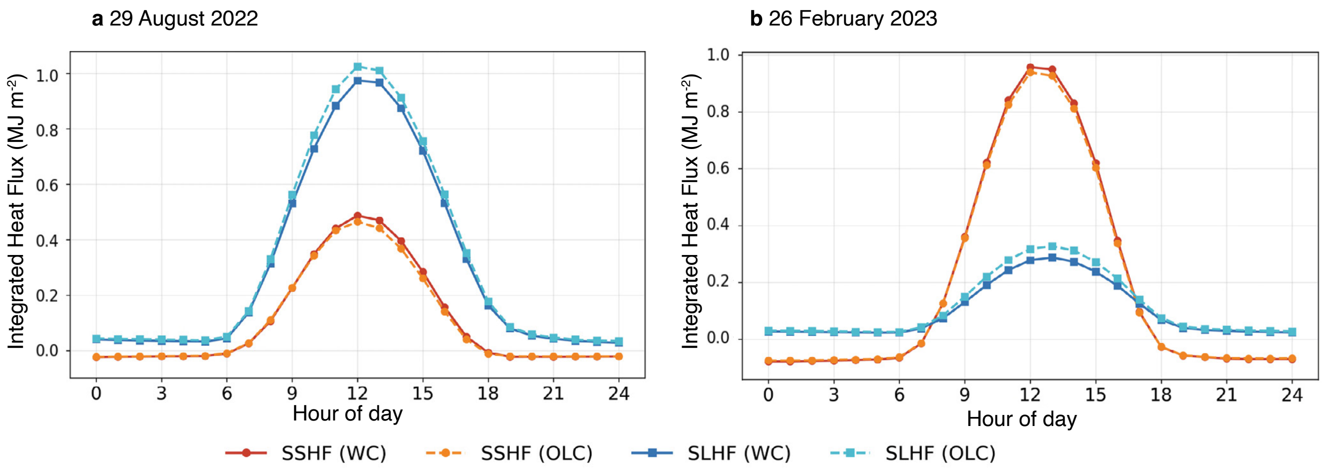

FIGURE 4 Hourly-integrated surface sensible and surface latent heat fluxes (as energy per hour). We illustrate the diurnal cycle of surface sensible heat flux (SSHF) and surface latent heat fluxes (SLHF) forecasts comparing the outer loop coupled (OLC, dashed lines) with the weakly coupled (WC, solid lines) as a function of the 'hour of day' on (a) 29 August 2022 and (b) 26 February 2023, averaged over the domain [5–20°N, 10°W–20°E] – also indicated as a dashed box in the spatial maps in Figure 5 in the left panels.

largely neutral impact on near-surface parameters, whereas the outer loop coupling configuration leads to a reduction in forecast errors.

To determine the optimal degree of coupling, we evaluated the impact of activating the LDAS in the first and additional later outer loops without specific tuning involved. The results demonstrated an overall positive impact on atmospheric forecast skill, particularly for T2m and lower-tropospheric temperature, when a higher degree of coupling was employed. Running the LDAS as part of the first three out of four outer loops produces the largest improvements in atmospheric forecast scores.

Improvements were especially pronounced in regions with sparse coverage of SYNOP T2m and RH2m stations. In these observation-sparse areas, atmospheric feedback plays a crucial role by providing more balanced conditions for the land analysis, which are subsequently passed back into the atmosphere. Hereby, atmospheric satellite observations can indirectly constrain the land variables in the absence of surface observations. This highlights the benefit of enhanced coupling. At the same time, some degradations were observed in skin temperature for configurations with increased coupling. These degradations were linked to a reduced amplitude of the diurnal cycle, indicating overly damped surface thermal variability. Sensitivity experiments showed that reducing the skin thermal conductivity over bare soil areas increases the amplitude of the diurnal cycle of skin temperature and can mitigate these degradations in some areas.

The experimental results guided the developments of outer loop coupled land–atmosphere data assimilation. As part of the CERISE project, we investigated the refined outer loop coupling system to produce coupled reanalysis prototypes as proof-of-concept in preparation for future generations of reanalysis systems beyond ERA6. The ERA-CERISE prototypes span boreal summer 2022 and winter 2022/23. Analysis fields are output hourly, consistent with standard reanalysis production – see summary in Table 1. In the selected setup, the LDAS runs within the first three outer loops, with progressively inflated observation errors applied to all observations assimilated into the SEKF.

Atmospheric forecast skill

We assessed the results of the experiments listed in Table 1 – comparing the performance of the ERA-CERISE summer and winter prototypes against the weakly coupled approach used in the current operational forecasting system. Compared to the weakly coupled reference, enhanced coupling generally reduces the forecast error for both T2m and 2-metre dewpoint temperature (D2m; Figure 2a), particularly in the tropics, as indicated by predominantly negative differences across nearly all forecast days. This improvement is statistically significant.

We find a spatially coherent positive impact on T2m, again with larger improvements in the tropics (Figure 2b). For RH2m, improvements are also widespread, with the strongest impact over the Gulf of Guinea region in southern West Africa. This example demonstrates that enhanced coupling promotes more active land–atmosphere interaction. Consequently, more moisture from the land enters the atmosphere, providing additional water vapor for cloud formation and precipitation. Both the land model and the LDAS are column-based, with propagation limited to the vertical. Feeding back land-surface conditions into the atmospheric 4D-Var, however, allows the information to be spread horizontally as well.

Impact on land variables

Figure 3 shows the soil analysis increments for winter 2022/2023 outer loop coupled data assimilation ERA-

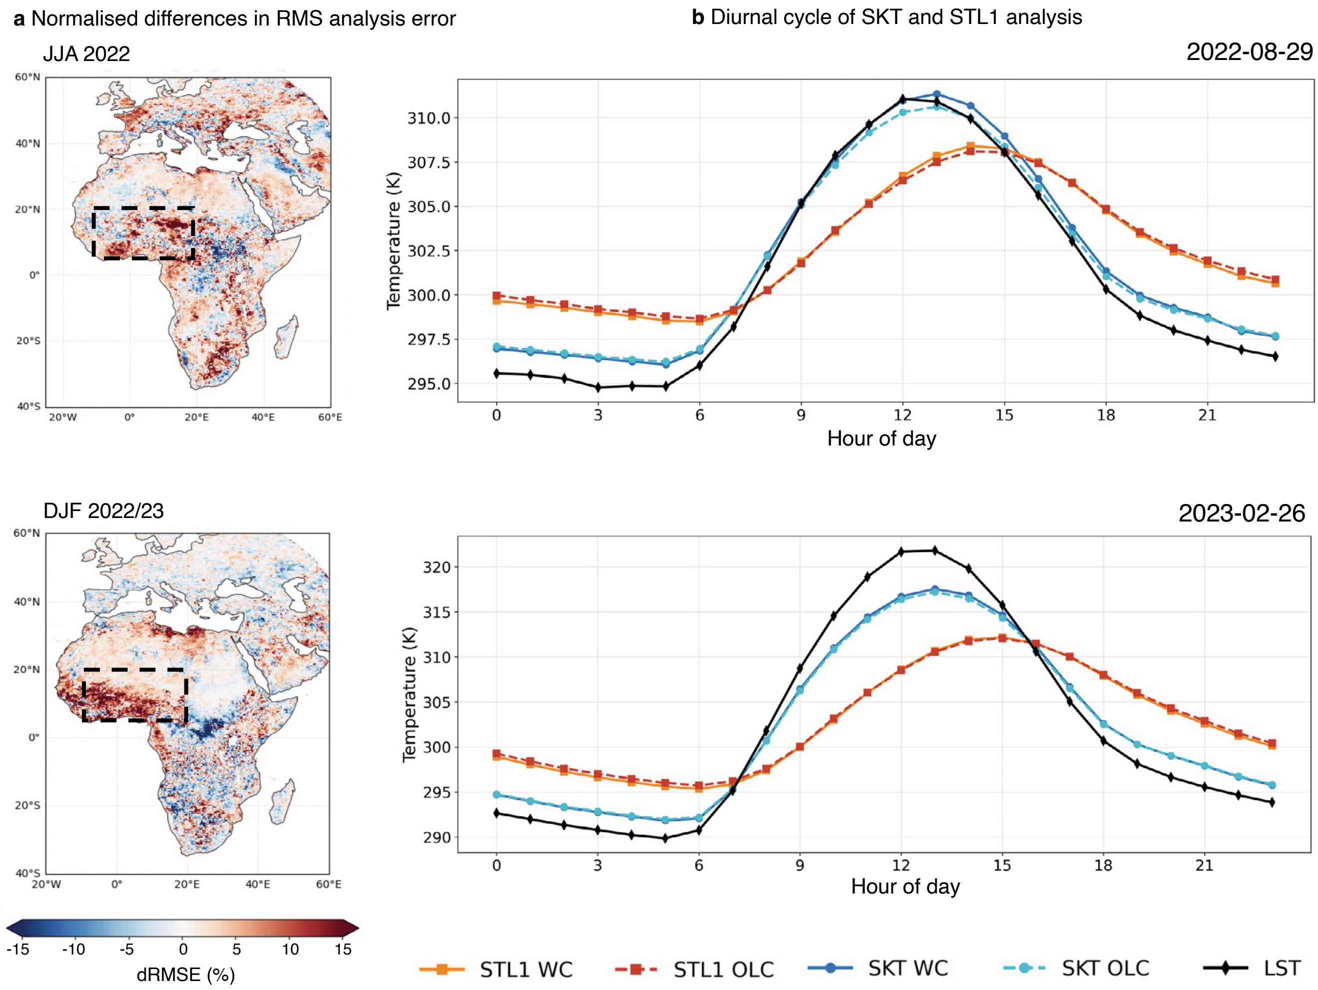

FIGURE 5 Verification of model skin temperature (SKT) analysis against SEVIRI-based land surface temperature (LST) retrievals for boreal summer (top) and winter (bottom), comparing the outer loop coupled (OLC) with the weakly coupled (WC) systems. (a) Spatial maps over Africa of three-month averages of normalised differences in RMS analysis error, where positive values denote error increase in the OLC system. (b) Diurnal cycle of SKT and first-layer soil temperature (STL1) analysis from the WC (solid lines) and the OLC (dashed lines) systems compared to LST observations (black curve) as a function of the 'hour of day' on 29 August 2022 (top) and 26 February 2023 (bottom). Values are averaged over the domain [5–20°N, 10°W–20°E] as indicated as the dashed box in the spatial maps.

CERISE prototype. The top panel shows that soil moisture increments are larger near the surface with smaller values in deeper layers, as expected. In regions of positive RH2m impact, such as West Africa in boreal winter (see right panel of Figure 2b), enhanced land–atmosphere coupling leads to larger positive increments in the top-soil layer. Increased soil moisture enhances evaporation, adding water vapor to the atmosphere – an effect relative to the weakly coupled system that can be observed across all forecast days – and thereby raises the dew point temperature. At nearly constant air temperatures, this higher dew point results in elevated relative humidity. Soil temperature shows larger increments at all soil layers throughout both periods (not shown). However, there are no observations assimilated with the capability to directly constrain subsurface variables.

Effect on the land–atmosphere interface

Enhanced coupling aims to improve the physical consistency of the model initial conditions at the surface–atmosphere interface. Ideally, updates to the land-surface state should translate into changes in surface fluxes and eventually elicit a downstream response in the atmospheric fields.

In the Gulf of Guinea domain, we assessed hourly-integrated forecasts of surface latent and sensible heat fluxes (Figure 4), averaged over the region, where we identified improvements relative to SYNOP RH2m. Sensible heat flux is typically negative at night, when the surface cools more rapidly than the overlying air. It is governed by the temperature difference between the surface and the air, increasing when the surface is warmer than the air and decreasing when the air is warmer than the surface. Latent heat flux depends on soil moisture availability and increases with wetter soils (or vegetation) due to stronger evapotranspiration. The land cover in this region is dominated by forest near the coast and savanna.

During the wet summer season, moisture-saturated soils lead to latent heat fluxes that clearly exceed sensible heat fluxes, whereas in the drier conditions in February, sensible heat flux is dominant.

Increased soil moisture from enhanced coupling means that the model allows for more evaporation and transpiration whenever water is available, which directly increases latent heat. With more energy being used for evaporation, less energy remains to heat the air, so sensible heat is reduced. Physically, this implies that enhanced coupling corrects for underestimated land–atmosphere moisture exchange in the model in both seasons, thereby correctly shifting the surface energy partitioning from sensible heat towards latent heat. This improves the representation of near-surface conditions, as seen in the RH2m forecast.

We validated the surface latent and sensible heat fluxes against in situ flux tower measurements. No overall significant impact was observed, particularly because surface flux measurements are inherently noisy and the observation network available is restricted to the European domain, where the SYNOP station network is already dense.

Skin temperature evaluation

Surface fluxes between the land and atmosphere are used in the IFS to compute the surface energy balance. The corresponding temperature at the land–atmosphere interface is represented by the skin temperature (SKT). We verify the SKT against land surface temperature (LST) retrievals from the Spinning Enhanced Visible and Infrared Imager (SEVIRI). In contrast to SKT, LST is measured by satellites and reflects the temperature of the radiating layer of the emitting surface, such as soil or vegetation canopy, with an effective depth that can vary depending on surface properties.

Figure 5 compares SKT with SEVIRI-based LST for boreal summer and winter, assessing both analysis error and diurnal cycle in the outer loop coupled and weakly coupled systems. The spatial distribution of normalised differences in RMS analysis error (Figure 5a) shows a mixed pattern in summer and more pronounced regional differences in winter, including over the region where relative humidity improvements are identified (see Figure 2b).

The diurnal cycle reveals a systematic underestimation (Figure 5b). Night temperatures are slightly overestimated (by about 2 K), while daytime peak temperatures are underestimated, particularly in winter (by up to about 5 K). Surface temperature reacts quickly to incoming radiation, leading to a larger diurnal cycle compared with temperatures in the lower atmosphere or in soil layers. Overall, these results indicate that the reduced diurnal amplitude is a model-related feature, and that outer loop coupling has only a minor effect in absolute terms.

Conclusion and outlook

As part of the CERISE project, we developed outer loop coupled land–atmosphere data assimilation as a major step toward enhanced consistency and physically balanced Earth system analyses. By repeatedly updating land initial conditions within the 4D-Var framework, we enable information to be exchanged between the atmosphere and land surface within the same data assimilation window, improving our representation of near-surface variables such as 2-metre temperature, dew point, and relative humidity.

Overall, we find that outer loop land–atmosphere coupling enhances forecast skill and consistency across components, providing a robust foundation for the operational forecasting and next-generation reanalysis systems toward ERA7 and beyond. An optimal degree of coupling that most benefits the atmospheric forecast has been identified. We are now focusing on operationalising the degree of coupling while applying safeguards – such as observation-error scaling and depth tapering of increments in the subsurface – to improve the physical realism of the land-surface state. Our enhanced coupling infrastructure also allows us to better exploit interface observations that are sensitive to both the land and atmospheric components, enabling them to simultaneously inform their respective analyses.

Funding acknowledgement

The CERISE project (grant agreement No 101082139) is funded by the European Union. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the Commission. Neither the European Union nor the granting authority can be held responsible for them.

Further reading

de Rosnay, P., **P. Browne**, **E. de Boisséson**, **D. Fairbairn**, **Y. Hirahara**, **K. Ochi et al.**, 2022: Coupled data assimilation at ECMWF: Current status, challenges and future developments. *Quarterly Journal of the Royal Meteorological Society*, **148**(747), 2672 –2702. *https://doi.org/10.1002/qj.4330*

Herbert, C., **P. de Rosnay**, **P. Weston** & **D. Fairbairn**, 2024: Towards unified land data assimilation at ECMWF: Soil and snow temperature analysis in the SEKF. *Quarterly Journal of the Royal Meteorological Society,* **150**(764), 4133–4155. *https://doi.org/10.1002/qj.4808*

Hersbach, H., **B. Bell**, **P. Berrisford**, **S. Hirahara**, **A. Horányi, J. Muñoz-Sabater et al.**, 2020: The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, **146**(730), 1999 –2049. *https://doi.org/10.1002/qj.3803*

Penny, S. G., **S. Akella**, **M. Buehner**, **M. Chevallier**, **F.Couillon**, **C. Draper et al.**, 2017: Coupled data assimilation for integrated earth system analysis and prediction: goals, challenges, and recommendations (No. GSFC-E-DAA-TN43810).

Building a data-driven Earth system model at ECMWF

Rachel Furner, Sara Hahner, Ewan Pinnington, Nina Raoult, Mario Santa-Cruz, Maria Luisa Taccari, Kenza Tazi, Lorenzo Zampieri

Predicting the weather requires more than just accurate modelling of atmospheric processes. Ocean temperature, currents and waves shape storm tracks, soil moisture influences heatwaves and regulates the land surface–atmosphere exchanges, sea ice modulates weather in polar regions and beyond, and river water levels are affected by rainfall, to name a few examples. Capturing these interactions is essential for realistic Earth system prediction, for both conventional and machine learning (ML) models.

After developing the Artificial Intelligence Forecasting System (AIFS) and bringing it into operations to produce daily 15-day weather forecasts, ECMWF is expanding the scope of data-driven modelling beyond the atmosphere to include key Earth system components.

In this article, we showcase how within the framework of the European Commission's Destination Earth (DestinE) initiative, individual ML models have been developed for several Earth system components. This marks a significant step towards an ML Earth system model that explicitly represents multiple components.

We highlight technical developments and scientific challenges for the following Earth system components:

- Sea ice, where we represent the thermodynamic evolution of this medium for an extensive set of variables (sea ice concentration, volume, albedo, and velocity) in response to atmospheric and oceanic forcings.
- Waves, where we learn ocean surface wave dynamics, enabling accurate prediction of potentially hazardous marine conditions.
- Oceans, where we model the evolution of the ocean throughout depth, including the response to atmospheric forcing. Global spatial fields for physical variables such as temperature, salinity, currents, and sea-surface height anomaly are predicted.

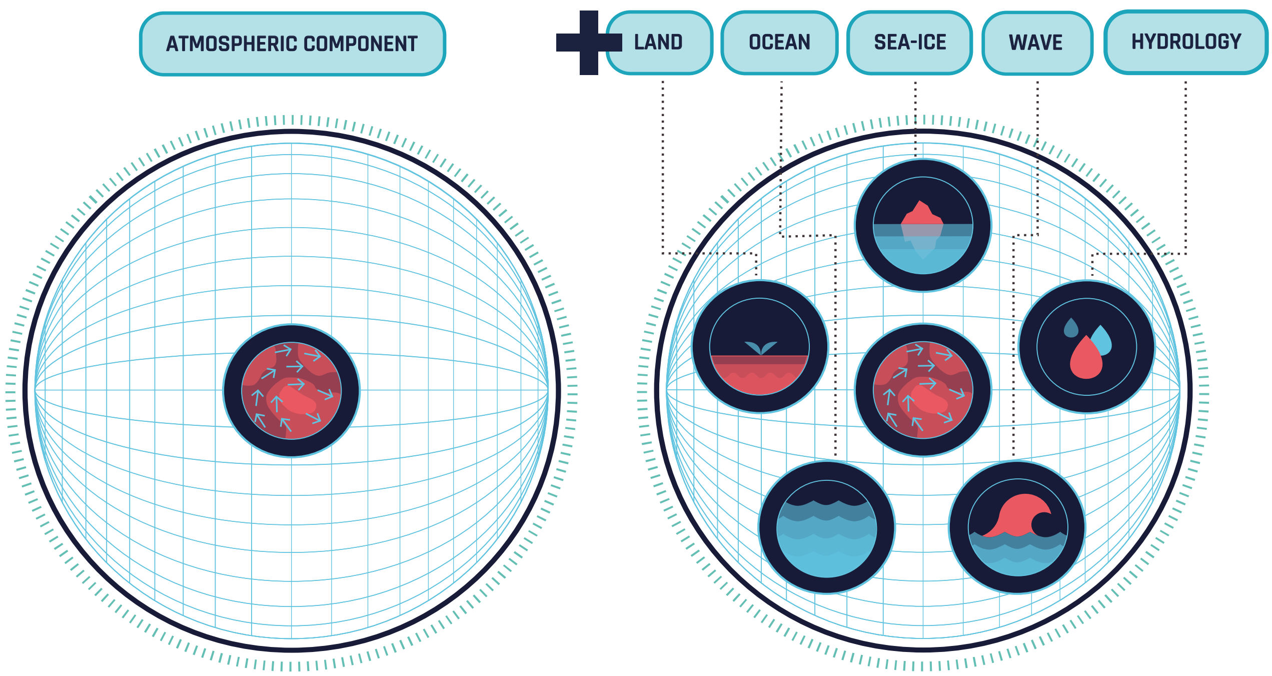

FIGURE 1 Building a data-driven Earth system model: new components for the land, ocean, sea ice, waves, and hydrology are being developed in DestinE to complement the already successful atmospheric component – the AIFS.

doi: 10.21957/c1dg-8k56

- Land, where we represent the land-surface response to atmospheric forcing, currently focusing on soil moisture and temperature, snow cover, skin temperature, and near-surface variables such as 2-metre temperature and dew point.
- Hydrology, where we model streamflow using meteorological forcings and catchment properties.

Each of these data-driven model components represents major advances in the respective domain. Together they build on decades of expertise in the physical modelling of the corresponding counterparts and lay the groundwork for a full Earth system model which captures these individual components and interactions between them.

Why more Earth system components matter

While the atmosphere is central to weather prediction, its dynamics are tightly coupled with processes in the ocean, land surface, sea ice, ocean waves and rivers (Figure 1). These interactions influence weather and climate across timescales ranging from hours to decades, and across spatial scales from small river catchments (2 km^2) to the world's oceans (107 km^2). Modelling these components not only increases the physical realism of weather predictions, but is also expected to improve accuracy, particularly for medium- to long-range predictions due to the long memory of land and ocean. Expanding the complexity of Earth system modelling has been a strategic priority within the AIFS ecosystem. By capturing more of the interactions that shape the planet's behaviour, we move closer to a fully coupled ML prediction model.

What are the challenges for Earth system components

Significant progress has been made in recent years in applying ML to atmospheric modelling for weather forecasting. Extending these advances to other Earth system components – such as the ocean or land surface – required overcoming several shared challenges. Observational coverage and data quality vary substantially across the different Earth system components, limiting the capacity of ML models to leverage observational, or observationally constrained, data. While the ERA5 reanalysis provides the backbone for atmospheric ML training, other components must rely on a mix of reanalyses, model output, and observations, which need to be carefully harmonised to ensure consistency. In addition, most ocean and land surface processes evolve more slowly and exhibit longer memory than the atmosphere. Training datasets spanning only a few decades might not be sufficient to fully capture their dynamics. Training strategies focused on learning a limited number of relatively short time steps might also not be suitable for learning these slow dynamics. Missing values are also common – for example, in time and space for hydrological observations, or points over land in the ocean, and wave fields – and therefore require explicit masking or gap filling. Furthermore, ML approaches and architectures that work well for the atmosphere may not be suitable for components with different spatial structures, temporal scales, or governing processes.

Beyond these shared challenges, each Earth system component has unique characteristics that require tailored approaches. Ocean surface waves are highly sensitive to local wind forcing, generating waves, as well as the propagation of swell generated by distant winds. Sea ice exhibits strong seasonal and interannual variability, particularly at ice edges. The 3D ocean operates on multiple temporal and spatial scales, with air–sea exchange impacting surface dynamics over short timescales, and deep ocean mixing influencing long-term climate. A particular challenge for land surface processes is the limited understanding of some of the physics – for example, the soil–plant–atmosphere interactions – which leads to uncertainties in the physical model, that is used as the basis of training data. Finally, many land surface observations are sparse and heterogeneous requiring model-based inference, for instance when estimating catchment-scale streamflow from limited data availability in hydrology. These differences highlight why the development of ML models beyond the atmosphere is non-trivial and why each Earth system component must be approached with domain-specific considerations.

Technical developments and innovations

Building ML components has required expanding the Digital Twin Engine with machine-learning pipelines built on the open-source Anemoi ecosystem, developed by ECMWF together with several national meteorological services across Europe. Anemoi provides software solutions for producing and handling AI-ready datasets, for end-to-end training, and for running weather and climate ML models.

Developing the Earth system components also required the creation of bespoke datasets for each component, illustrated in Figure 2. Land and waves components were built from synthetic model runs and specially produced hindcasts. The hydrology component, which focuses on streamflow, used observational datasets, while ocean components (sea ice, surface and 3D ocean) used datasets derived from ORAS6, the state-of-the-art ECMWF ocean reanalysis product.

Addressing many of the challenges discussed above relied on a range of technical innovations in the Anemoi

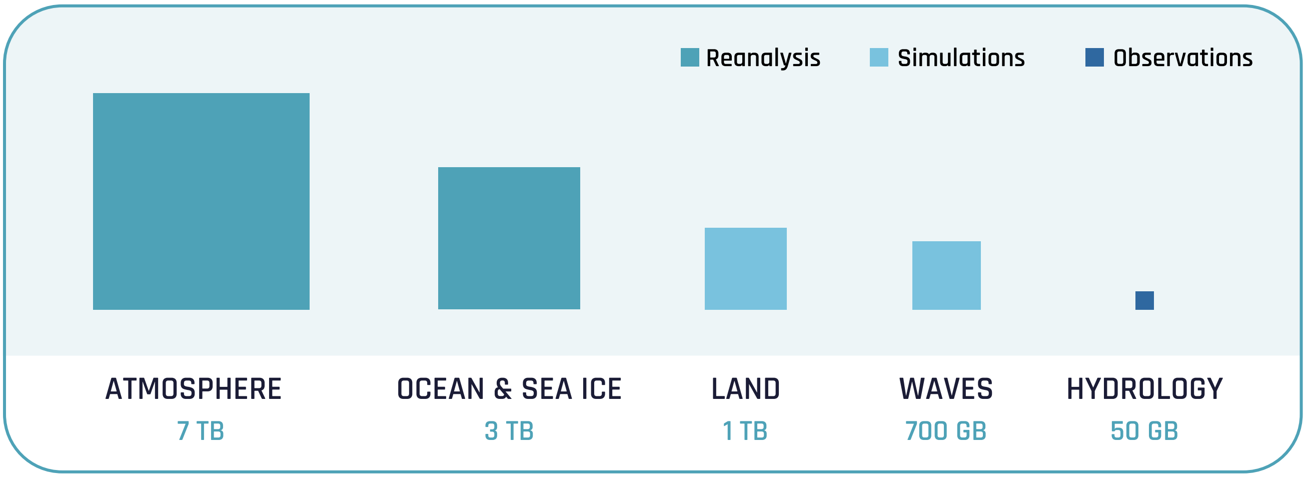

FIGURE 2 Size and provenance of the datasets used for training different Earth system components.

ecosystem. Originally developed for atmospheric ML modelling, the framework has since been adapted and extended to support the new Earth system components through a close collaboration between ML specialists, software engineers and physical modellers.

To support diverse data sources, *anemoi-datasets* was, for example, extended to support different data formats and to create additional variables where needed. This includes representing cyclic wave direction continuously using cosine and sine components, as well as handling not only time-varying fields, but also static variables (such as ocean bathymetry) and climatological fields (such as leaf area index for land-surface processes).

While *anemoi-datasets* is naturally optimised for gridded data, a different approach was required for hydrology, where streamflow is often modelled using catchment boundaries that do not fit neatly onto a regular grid. In this case, earthkit-hydro – a hydrology-focused toolkit within ECMWF's earthkit (*https://earthkit.ecmwf.int*) ecosystem – was used to efficiently perform catchment-based spatial aggregation and preprocessing. This approach enables the integration of real-time observations and provides hydrologically meaningful inputs for ML models.

The development of land and marine components means we now encounter missing data over the land and ocean, respectively. This required robust handling of missing values in *anemoi-core* and *anemoi-transforms,* since most ML models cannot natively process them. During training, missing values in the input fields are masked to prevent numerical instability or biased learning. The model is then allowed to generate predictions over the full domain, but values in masked regions are not considered in the training loss calculation. If a field is not part of the model input, its mask can be inferred from related variables – for example, snow cover, which is only part of the model output, uses the missing-value mask of snow depth, as illustrated in Figure 3.

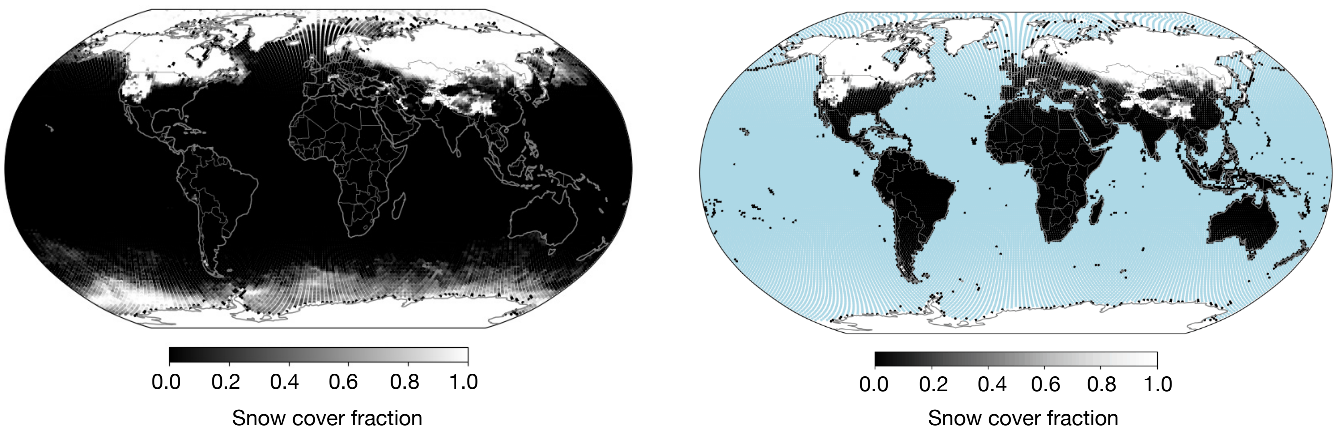

FIGURE 3 Illustration of how missing values are handled in Anemoi While the model is allowed to predict snow cover over the oceans (left), these regions are masked and filled with missing values in the output (right; blue indicates masked values).

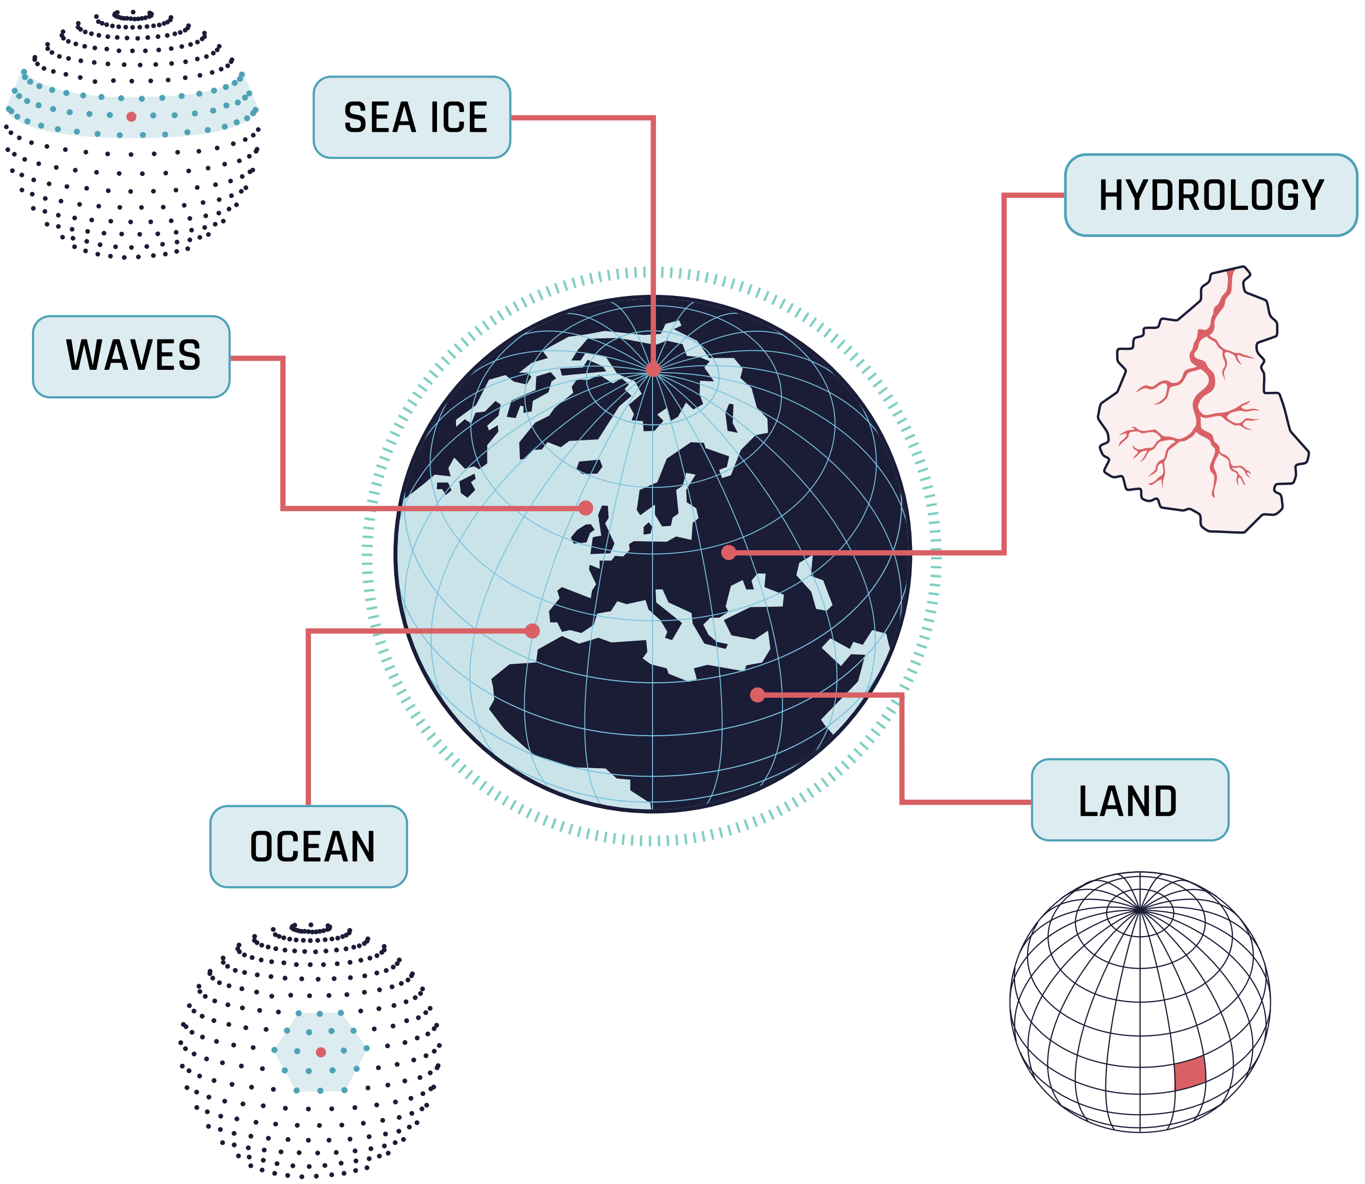

FIGURE 4 Diversity of model architectures across Earth system components. The figure illustrates how information can propagate in the different models, highlighting the locations/data which can be used when making predictions for a given point. During training, the model learns how much attention to give to each of the potential inputs, i.e. the attention mechanism.

Further innovations include the development of component-specific ML model architectures and exploration of attention mechanisms tailored to the characteristics of individual Earth system components (Figure 4). For example, land-surface processes are modelled using multilayer perceptrons (MLPs), reflecting their largely local and column-based nature, while hydrological processes along river networks are represented using long short-term memory (LSTM)-based models to capture directed flow and temporal memory. For ocean components, the core encoder–processor–decoder architecture of the AIFS is retained, with alternative attention mechanisms used to better represent localised ocean dynamics.

A key development to enable skilful Earth system component predictions has been the careful scaling of variables and weighting of loss terms to account for differences in magnitude and characteristic timescales. This helps to ensure that both fast processes (e.g. near-surface fluxes) and slow processes (e.g. soil moisture evolution and ocean heat uptake) are learned effectively. Component-specific normalisation is applied to reflect differences in temporal and spatial scales, including tendency-based scaling for the land and ocean components, and loss normalisation by catchment size in the hydrology model.

Ensuring physical consistency across components remains a central challenge and has been addressed through a combination of architectural choices and post-processing constraints. One example for this can be found in the sea-ice model, where sea-ice concentration constrains other sea-ice variables such as ice velocities, as in the physical model.

A key additional capability is the ability to run Earth system components as standalone ML models forced by atmospheric output from the AIFS. This modular approach enables targeted experimentation and development of individual components informed by physical domain expertise. It also represents a first step toward progressively integrating ML components into a fully coupled Earth system modelling framework.

Next steps for data-driven Earth system modelling

As part of DestinE Phase 2, successful prototypes were developed for various Earth system components (waves, ocean, land, sea ice, hydrology). These are further described in a number of dedicated DestinE blog posts and videos on the ECMWF DestinE website (*https://destine.ecmwf.int/ml-earth-system-components/*). The models were built by ECMWF leveraging in-house domain-specific knowledge and making use of EuroHPC supercomputing resources. These prototypes will continue to evolve, co-developed by domain and machine learning experts to ensure efficiency, scientific robustness, physical consistency, and relevance for operational Earth system prediction.

Looking ahead, two complementary pathways are being pursued for advancing data-driven Earth system modelling. Within the framework of DestinE, Earth system components will be coupled into a full Earth system model, similar to the way physics-based weather forecasting and climate models are coupled.

In parallel, the AIFS Single and Ensemble (ENS) configurations used operationally at ECMWF were enhanced by including fields representing the land surface and ocean waves in AIFS v2.0, and variables representing sea ice and surface ocean will be added to future releases.

Together, these efforts lay the foundations for comprehensive, fully data-driven Earth system prediction capabilities that combine atmospheric ML with specialised ML representations of the ocean and land, through targeted processes such as sea ice, ocean waves and currents, land-surface interactions, and streamflow all within the DestinE initiative.

Note: This work was co-authored with equal contributions by all contributors. Authors are listed in alphabetical order.

ECMWF publications

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

Technical Memoranda

936 **Pan, W., M. Bonavita, M. Chrust, & E. Hólm:** Data-Driven Emulation of Background-Error Variance in Variational Data Assimilation. *February 2026*

ESA or EUMETSAT Contract Reports

Jana Kolassa, J., P. Weston, K. Salonen & P. de Rosnay: Annual SMOS brightness temperature monitoring report - 2024/25. *May 2026*

Lean, K., N. Bormann, S. Pandítharatne: Evaluation of hyperspectral microwave observations for global NWP: Representation of expected uncertainties. *March 2026*

Jana Kolassa, J., P. Weston, K. Salonen & P. de Rosnay: Quarter 4 2025: Operations Service Report. *January 2026*

ECMWF Calendar 2026

Jun 1–4	Using ECMWF's Forecasts (UEF2026)
Jun 3	C3S User Workshop with the Hungarian User Community
Jun 3–5	9th C3S General Assembly
Jun 8–9	DestinE Co-Design Workshop
Jun 9	DestinE Networking Session
Jun 8–9	DestinE Co-Design Workshop
Jun 9–10	5th Destination Earth User eXchange
Jun 10	DMCSEE and C3S joint workshop with the Southeastern Europe user community
Jun 10–11	Copernicus pan-Arctic Regional Reanalysis (CARRA2) User Workshop
Jun 10–12	ICAP 2026 Meeting
June 15–18	Online Training course: Advanced HPC training
Jun 16–17	Council
Jun 16–18	Discover Anemoi: Webinar Series
Jun 18	Informal Seminar: Evaluating Extreme Precipitation Forecasts
Jun 18	AI Weather Quest: MAM Awards Webinar
Jun 23–24	CATRINE project General Assembly
Jun 29–Jul 2	22nd International Workshop on Greenhouse Gas Measurements from Space (IWGGMS-22) for better weather forecasts
Jun 29–Jul 3	2026 AR Recon workshop and 2nd Observational campaigns workshop
Jul 6–9	ECMWF Machine Learning Pilot Project All Hands 2026
Jul 15	Informal Seminar: NVIDIA Science
Sep 10–11	CaMa-Flood developer/user international meeting 2026
Sep 14–18	Annual Seminar 2026
Oct 5–7	Scientific Advisory Committee
Oct 5–8	Training course: Use and interpretation of ECMWF products
Oct 8–9	Technical Advisory Committee
Oct 12–16	Online training course: Introduction to ECMWF Data and Archive Services
Oct 19–23	Training course: Predictability and ensemble forecast systems
Oct 20–21	Finance Committee
Oct 21	Policy Advisory Committee
Oct 26–30	Joint CP-TT/DA-TT Workshop
Nov 3–5	2026 Workshop on Machine Learning for Land and Hydrology
Nov 4–6	Visualising Climate 2026
Nov 9–12	NWP-SAF workshop on exploiting satellite cloud and precipitation observations for weather forecasting
Nov 16–20	Training course: Data-driven Earth system forecasting
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. 187** | Spring 2026

European Centre for Medium-Range Weather Forecasts

www.ecmwf.int