ECMWF
European Centre for Medium-Range Weather Forecasts

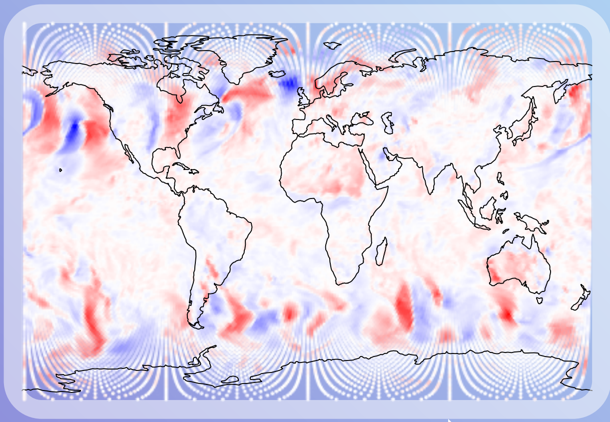

Newsletter

No. 181 | Autumn 2024

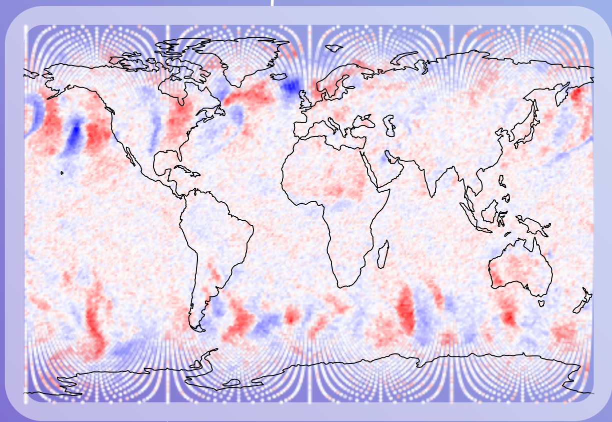

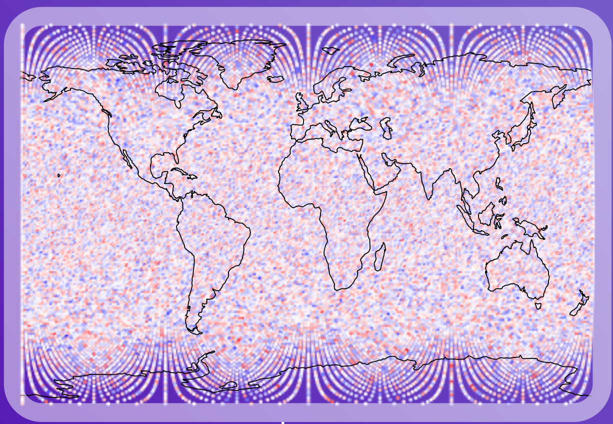

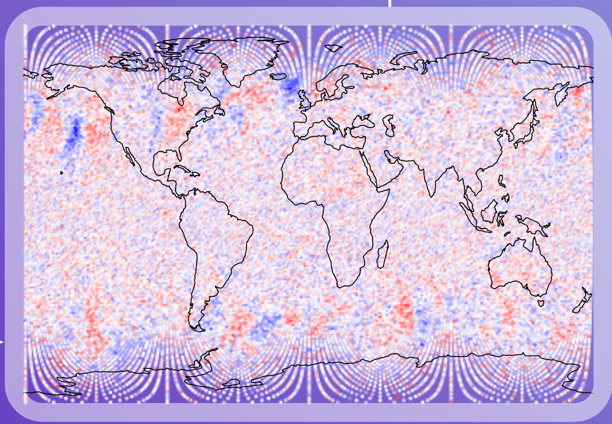

Publication policy

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

The ECMWF Newsletter is not peer-reviewed.

Any queries about the content or distribution of the ECMWF Newsletter should be sent to Georg.Lentze@ecmwf.int

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*

Machine learning ensembles

A big drawback of single, deterministic forecasts is that they tell us nothing about the certainty of the predicted outcome: is it quite certain, in other words is there quite a narrow band of possible outcomes, or is it quite uncertain, in other words is there a rather broad band of such outcomes? This is why ensemble forecasts were introduced: a whole range of forecasts are produced with slightly different initial conditions and approximations in the model. The result enables us to draw conclusions on the probability that particular outcomes will materialise. Ensemble forecasts have been with us since 1992. They were initially introduced at a coarser resolution than a single 'high-resolution forecast'. However, since June 2023 the grid spacing of all our global medium-range forecasts, produced operationally by the Integrated Forecasting System (IFS), has been 9 km. This change in emphasis in favour of ensemble forecasts has now also been applied to the experimental machine learning forecasts we are producing with the Artificial Intelligence Forecasting System (AIFS). As described in this Newsletter, at this stage two methods have been developed to produce AIFS ensemble forecasts. They have both been found to be similarly skilful, and we have made forecasts from one of them available as open charts under ECMWF's open data policy.

The task now will be to decrease the horizontal grid spacing of these forecasts, which is still rather coarse at about 111 km. This compares with a grid spacing of currently 28 km for deterministic AIFS forecasts, which is set to go down further. The number of AIFS ensemble members is currently 51, the same as for our medium-range IFS forecasts. However, it could be higher in the future because AIFS forecasts can be produced using considerably less computing power than traditional, physics-based forecasts. Meanwhile, ECMWF has teamed up with Member States in an initiative to create machine learning weather forecasting systems, called Anemoi.

This Newsletter also provides an overview of the changes to be introduced in the IFS next month, when the operational forecasting system is upgraded to Cycle 49r1. This upgrade particularly improves 2 m temperature and 10 m wind speed forecasts. One reason for these improvements is the assimilation of 2 m temperature observations, but upgrades to the data assimilation methodology and improvements to the land surface model also play a role. IFS Cycle 49r1 includes a new and fundamentally different scheme for model uncertainty in the operational ensemble forecasts. This development, which has been in the making for years, is presented in a separate article. There are also updates from the two EU Copernicus services we run. The Copernicus Climate Change Service (C3S) has developed a tool to explore climate change, and the Copernicus Atmosphere Monitoring Service (CAMS) has monitored high levels of wildfire emissions in Canada. Progress has also been made in ECMWF's contribution to the EU's Destination Earth initiative, which is detailed in an article that presents evaluation results. Destination Earth also uses machine learning to provide ensemble capabilities at high resolution, which shows that this method has wide applicability in what we do.

Florence Rabier
Director-General

Contents

Editor Georg Lentze • **Typesetting & Graphics** Anabel Bowen • **Cover** See figure on page 33

doi: 10.21957/de720h3x9p

Extremely warm summer in southeast Europe

Linus Magnusson, Rebecca Emerton

Summer 2024 was the warmest on record for Europe as a whole, albeit with large variability across the continent, according to ECMWF's ERA5 reanalysis from January 1940 to the present. While parts of western Europe saw near-average or below-average temperatures, the summer was characterised by much warmer-than-average temperatures across most of eastern Europe. As an example, Cuprija in Serbia had 65 consecutive days with a maximum temperature above 30°C. In this article, we focus on a region covering a large part of southeastern Europe (39–46°N, 15–30°E), much of which saw its warmest summer in the ERA5 record.

The sub-seasonal to seasonal scale

Here we focus on predictions at the sub-seasonal scale up to six weeks ahead, and at the seasonal scale four months ahead. Comparing composites of 2-metre temperature weekly anomalies for 6 June to 25 August from ERA5 and ECMWF sub-seasonal forecasts with different lead times, the pattern over Europe was well captured in week-2 forecasts (see the first image). In week-6 forecasts the signal was weaker, as expected, but it captured the pattern, with southeastern Europe having a stronger anomaly than western Europe. However, the extension of the warm anomaly to northeastern Europe was not well captured.

The seasonal forecast from 1 May valid for June–July–August also captured a signal that southeastern Europe would have a stronger warm anomaly than western Europe. However, in the seasonal forecast, western Europe was predicted to be warmer than average, too, which was not the case in the outcome. The seasonal forecast also had a strong cold anomaly in the northern Gulf of Bothnia in the Baltic Sea, which was a remnant of the cold winter in northern Europe, but this water body became warmer than average.

Turning to the week-to-week variability for land-points in the box over southeastern Europe, we see significant variability during the summer in the ERA5 reanalysis, but all weekly means had a warm anomaly (see the second image). The largest anomalies are found for the 2nd and 3rd week of July. The week-to-week variability was very well captured in week-1 forecasts. Week-2 forecasts captured the variability but with a somewhat lower amplitude in the ensemble mean. The week-3 forecasts were more problematic, especially for the onset of the hottest period in the middle of July. For the 2nd week of July, the ensemble mean from 26 June predicted a slight cold anomaly for a large part of southeastern Europe. The outcome was also warmer than the 95th percentile of the ensemble

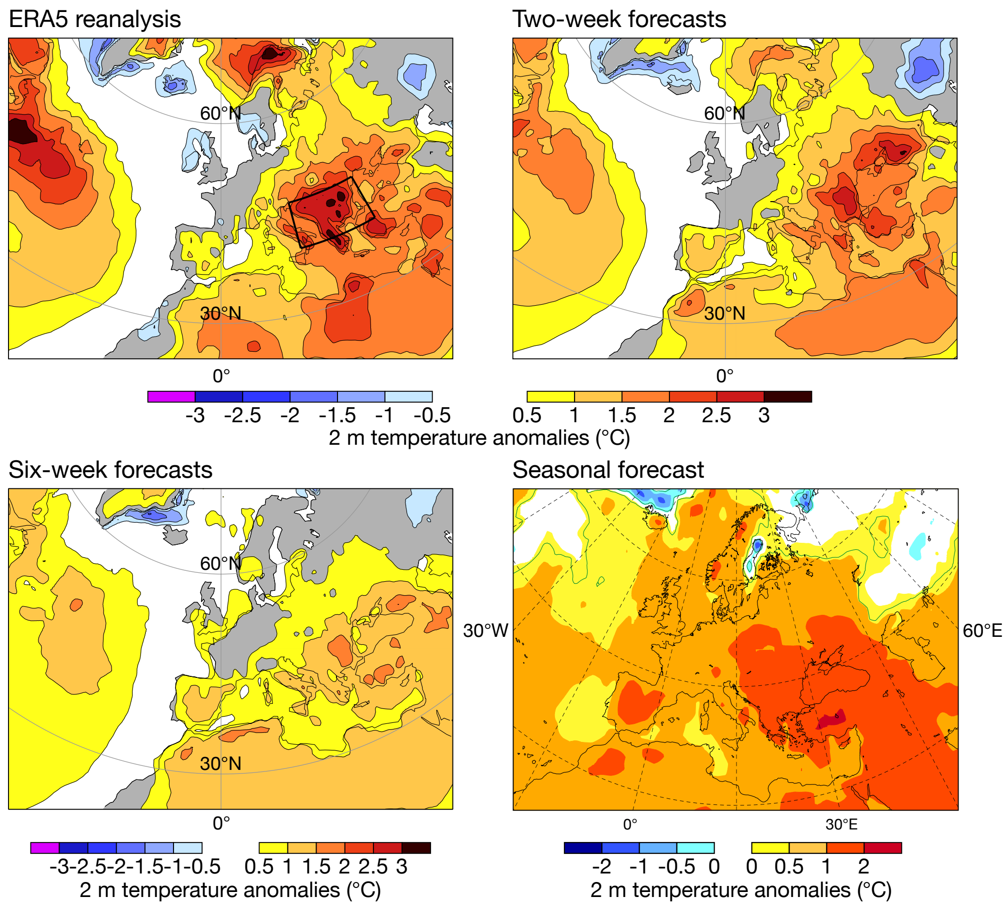

Temperature anomalies in reanalysis and forecasts. Two-metre temperature anomalies for 6 June – 25 August from ERA5 (top left), composite of 2-week forecasts (top right) and 6-week forecasts (bottom left) from sub-seasonal forecasts, and seasonal forecast from 1 May for June–July–August (bottom right).

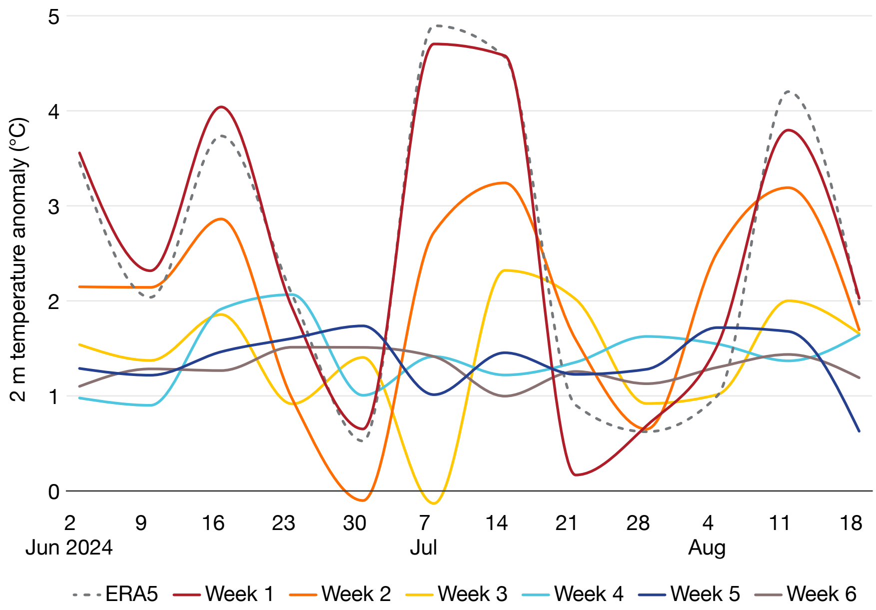

Weekly temperature anomalies in southeastern Europe. Weekly mean 2-metre temperature anomalies for land areas in southeastern Europe, in the box highlighted in the first plot of the previous figure, from ERA5 (dashed) and ensemble means from sub-seasonal forecasts with different lead times.

(not shown). For week-4 to week-6 forecasts, we find very little intra-seasonal variability, but all forecasts predicted warm anomalies of 1–2°C.

Heat stress

Associated with the warm anomalies, southeastern Europe also experienced well-above-average heat stress based on the Universal Thermal Climate Index (UTCI), which considers temperature, humidity, wind speed, sunshine and heat emitted by the surroundings, and how the human body responds to different thermal environments. It indicates weather-induced outdoor thermal stress in humans by classifying UTCI values into ten different heat and cold stress categories, with units of °C representing a 'feels-like temperature'. An ERA5 dataset of UTCI is available on the website of the EU-funded Copernicus Climate Change Service (C3S) implemented by ECMWF (https://doi.org/10.24381/cds.553b7518).

Based on UTCI values calculated from ERA5, the UTCI anomaly in southeastern Europe for the summer as a whole was 3.3°C above average, relative to the 1991–2020 reference period. The entire period from 28 May to 31 August saw daily maximum feels-like temperatures averaged over southeastern Europe exceeding the threshold for 'moderate heat stress' (see the third figure). Sixty-six days during summer (June–July–August) exceeded the 'strong heat stress' threshold (32°C), which can be compared to an average of 29 days per summer in the period 1991–2020. The 'very strong heat stress' threshold (38°C) was reached on 17 July (38°C) and 13 August (38.3°C). At the peak of the heat stress on 13 August, the feels-like temperature was up to 10°C higher than the 2 m temperature in some locations, particularly in the east of the region and in southern Italy (not shown). This indicates the importance of considering other environmental factors beyond temperature for assessing the impacts of heatwaves. Operational medium-range forecasts of UTCI, alongside other thermal comfort indices such as heat index, wind chill factor and wet bulb globe temperature, will become available in Cycle 49r1 of the Integrated Forecasting System (IFS), which is to be made operational this November.

Further information on surface air temperature anomalies in Europe and the world in August 2024 can be found in the corresponding C3S climate bulletin: https://climate.copernicus.eu/surface-air-temperature-august-2024

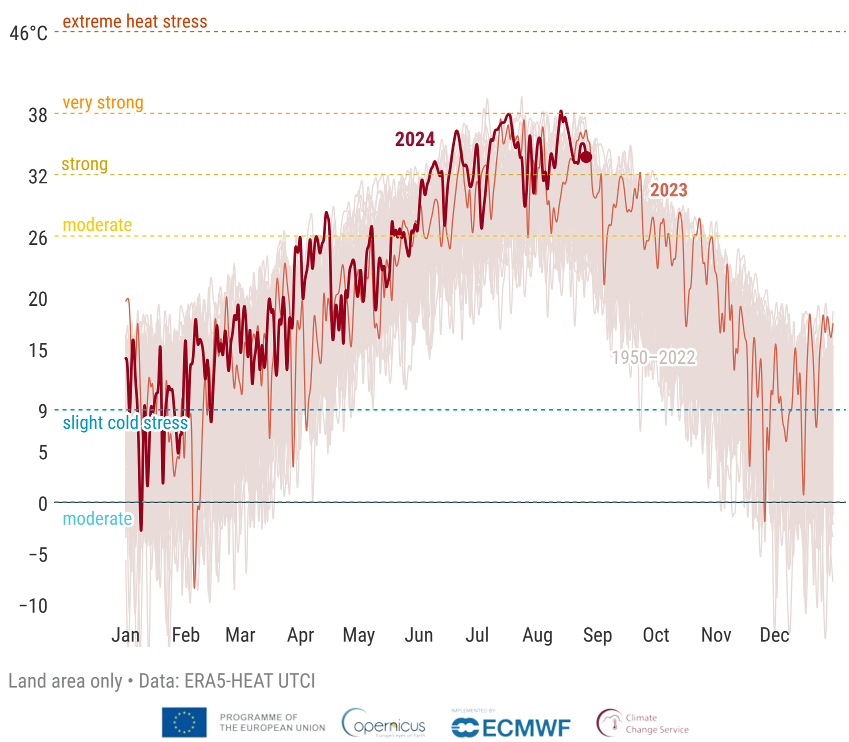

Feels-like temperature over southeastern Europe. Average daily maximum feels-like temperature (°C) in land areas over southeastern Europe, in the box highlighted in the first figure, indicating heat and cold stress based on the ERA5-HEAT Universal Thermal Climate Index.

State of wildfires 2023–24

Francesca Di Giuseppe, Joe McNorton, Mark Parrington, Anna Lombardi

The State of Wildfires report (https://essd.copernicus.org/articles/16/3601/2024/), the first edition of a new systematic annual review of wildfires across the globe, was published in the journal *Earth System Science Data (ESSD)* in August. The publication takes stock of extreme wildfires of the 2023–2024 season, explains their causes, and assesses whether events could have been predicted. It also evaluates how the risk of similar events will change in future under different climate change scenarios.

The report, which will be published annually, is co-led by authors from ECMWF, the UK Met Office, the University of East Anglia (UK), and the UK Centre for Ecology & Hydrology, with the support of more than 40 fire scientists from five continents.

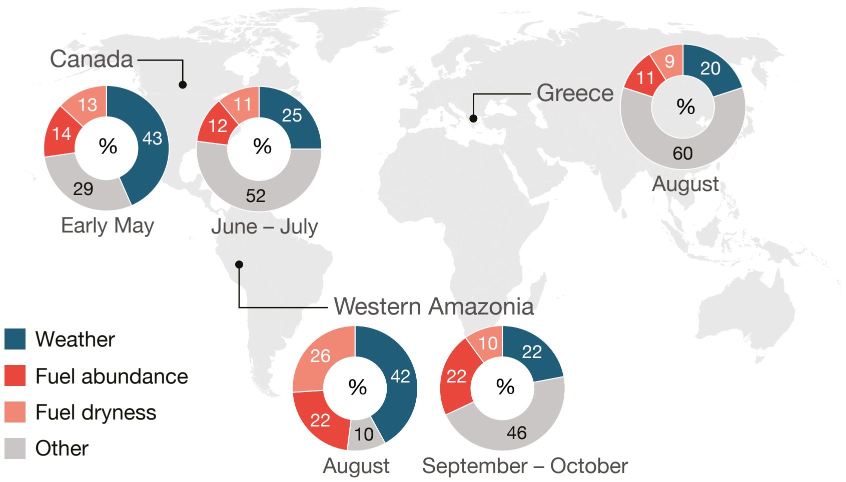

Cause attribution. This is a summary of the causes of wildfires in three regions in 2023. It highlights the shifts in the drivers of fire activity during the extended fire seasons in Canada and Western Amazonia, where the influence of natural factors (weather and fuel) decreases as the fire season progresses. The anomalous event in Greece also stands out, with 60% of the causes remaining unexplained, likely linked to human influence.

Main outcomes of this year's report

Globally, during the March 2023 to February 2024 fire season, wildfires released 2.4 Gt of carbon (C), or 8.8 Gt of carbon dioxide (CO_2), into the atmosphere – an amount equivalent to the combined annual anthropogenic emissions of the USA, the EU, and South America.

Notable fire events included a record-breaking fire season in Canada, the largest recorded wildfire in the EU (Greece), El Niño-driven drought fires in western Amazonia and northern South America, and deadly, fast-moving fires in Hawaii and Chile. All major wildfires combined claimed over 250 lives and led to the evacuation of at least 250,000 people worldwide.

The report highlights how human-driven climate changes have made these extreme fire events more likely. An attribution analysis led by the UK Met Office shows how human influence has made fire weather conditions witnessed during parts of the 2023–2024 fire season three times more likely in Canada, 20 times more likely in western Amazonia, and twice as likely in Greece.

The report also explores the likely scenarios if human emissions continue at their current pace or decrease thanks to reduction efforts. Even assuming emission reduction initiatives are in place, a Canadian born today is more likely than not to experience another fire season of similar magnitude to last year's within their lifetime. In comparison, someone born in the 1940s would have had a one-in-ten chance of witnessing such an event.

ECMWF's contribution

ECMWF contributed to diagnosing the drivers of wildfires, assessing their predictability, and evaluating emissions related to these events.

The Centre's fire weather forecast identified which events could be predicted and at what lead time, proving forecasting skills at the seasonal timescale.

The Probability of Fire (PoF) model used by ECMWF provided valuable insights into the drivers of key wildfire events (see the figure). The model concluded that early-season wildfires in Canada and Western Amazonia were primarily driven by extreme fire weather conditions. However, persistent wildfires later in the season were unlikely to have been caused by anomalous weather or fuel conditions alone. This suggests that several factors unaccounted for by the model may have contributed to the prolonged fire season. These factors may include limited fire management resources following the unprecedented early season activity, continued fire spread not captured by the model, land management practices, or the re-emergence of dormant early-season fires.

Data from the EU's Copernicus Atmosphere Monitoring Service (CAMS), which is implemented by ECMWF, was also used extensively in the report. The reported fire emission totals were based on the mean value between CAMS Global Fire Assimilation System (GFAS) and Global Fire Emissions Database (GFED) datasets. Additional context for the impact of fire emissions on air quality was provided with the CAMS reanalysis of global atmospheric composition, showing that Canada had the highest annual average surface concentrations of particulate matter less than 2.5 micrometres in diameter (PM2.5) in 2023. Meanwhile, many regions of North America experienced concentrations above the US Environmental Protection Agency's exposure threshold of 35 µg/m^3 for more than 10 days, increasing to more than 120 days close to the where the fires were burning in western Canada.

Media impact

Officially published on Wednesday, 14 August 2024, the report attracted significant media attention, being featured in over 600 items of international coverage. It has been promptly accepted as a scientific reference source and as a basis for several Wikipedia page updates documenting the extreme wildfires of the 2023–2024 season. It now ranks as ESSD's 6th most impactful publication, according to Altmetric. The Altmetric score even surpassed those of other major reports, like the Global Carbon Budget and the 'Indicators of Global Climate Change', a keystone publication by Intergovernmental Panel on Climate Change (IPCC) Working Group I authors. These achievements are even more impressive given this was the inaugural edition of the report, and we are still in the process of establishing its reputation.

Monitoring the 2024 Canada wildfires in CAMS

Mark Parrington, Enza Di Tomaso

During the summer of 2024, a large number of wildfires burned across Canada, primarily the western part of the country. According to estimated biomass burning emissions from the Copernicus Atmosphere Monitoring Service (CAMS), implemented by ECMWF for the EU, they resulted in the second-highest annual total emissions since the start of systematic monitoring in 2003. They were only surpassed by the historic Canadian wildfires which burned from May to September 2023 (see the first figure).

Location of fires

The first significant fires of 2024 occurred in May in western Canada. The British Colombia Wildfire Service attributed them to the surface reignition of holdover fires which had been smouldering underground through the winter. Most of the fires were relatively localised to the northeastern corner of British Columbia but resulted in the evacuation of thousands of people. The fires continued in this area through June. The fire count increased significantly through July as large wildfires developed across a wide region of boreal North America, covering Alaska and Canada's western territories and provinces.
The Northwest Territories, as in 2023, contributed the highest number of fires of any region of Canada through August. Further east during August, large fires also developed across Saskatchewan, Manitoba and Ontario.

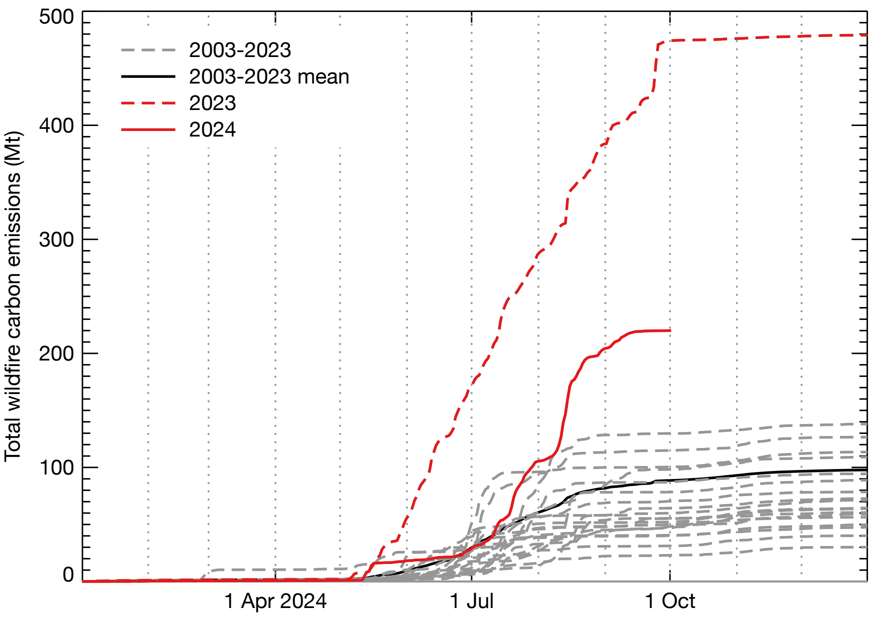

Canada wildfire carbon emissions. GFASv1.2 daily total cumulative wildfire carbon emissions for Canada in 2024 (red solid line) compared to 2023 (red dashed line) and the data for 2003–2022 (grey dashed lines, mean in black solid line).

Monitoring of emissions

As part of CAMS, the Global Fire Assimilation System (GFAS) provides near-real-time monitoring of daily global wildfire locations and emissions. GFAS currently merges fire radiative power observations from the two Moderate-resolution Imaging Spectroradiometer (MODIS) instruments on the NASA Terra and Aqua Earth Observation System satellites. Emissions of a variety of smoke constituents, including carbon monoxide (CO), carbon dioxide (CO_2), methane (CH_4), black carbon and organic carbon aerosols, nitrogen oxides (NO_x), and non-methane hydrocarbons (NMHCs), are estimated in GFAS. They are subsequently used as lower boundary conditions for the operational 5-day forecasts and analyses of global and regional atmospheric composition produced by CAMS. The GFASv1.2 dataset provides a daily time series that starts in 2003 at a spatial resolution of 0.1 degrees, allowing us to compare current emissions with those in previous available years.

Routine analysis of the GFAS data and CAMS forecasts hinted throughout the summer that the 2024 fire emissions were going to be at an extreme level for the second year in a row. Overall, the total emissions for Canada were established in early August as the second-highest annual total of the past 22 years (see the first figure). However, at the territory/province level different patterns emerged. Twenty-two-year records in monthly total emissions were set for British Columbia (May), Alberta (July), the Northwest Territories (August), Saskatchewan (August), Manitoba (May and August) and Ontario (August). The highest August monthly total emissions were also recorded for Canada.

Effects on air quality and temperature

The scale of the 2024 wildfires resulted in severely degraded air quality across western Canada, which also spread as far as eastern Canada and the USA. One notable atmospheric impact, resulting from the increased fire emissions in the Northwest Territories during August, was an episode of long-range smoke transport across the North Atlantic (see the second figure). This was detectable as enhancements to surface PM2.5 concentrations across parts of western Europe. Analysis of the impacts on numerical weather prediction of this smoke transport, comparing the interactive aerosols in the CAMS operational system with a corresponding control run based on the CAMS IFS configuration, but using climatological aerosols, indicated a reduction of up to 2°C in the daily mean 2 m temperature from France to Scandinavia.

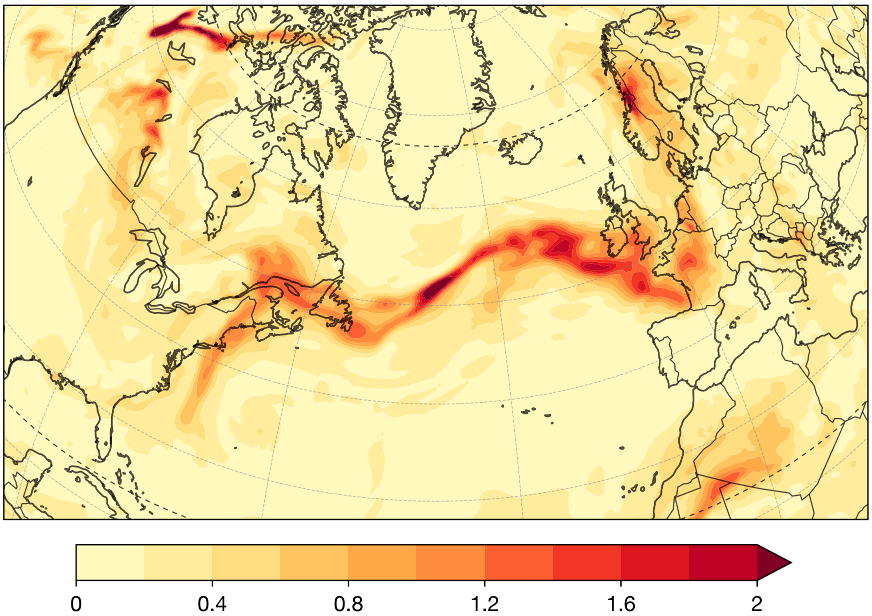

Aerosol optical depth analysis. Total aerosol optical depth analysis at 550 nm valid for 18 August 2024 at 12 UTC, showing a smoke plume crossing the North Atlantic.

Why we monitor wildfire emissions

Canada fires in 2024 have clearly been record-setting and impacting different scales, from local air quality degradation in downwind regions to long-range effects at both the continental and inter-continental scales. CAMS near-real-time monitoring of wildfire emissions around the world and the impacts on atmospheric composition are essential for evaluating the CAMS forecast performance. It also enables applications of CAMS open data on significant episodes which can lead to degraded air quality in near real-time.

Introducing Anemoi: a new collaborative framework for ML weather forecasting

Jesper Dramsch, Baudouin Raoult, Matthew Chantry (all ECMWF), Teresa García (AEMET, Spain), Leif Denby (DMI, Denmark), Florian Prill (DWD, Germany), Niko Sokka (FMI, Finland), Antonio Vocino (ITAF Met Service, Italy), Jasper Wijnands (KNMI, the Netherlands), Thomas Nipen (MET Norway), Carlos Osuna (MeteoSwiss), Sara Akodad (Météo-France), Michiel Van Ginderachter, Dieter Van den Bleeken (both RMI, Belgium)

A range of national meteorological services across Europe and ECMWF are pleased to announce the launch of Anemoi, a Python-based framework for creating machine learning (ML) weather forecasting systems. Named after the Greek gods of the winds, Anemoi is a collaborative, open-source initiative involving the Spanish State Meteorological Agency (AEMET), the Danish Meteorological Institute (DMI), the German National Meteorological Service (DWD), the Finnish Meteorological Institute (FMI), the Italian Air Force Meteorological Service (ITAF Met Service), the Royal Netherlands Meteorological Institute (KNMI), MET Norway, Météo-France, MeteoSwiss, Belgium's Royal Meteorological Institute (RMI) and ECMWF. It has the potential to democratise access to and further develop data-driven weather forecasts.

The goal of Anemoi is to provide the key building blocks to train state-of-the-art data-driven models and run them in an operational context. As a framework it seeks to handle many of the complexities that meteorological organisations will share, allowing them to easily train models from existing recipes but with their own data.

Components

Anemoi comprises an ecosystem of Python packages, which cover the full life cycle of data-driven modelling. These Python packages each address a crucial aspect of the artificial intelligence (AI) weather forecasting pipeline:

- **Anemoi Datasets:** This component generates ML-optimised datasets from various sources and data formats of meteorological data and observations, complete with rich metadata. For example, an optimised subset of ECMWF's ERA5 reanalysis or of a historical operational analysis dataset. It streamlines the often cumbersome process of data preparation, ensuring that high-quality, consistent, and optimised data is available for model training. (https://anemoi-datasets.readthedocs.io)

- **Anemoi Training:** Understanding that flexibility is key in ML model development, this module allows users to modify most parts of the training through configuration files, without needing to alter the underlying code. This approach democratises access, enabling meteorologists without deep coding expertise to experiment with advanced data-driven weather prediction models. (https://anemoi-training.readthedocs.io)
- **Anemoi Models:** This package houses the model code, designed with efficiency and minimal dependencies in mind. It ensures that the transition from development to deployment is as smooth as possible. (https://anemoi-models.readthedocs.io)
- **Anemoi Inference:** Building on ECMWF's experience with the AI Models tool, which has been running daily experimental ML forecasts for over a year, Anemoi Inference enables fast operational deployment of trained models. This component is crucial for integrating ML forecasts into time-sensitive operational workflows. (https://anemoi-inference.readthedocs.io)
- **Anemoi Graphs:** Supporting custom graph generation, this module is particularly exciting for researchers exploring novel graph architectures. It already supports multi-scale GraphCast-like graphs and stretched-grid graphs showcased recently by MET Norway, with more innovations on the horizon. Graphs can be easily visualised to understand the connectivity. (https://anemoi-graphs.readthedocs.io)

Additional utility tools exist in supporting repositories beyond these core modules. Anemoi builds on top of

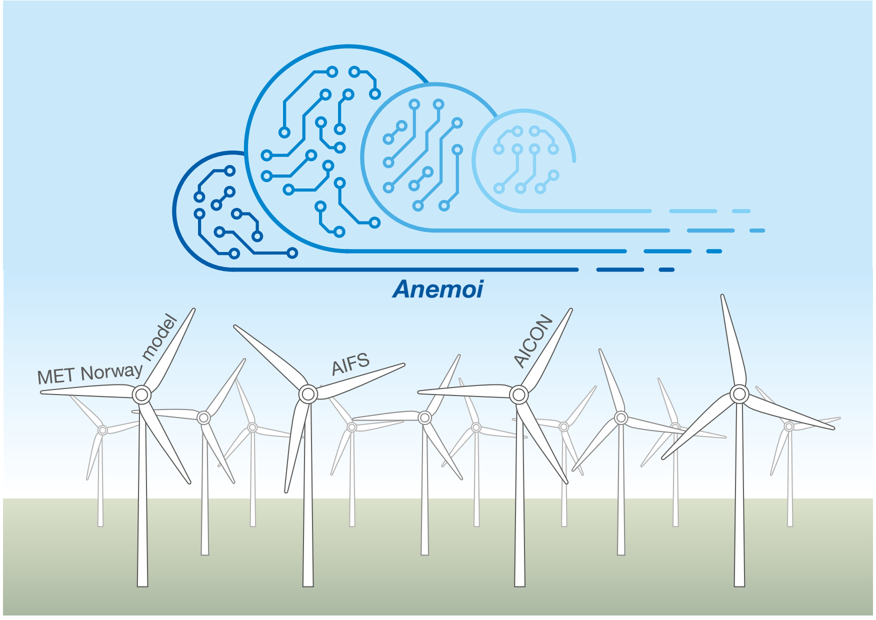

Current contributors to Anemoi

A large number of people across the organisations mentioned in this article have already contributed to the Anemoi codebase, and we'd like to thank them by naming them here:

Ana Prieto Nemesio, Baudouin Raoult, Cathal O'Brien, Christian Lessig, Ewan Pinnington, Florian Pinault, Gabriel Moldovan, Gareth Jones, Gert Mertes, Harrison Cook, Helen Theissen, Håvard Homleid Haugen, Jakob Schlör, Jan Polster, Jesper Dramsch, Kasper Hintz, Leif Denby, Magnus Sikora Ingstad, Marek Jacob, Mariana Clare, Mario Santa Cruz, Matthew Chantry, Meghan Plumridge, Michiel Van Ginderachter, Mihai Alexe, Ophelia Miralles, Rilwan Adewoyin, Sándor Kertész, Sara Hahner, Simon Lang, Simon Kamuk Christiansen, Steffen Tietsche.

ECMWF's Artificial Intelligence Forecasting System (AIFS), expanding that codebase to enable wider functionality for a greater range of users. Future AIFS models will be trained using Anemoi.

Collaborative and data-driven innovation

Anemoi represents a significant European effort to pool expertise and resources, at a time of rapid changes in weather forecasting. It is not just a technical framework but a philosophy of open collaboration, which we believe is essential for tackling the complex challenges of modern meteorology. The code is freely available on github, under a permissive licence, meaning that anyone can contribute to its development or use it for their activities.

AEMET, DMI, DWD, FMI, ITAF Met, KNMI, MET Norway, Météo-France, MeteoSwiss, RMI and ECMWF have joined together to further develop this code, with some using it as a basis for their data-driven activities. It supports collaboration across borders, leveraging the best tools and knowledge to develop new forecasting methodologies. This cooperative approach ensures that Anemoi can adapt to evolving scientific and technological advancements and will facilitate the integration of ML algorithms into existing forecasting workflows. Anemoi offers a platform for experimentation and development, encouraging stakeholders to test novel techniques and share results.

We look forward to more contributions from the organisations currently contributing to Anemoi as well as others, as we together develop a tool by the community and for the community.

The vision of Anemoi

The development of Anemoi aligns with ECMWF's commitment to improving forecast accuracy and reliability in the face of changing climate conditions.

By utilising ML weather models, Anemoi aims to significantly improve short-term and long-term forecasts, helping societies better prepare for extreme weather events.

Looking ahead, the project plans to expand its network of collaborators, welcoming new partners to contribute to this innovative initiative. This expansion will not only increase the diversity of input data but also drive further advancements in ML applications for meteorology, including new AI real-time services to support national meteorological services.

We invite the global meteorological community to join us in this exciting journey. Whether you are a seasoned ML practitioner, a meteorologist curious about new technologies, or a student looking to shape the future of weather forecasting, there might be a place for you in the Anemoi community.

To stay updated on Anemoi contributions from ECMWF, visit the AIFS blog: https://www.ecmwf.int/en/about/media-centre/aifs-blog.

Solar eclipses in IFS forecasts and (re)analyses

Philippe Lopez

Every year, two to five solar eclipses (either total, annular or partial) affect our planet for up to six hours at a time. As documented in the literature using both observations and numerical simulations, the impact of an eclipse day on meteorological conditions can be significant, even in regions where the sun is only partially eclipsed by the moon. Indeed, during an eclipse solar radiation can locally be partially obstructed for up to three hours. The moving region within which obstruction occurs is up to 6,500 km wide. The strongest meteorological effects are found over land within the planetary boundary layer, especially close to the surface, with a local cooling of up to 7°C and a local reduction in wind speed. These low-level meteorological impacts are strongest in fair-weather conditions and for high solar elevations, and they can persist for several hours after the eclipse. Furthermore, temperatures in the stratospheric ozone layer between 15 and 50 km altitude can also drop by several degrees Celsius, through a reduction in the ozone heating rate and, to a lesser extent, changes in chemical processes involving ozone. All these effects will be included in the analysis and in forecasts of ECMWF's Integrated Forecasting System (IFS) from Cycle 49r2. They will be made operational next year with the upgrade to Cycle 50r1. They will also be included in the next global reanalysis, ERA6.

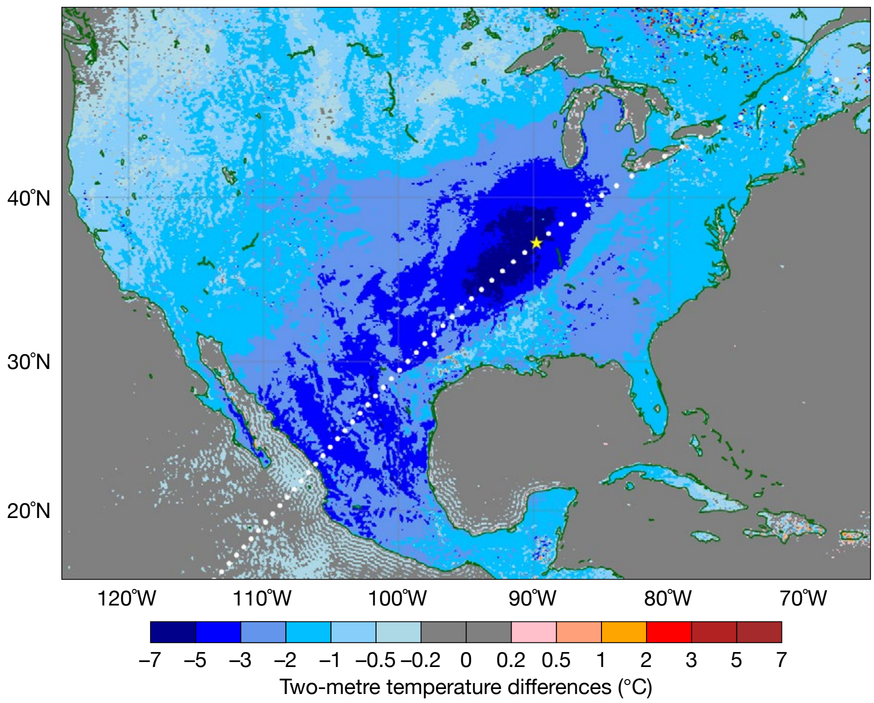

Effect of a solar eclipse on an IFS forecast. Two-metre temperature differences (°C; colour shading) between two 9 km IFS 19 h forecasts with and without solar eclipses included. The valid time is 19 UTC on 8 April 2024, during the total solar eclipse that crossed North America from the Pacific coast of Mexico towards Labrador. Small white dots show the eclipse path, while the yellow star indicates the eclipse's centre at 19 UTC.

Computation of solar eclipses

Until now, solar eclipses have been neglected in the IFS. This could occasionally lead to unwelcome large-scale errors, especially in low-level atmospheric temperatures, lasting for several hours. As a result, both analyses and forecasts were degraded, not only in ECMWF's daily operations but also in reanalyses of past weather, such as ERA5. Although the effects of solar eclipses are not visible in the long-term statistical evaluation of the IFS due to their relative rarity, it is important to capture their effects to improve the forecasts when they do occur. The IFS has therefore been modified to account for the effects of solar eclipses using accurate astronomical computations of the sun's position (VSOP87D solutions from Bretagnon and Francou, 1988) and of the moon's position (ELP-MPP02 solutions from Chapront and Francou, 2003). Over the period 1900–2100, the accuracy of the eclipse's central location on Earth as predicted by the IFS is typically within a couple of kilometres, compared to detailed calculations by NASA. Such a level of accuracy seems adequate for ECMWF's operational 9 km resolution, but also for kilometre-scale simulations used in ECMWF's contribution to the

EU's Destination Earth initiative. Larger errors may occur for solar elevations lower than a few degrees, because in these cases the lunar shadow is very elongated. However, in these situations the radiative impact of the eclipse is usually negligible. The fraction of the sun eclipsed by the moon is computed at each location on Earth and used to reduce the incoming solar radiation at the top of the atmosphere, at each model time step.

Impacts on weather forecasts

The dramatic meteorological impact of a total solar eclipse in the IFS is illustrated with an event over North America on 8 April 2024. The figure shows a map of 2-metre temperature differences between two 9 km IFS 19 h forecasts with and without the solar eclipse included. This was a total solar eclipse that crossed North America from the Pacific coast of Mexico towards Labrador. The figure highlights a drop of up to 6°C in predicted 2-metre temperature near the eclipse's maximum over land. A cooling of comparable magnitude was observed using ground station 3-second measurements kindly provided by the Purdue University Mesonet (Indiana, USA). It is worth emphasising that the sizeable effects of the eclipse are not confined to the 180-km wide band of totality. They also affect the much wider penumbral region, with a cooling exceeding 1°C over most of the USA and Mexico. The asymmetry in the cooling over the central USA is due to the presence of clouds east of the eclipse path. Another particularly strong impact in the studied case is the reduction in predicted low-level wind speeds by up to 3 m/s, in agreement with observations. Over sea and large lakes, the impact of the eclipse is negligible, mainly because of the slow thermal response of water to changes in radiation.

Inclusion in data assimilation

The handling of solar eclipses in the IFS has also been added to 4D-Var data assimilation (DA), since their absence could lead to undesirable, excessive observation–background departures, especially for temperature. The inclusion of eclipses in both 4D-Var trajectories and minimisations improves the assimilation process, not only because the fit between model and observations gets closer, but also because fewer observations are rejected through quality control in regions affected by the eclipse. Further benefits should be expected when solar reflectance satellite observations start to be assimilated, due to their direct strong dependence on incoming solar radiation.

The modernisation of the Data Stores at ECMWF

Angel López Alós

The Copernicus Climate Data Store (CDS) and Copernicus Atmosphere Data Store (ADS) services are the backbone technical components supporting the implementation of the EU's Copernicus Climate Change Service (C3S) and Atmosphere Monitoring Service (CAMS) run by ECMWF. These CDS and ADS services have recently been modernised and brought together in a Data Stores Service. This encompasses a wide range of software and data products, which now integrates with, benefits from and contributes to the more comprehensive ECMWF Software EnginE (ESEE). The ESEE is used to provide transversal data provision and workflow services across the portfolios of ECMWF, Copernicus, and the EU's Destination Earth initiative, to which ECMWF contributes. The Data Stores Service is now hosted and running on the ECMWF Common Cloud Infrastructure (CCI).
The modernised Data Stores have evolved from the original CDS infrastructure and capitalised on the experience and know-how gained from it. At the same time, they rely on open-source and cutting-edge technologies fully aligned with ECMWF's Strategy.

Key components

The operational Data Stores Service is split into two main layers, which have different functions and cloud requirements: Data Repositories and Software Services. These layers are integrated by modular components sharing common interfaces and acting jointly to perform the full range of capabilities exposed to users. Similar to what happens for other services provided at ECMWF, such as the European Weather Cloud (EWC), the modernised Data Stores Service is deployed and runs on the CCI. This ensures the elastic provision of resources for further scalability of different components, facilitates integrated management across hardware and software, and strengthens synergies with different EU cloud-related initiatives in which both ECMWF and EUMETSAT participate, such as WEkEO. Layers of the Service can be described as follows:

- **Data Repositories:** These are the foundational base of the Data Stores Service. This layer encompasses a broad range of distributed data products made available to users as part of the Services catalogue portfolio. **External** repositories are those hosted outside the CCI, while **internal** repositories are hosted within the CCI. Internals are managed as part of the Data Stores and are optimised to be accessed by its different components. The internal data repositories include an instance of ECMWF's online Meteorological Archival and Retrieval System (MARS), where a subset of the most requested variables from the core MARS archive are regularly uploaded. It also includes an observations repository; other small on-disk datasets from C3S and CAMS; and an experimental ARCO (Analysis Ready, Cloud Optimized) data lake to improve visualisation and interactivity of catalogued data on the WEkEO platform and as an extension to address the needs of demanding machine learning/artificial intelligence (ML/AI) and visual-interactive applications.
- **Software Services:** This layer

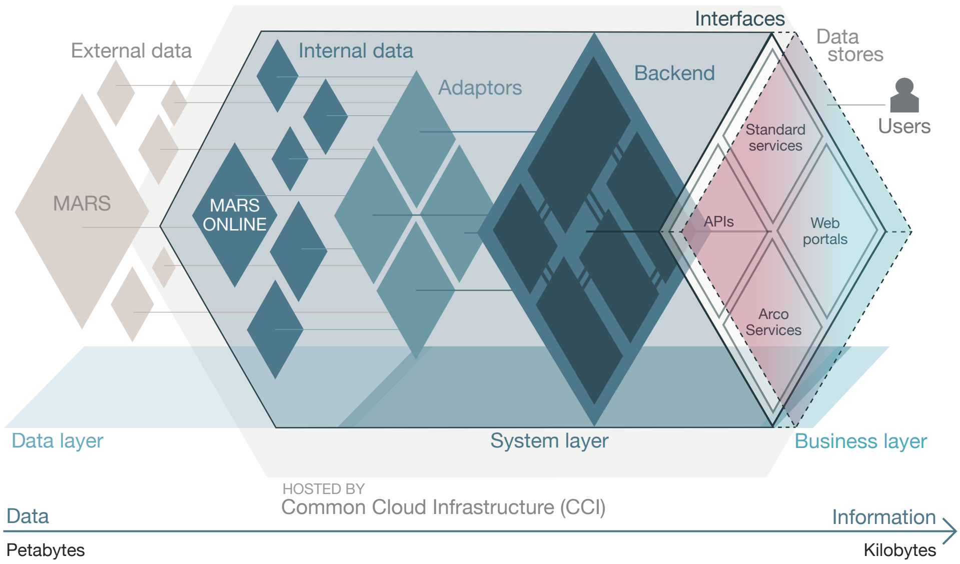

ECMWF Data Stores. The image shows the main elements of the Data Stores: the data layer, the system layer, and the business layer.

integrates all the different software components which deploy and run together to support the operational functioning of the Data Stores Service. The functional baseline of the Data Stores is to provide a seamless user journey from searching, discovering and sub-setting to retrieving data, via interactive and programmatic interfaces. These services encompass different software applications, which deploy and run on dedicated servers and clusters within the CCI. They include the Core Data Stores Engine, the Evaluation and Quality Control (EQC) function, monitoring and metrics, the Copernicus Services web portals, interactive applications, different micro-services deployed and operated by third parties as part of contractual agreements, and soon a JupiterHub development environment where users will be able to perform computation and visualisation on top of the data, using a set of preconfigured expert tools.

Data Stores software infrastructure is based on a plug-in architecture, which makes it possible to share components to power other platforms, but also to integrate third-party components to complement functional areas of the system. This layer supports interactive and programmatic interfaces (APIs), which are highly configurable and serve as an entry point for users to the data repositories and other standard services offered by the system. Most of this software has been developed directly on the CCI as part of the modernisation project. Data Store software components are cloud-oriented and ready to deploy and run in different clouds as backend engines, ‘powering’ other services and platforms.

The CDS for C3S, the ADS for CAMS, and the recently launched Early Warning Data Store (EWDS) for the EU’s Copernicus Emergency Management Service (CEMS) are the well-known public-facing interfaces of the Data Stores Service.

On the periphery of the Data Stores, but closely interlinked with them, there is a layer made of components which complement the content and functional scope of the Data Stores. Of special relevance are the following:

- **earthkit:** the ECMWF open-source Python code repository offering a broad set of expert libraries optimised to work with ECMWF and Data Store data resources. Within the Data Stores, earthkit libraries are used to foster data formatting, processing and visualisation capabilities.
- **Visual Interactive Content (VIC):** this includes a broad set of user-oriented applications and training material that showcases or makes use of the full range of data resources and capabilities of the Data Stores and earthkit via friendly interfaces. VICs can deploy and run anywhere. This enables user communities to discover, learn and interact with the available data and functionalities. The recently launched Copernicus Interactive Climate Atlas is an example of a VIC.
- **External platforms:** this includes a very broad ecosystem of platforms and infrastructures ‘powered by’ the Data Stores. These platforms may interact with or consume data resources via the exposed interfaces, embed technical components or integrate VIC as part of their portfolio.
  The Copernicus WEkEO DIAS Platform – a partnership of ECMWF, the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT), Mercator Ocean and the European Environment Agency (EEA) – is an example of such a platform.

Like an iceberg

The ECMWF Data Stores Service is a complex, multi-layer system that can be conceptually likened to an iceberg: simple interfaces on the surface, with a robust and scalable backend underneath providing seamless access to a broad set of catalogued resources.

The modernisation of the former Data Stores has touched all the different layers of the architecture. It became necessary to overcome the obsolescence of former components and make the system evolve in a changing and highly demanding environment. The Data Stores Service will thus play a central role as an efficient, versatile and trusted transmission link between big, distributed and heterogeneous data sources and increasingly sophisticated user requirements, driven by the development of new technologies and the need for immediate data and information.

Increasing the lead time for early warnings

Christian M. Grams (MeteoSwiss), Joshua Dorrington (University of Bergen, Norway), Federico Grazzini (ARPAE-SIMC, Italy), Linus Magnusson, Frédéric Vitart (both ECMWF), Marta Wenta (AXPO Solutions AG, Switzerland)

ECMWF has taken part in a project that aimed to increase the lead time for early warnings by ensuring that the forecasting of extreme events is dynamically informed. In this 'Transfer Project' of the collaborative research centre Waves to Weather (SFB/TRR165) funded by the German Research Association (DFG), two forecast prototypes, DOMINO and MaLCoX, were developed. It was realised together with scientists at the Karlsruhe Institute of Technology (KIT, Germany), the University of Munich (Germany), and the Italian regional meteorological service of Emilia-Romagna (ARPAE). ECMWF's role came about through its collaboration with ECMWF Fellow Christian Grams and his group at KIT. The central aim was to make use of the wealth of knowledge about dynamical precursors for extreme events. While the prototypes showcase heavy precipitation events in northern Italy, they are generally applicable to any type of weather event for which some understanding of dynamical precursors is available.

Forecasting extreme weather events is the most important task of operational numerical weather prediction (NWP). However, predicting details of local extremes beyond 1–3 days in advance remains challenging. Research has shown that most events have specific large-scale dynamical precursors, which depend upon the type of extreme weather, the region, and the season. For instance, springtime heavy precipitation events in northern Italy are typically embedded in a highly amplified Rossby wave pattern downstream of a trough in a region of enhanced integrated water vapour transport (IVT). Such dynamical precursors can provide a narrative chain of the unfolding of an extreme event. However, operational forecasting procedures do not yet systematically incorporate this knowledge.

DOMINO

The first tool, DOMINO (Dorrington et al., 2024a, https://doi.org/10.1002/qj.4622), focuses on the large-scale circulation patterns that modulate the likelihood of extreme events. Given a list of extreme event dates, the framework automatically identifies the precursor patterns in any set of candidate variables in ECMWF's ERA5 reanalysis, e.g. 500 hPa geopotential height (Z500) or IVT. Based on these patterns, DOMINO computes a standardised 'activity index', which estimates the probability of an extreme. An elevated precursor activity occurred for northern Italy during 15–17 May 2023, resulting in devastating floods in the Emilia-Romagna region. The medium-range ensemble precipitation forecast of ECMWF's Integrated Forecasting System (IFS) indicated the extreme event with a lead time of about 3 days. However, only 1.5 days before the event (forecast initialised on 14 May) did the ensemble centre around the observed precipitation (see the first plot of the figure and Dorrington et al., 2024b, https://doi.org/10.5194/nhess-24-2995-2024). In contrast, already from 8 May the Z500 and IVT precursors showed elevated risk with the ensemble

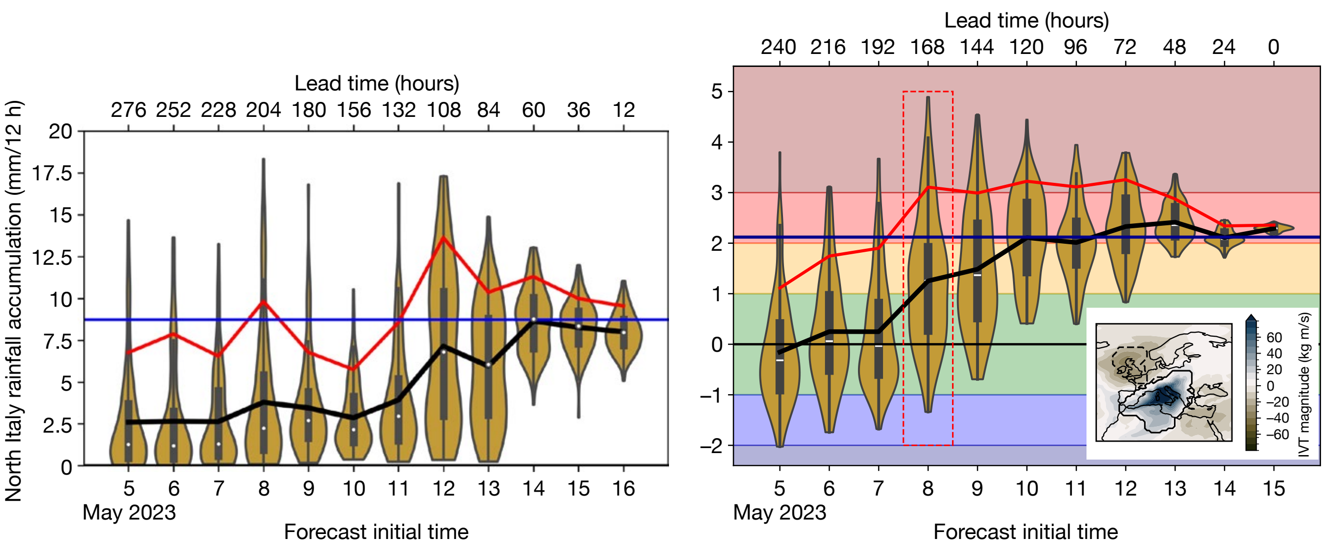

Ensemble precipitation and precursor index forecasts. We show on the left an ensemble forecast of accumulated precipitation from 12 UTC on 15 May to 00 UTC on 16 May (area averaged in mm/12 h, ensemble distribution as box-and-whisker-violin, ensemble mean in black, high-resolution forecast in red, estimate of actual value in blue), and on the right an ensemble forecast for the IVT precursor index on 15 May with forecast indices color-coded (-1–1 green = no risk, >1 yellow = watch, >2 red = alarming). The inset shows the climatological IVT anomaly during springtime heavy precipitation events in northern Italy. Adapted from Figures 4, 12, and 13 of Dorrington et al., 2024b, https://doi.org/10.5194/nhess-24-2995-2024, published under CC BY Licence 4.0.

median slightly above a pre-warning level (yellow >1) and 25% of the ensemble members above 2 (red) for IVT (see the second plot of the figure). Thus the dynamical precursor ensemble forecast indicated the event 2–5 days earlier compared to direct precipitation forecasts, extending the lead time for early warnings from 3 days to up to 8 days. DOMINO also makes it possible to explore a potential predictability barrier for an event and to assess – a priori – when a forecast becomes more reliable (details in Dorrington et al., 2024b).

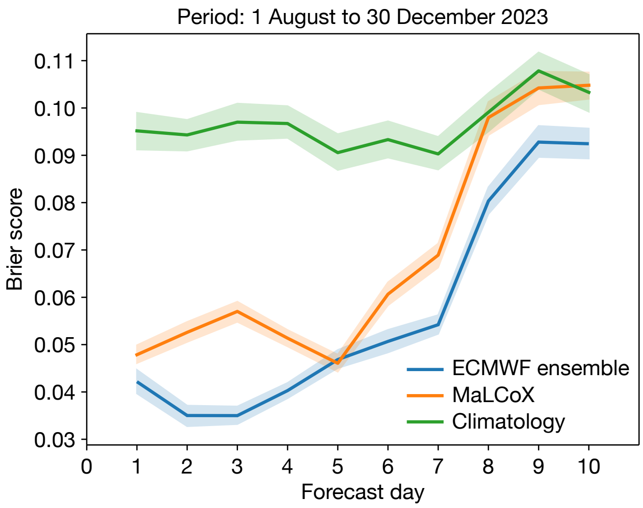

Performance of MaLCoX. Brier score for precipitation forecasts of MaLCoX (orange) compared to the direct ensemble output (blue) and climatology (green). The shading represents 95% confidence intervals.

MaLCoX

The second tool, MaLCoX (Machine Learning model predicting Conditions for eXtreme precipitation, Grazzini et al., 2024, https://doi.org/10.1002/qj.4755), recognises favourable synoptic conditions leading to precipitation extremes and subsequently classifies extremes into three different categories according to the presence of convection. It is tailored for extreme precipitation events in northern and central Italy but is extensible to other regions and timescales. MaLCoX, which is based on random forest architecture, uses different groups of predictors, including local (point-based) predictors and the precursor indices from DOMINO as non-local (field-based) predictors. The non-local predictors are particularly relevant and increase the skill of MaLCoX for medium-range lead times > 6 days.

MaLCoX has been implemented operationally at ARPAE Emilia-Romagna (Bologna) and shows comparable skill to direct precipitation output from the ensemble, in particular in situations with strong dynamical forcing (see the second figure). This is remarkable as MaLCoX is solely trained on the control forecast and currently uses no dynamical ensemble information as input. An advantage over ensemble forecasts comes from interpretability. MaLCoX is able to attribute its decision (extreme ‘yes’ or ‘no’, and category) to individual components, thus providing forecast products which allow dynamical insights (e.g. anomalous IVT values vs. wave amplitude). Understanding why a model makes a certain prediction can be as crucial as the prediction's accuracy, especially when forecasters have to face rare conditions and they need to gain trust in model output.

Applicability

According to discussions with forecasters at ARPAE Emilia-Romagna and MeteoSwiss, such early warnings based on dynamical precursors are particularly useful for ‘low regret’ actions, for which measures to mitigate negative impacts can be taken at a low cost and some false alarms are acceptable. The MaLCoX suite at ARPAE Emilia-Romagna (Bologna) is already fine-tuned for heavy precipitation events in northern Italy, and it runs operationally at ARPAE. For more information, please contact Frederico Grazzini.
The DOMINO workflow is implemented on ECMWF’s ecgate server and is available to ECMWF Member States as an easy-to-use preconfigured software package.
For more information, please contact Linus Magnusson, Christian M. Grams, or Joshua Dorrington.

Code for Earth 2024 – successful completion of 13 challenges

Athina Trakas, Esperanza Cuartero

Code for Earth is an ECMWF activity focused on innovation, collaboration and open-source development.
It bridges the gap between creative minds and real-world challenges. Since 2018, selected developer teams and individuals from outside ECMWF have come together each summer with ECMWF mentors and partners on innovative projects in data science; weather, atmosphere, climate or other Earth sciences; visualisation; and more. So far, the programme has produced over 50 open-source projects, some of which have been beneficial to activities at ECMWF and its Member and Co-operating States. For example, the Jupyter Notebooks developed in one of the 2023 projects will be used in core ECMWF training resources, and code developments in land verification from another 2023 project have been incorporated into the LANDVER package. In this year’s edition, Code for Earth has seen the successful completion of 13 challenges under the headings of data visualisation, machine learning, and software development.

Code for Earth in numbers

Started in **2018**

Has produced **59** projects

Has involved **107** mentors

Has attracted **126** participants

Code for Earth 2024

Each year, the programme outlines a set of challenges and invites individuals and developer teams to submit proposals through an open Call for Participation. Once selected, teams work on their chosen challenge during a four-month coding phase, transforming the initial ideas into actionable projects as they develop their solutions.

This year's Code for Earth edition benefited from expanding its impact through strengthened partnerships, in particular with IFAB, the International Foundation Big Data and Artificial Intelligence for Human Development, and additionally through so-called Joint Challenges. These were developed jointly with the following partner organisations: the European Environment Agency, the University of Reading (UK), the University of Bonn/ Center for Earth Observation and Computational Analysis (Germany), and the Helmholtz-Zentrum Hereon (Germany).

From the proposals submitted, 13 talented developer teams were selected to tackle challenges faced by scientists and technical specialists working on real-world issues. Each team was guided by an ECMWF mentor and, for Joint Challenges, mentors from partner institutions – all experienced specialists in their respective areas. The teams worked on the following projects:

Stream: Data visualisation and visual narratives for Earth science applications

- **CAMS Verisualiser: A web application for interactive verification results.** This project developed a web application with a Python backend and JavaScript frontend to visualise CAMS verification data interactively, allowing a user to change scale, zoom in, etc. CAMS is the EU's Copernicus Atmosphere Monitoring Service, implemented by ECMWF.
- **CAMS-nb-Charts: Jupyter notebooks for CAMS web charts.** CAMS-nb-Charts created reusable Jupyter notebooks to programmatically reproduce CAMS forecast charts. Using a flexible template and configuration file, the project enables easy customisation and data visualisation across various variables.
- **vAirify: An air quality dashboard.** The project developed a dashboard that compares CAMS air pollution forecasts with in-situ measurements. Forecasters can visualise variations and zoom into local areas via an intuitive map view, helping them to improve model accuracy.
- **Tales of Dry Lands: Contextualising Earth's water story.** This project used Python notebooks to help users understand and visualise climate data related to droughts. It simplified complex drought patterns, making them accessible to a broader audience.
- **SunVizor: Visualising ECMWF and Copernicus data for the renewable energy community.** SunVizor developed a user-centric web application to visualise and compare solar energy data, addressing the need for better tools in the solar energy sector.

Stream: Machine learning for Earth science applications

- **HydroGap-AI: Bridging gaps in streamflow observations with ML-driven solutions.** The project focused on filling gaps in daily streamflow time series using advanced machine learning techniques. It provides an open-source Python package to ensure more accurate and continuous hydrological data.
- **Project Polly: A natural language processing interface to extract complex features from weather datacubes.** The project developed a Natural Language Model (NLM) capable of extracting specific weather information from large datasets. It integrates with a chatbot to assist both technical and non-technical users in accessing complex weather data.
- **XAI for weather forecasting models.** The project focused on analysing training phase data of AI weather forecasting models (PanguWeather and ECMWF's Artificial Intelligence Forecasting System – AIFS). It explored input–output relationships for process-based analysis to improve forecasting accuracy.
- **ML-BEES: Using machine learning to emulate the Earth's surface.** This project evaluated and improved the performance of ECMWF's land surface ML model prototype by validating its accuracy and comparing its output to in-situ observations. It developed an

Group photo. Participants in the Code for Earth Final Event in Reading, UK, in September 2024.

evaluation framework for land surface model emulators.

- **KGB-TruthGuiding: Knowledge graph generation for enhanced chatbot and scientific literature synthesis.** The project built a knowledge graph from scientific documents to enhance Large Language Models (LLMs) and chatbots. It aimed to improve the ECMWF-assistant chatbot by providing more interactive, explainable and engaging responses.

Stream: Software development for Earth science applications

- **Optimising CDSAPI Datasets Retrieval: Advance user capabilities to handle data constraints when using CDSAPI.** The project developed a mechanism to optimise data retrieval requests via the Climate Data Store API (CDSAPI). A key feature was to break down user requests into multiple valid sub-requests that adhere to data constraints.
- **CDSAPI Request Check.** The second project on advancing user capabilities when using CDSAPI produced a Python library that validates CDSAPI data requests in advance, increasing the chances of successful data retrieval, reducing time, and minimising server costs.
- **AirQuality Urban View: Regional to urban air quality mapper.** The project achieved downscaling regional air quality data to urban levels and visualising it through intuitive maps, providing precise air quality insights for European cities.

In 2024, the Final Event celebration took place at the end of September at ECMWF's headquarters in Reading, UK. All teams showcased their projects and solutions. These innovative outcomes are available on GitHub (https://github.com/ECMWFCode4Earth) and may be incorporated into operational systems. The participants' and mentors' hard work and commitment played a key role in another successful Code for Earth edition.

A look to the future

Looking forward to 2025, the programme will keep evolving to address new challenges and partnerships. With rapid advancements and shifting user needs, Code for Earth will identify key opportunities and turn them into challenges for the 2025 edition. For more details and feedback, visit our website (https://codeforearth.ecmwf.int/) or contact us directly at codeforearth@ecmwf.int.

Evaluating km-scale simulations in Destination Earth

Estíbaliz Gascón Salvador, Michael Maier-Gerber, Benoît Vannière, Žiga Zaplotnik, Tobias Becker, Linus Magnusson, Matthieu Chevallier, Irina Sandu

As part of the European Commission's Destination Earth (DestinE) initiative, ECMWF is developing a digital twin focused on weather-induced extreme events. This digital twin (Extremes DT) includes a global component, created using ECMWF's Integrated Forecasting System (IFS) at 4.4 km resolution, and a regional component over Europe, at 500 to 700 m resolution, developed by a consortium led by Météo-France involving several national meteorological services.

ECMWF has been running the global Extremes DT daily for over a year. To assess its performance, ECMWF conducted a preliminary evaluation of five-day 4.4 km forecasts initialised daily from 1 September 2023 to 31 August 2024 from the ECMWF operational analysis. This evaluation provided valuable insights into the added value of km-scale simulations, particularly regarding the representation of extreme weather events, and for future model improvements.

The evaluation shows that the global Extremes DT improves the prediction of certain extreme phenomena compared to ECMWF's 9 km operational forecasts. For example, the 4.4 km resolution enhances the intensity prediction of tropical cyclones (TC) in general, but particularly for some rapidly intensifying TCs. It also improves the representation of extreme precipitation in mountainous areas, thanks to better orographic representation.

Enhanced precipitation forecasts over complex orography

In terms of extreme precipitation, this one year of data has demonstrated that the 4.4 km resolution of the Extremes DT brings an improvement of the global forecast performance in complex orographic areas, compared to the 9 km resolution. This is noticeable across most percentiles of the precipitation distribution, as shown in the first figure. Both the overestimation of small 24 h precipitation values and the underestimation of very large precipitation values observed in the 9 km forecast are reduced, which leads to a smaller absolute bias.

Increasing the resolution from 9 km to 4.4 km also significantly reduces the occurrence of false alarms for extreme convective precipitation over small but complex orographic islands (i.e. Tenerife). These false alarms are due to an overestimation of the moisture convergence closure of the convective parametrization when small mountain ranges with steep slopes are under-resolved. Higher resolutions mitigate this issue.

Higher resolution enables more accurate tropical cyclone intensity prediction

The evaluation of all TCs occurring globally in the one-year verification period clearly demonstrates a great benefit of running the IFS at higher resolution for TC intensity prediction. While the TCs are always 12–15 knots weaker on average than observed in 9 km forecasts up to 120 h ahead, the initially equally large bias in the 4.4 km forecasts steadily reduces until it almost vanishes on day five (see the second figure). This confirms our experience from monitoring forecast performance for individual TCs during that one year: DestinE forecasts were usually better at predicting both TC intensification and peak intensity.

The fact that the bias at initialisation time remains almost unchanged highlights that the resolution increase in the forecast model should ideally be accompanied by a higher-resolution 4D-Var analysis. Initial tests in this direction are promising as they show, for example, that the explosive rapid intensification of TC Otis (2023) could have been accurately predicted if the DestinE forecast had been initialised from a higher-resolution 4D-Var with a 12-hour window, together with a reduced dependence on the convection scheme.

Addressing remaining challenges

Despite promising results from this one-year evaluation of Extremes DT daily forecasts, some challenges remain. An example is the overestimation of very cold 2 m temperatures globally in very stable boundary layers. This is likely caused by increased vertical wind shear due to better resolved orography. However, this aspect has significantly improved in DestinE 4.4 km simulations carried out with the new IFS Cycle 49r1, which is to be implemented in November 2024. This is thanks to the inclusion of 2 m temperature in data assimilation and enhanced 2 m diagnostics. Some issues with convective precipitation still persist, such as the failure to propagate marine convective precipitation further inland. Tuning the cloud base convective mass flux in the convective parametrization has shown potential to help, but as it causes some forecast degradation in the tropics, it still requires further investigation. In view of the enhancements that high-resolution 4D-Var has brought in the case of TC Otis, as mentioned above, we will undertake a more comprehensive investigation into its added value for the representation of extreme weather events. These aspects will be key points of investigation for Phase II of DestinE (2024–2026).

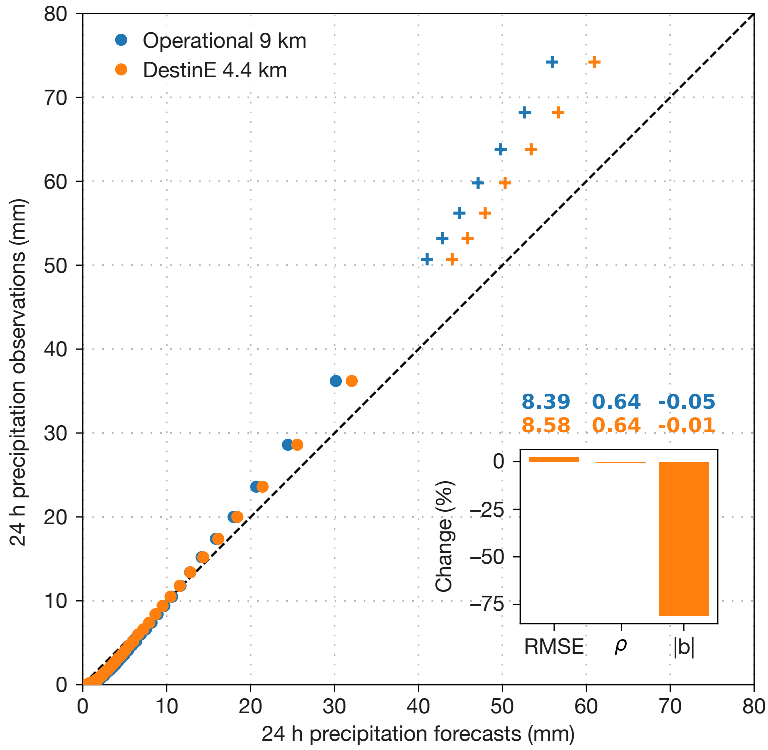

Comparison of precipitation forecasts. We compare one year of operational (9 km, blue) and DestinE (4.4 km, orange) 24 h precipitation forecasts valid on day 4. Specifically, we show the results of a quantile–quantile evaluation against SYNOP weather station observations in areas with complex orography. Dots are 1–98 percentiles, and crosses are percentiles from 99 to 99.9. The closer the dots and crosses are to the diagonal line, the less the forecast distribution diverges from the observation distribution. The inset bar plot shows the changes of root-mean-square error (RMSE), correlation ρ, and absolute bias |b| for forecasts at 4.4 km resolution with respect to the 9 km forecast. Numbers above the bar plot show raw verification indices for each forecast model.

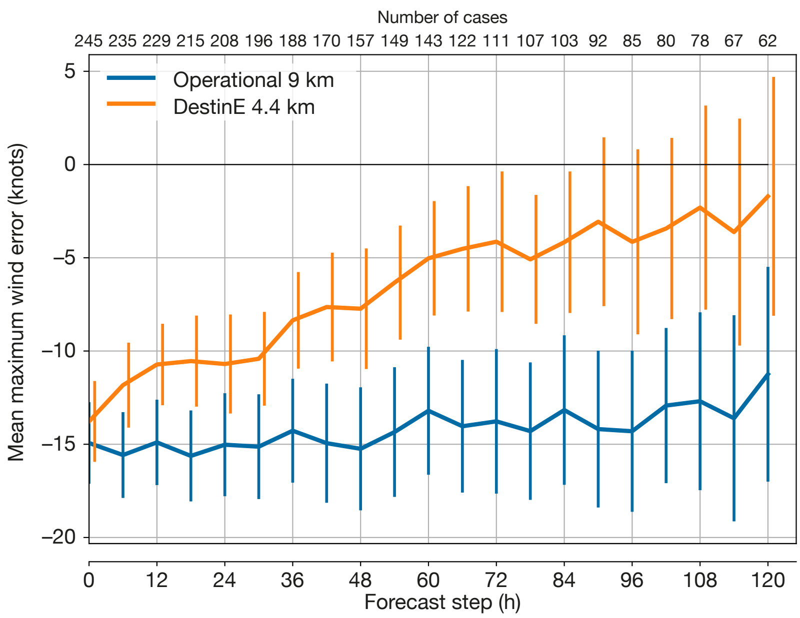

Comparison of TC intensity forecasts. We compare one year of mean maximum wind speed errors of TCs from all ocean basins in operational (9 km, blue) and DestinE (4.4 km, orange) forecasts. The error is shown as a function of lead time. Vertical bars indicate 2.5%–97.5% confidence intervals. IBTrACS observations serve as reference, and a 17 m/s intensity threshold is applied to the observations to only consider TCs of at least tropical storm intensity. The numbers on top indicate the number of TCs found per lead time in both operational and DestinE models.

IFS upgrade improves near-surface wind and temperature forecasts

Christopher D. Roberts, Bruce Ingleby, Alan Geer, Elias Hólm, Martin Janousek, Fernando Prates, Mark Rodwell

The latest update to the ECMWF Integrated Forecasting System (IFS) is due to be implemented operationally in November 2024. IFS Cycle 49r1 is a major upgrade to the IFS model and associated data assimilation system. Among many other changes, it includes the assimilation of 2 m temperature observations; increased resolution and soft re-centring of the Ensemble of Data Assimilations (EDA); activation of the Stochastically Perturbed Parametrizations (SPP) scheme for model uncertainty in all ensemble applications; extended use of microwave data over sea ice; higher-resolution data from infrared satellite sounders; and improvements to land-surface modelling and assimilation methodology. This upgrade substantially improves 2 m temperature and 10 m wind speeds, particularly for the winter months in the northern hemisphere. Cycle 49r1 also brings changes to the resolution of the ocean wave model and the frequency of medium-range and sub-seasonal re-forecasts, which will now run on fixed days of the month.

Forecast model

An essential ingredient of a reliable ensemble forecast system is an accurate representation of the uncertainties associated with parametrized physical processes. One of the major changes in IFS Cycle 49r1 is the activation of the Stochastically Perturbed Parametrizations (SPP) scheme, which replaces the effective and long-serving Stochastically Perturbed Parametrization Tendencies (SPPT) scheme in all ensemble configurations. The main impacts of the switch to SPP include:

- improved physical consistency due to local conservation of moisture and energy, which was not respected by SPPT
- an improved representation of uncertainties in the near-surface boundary layer, which contributes to improved probabilistic skill of 2 m temperature and 10 m winds in medium-range forecasts
- improvements to the spread–error relationship of sub-seasonal forecasts of the Madden–Julian Oscillation (MJO), and
- an increase in the frequency of tropical cyclones (TCs).

Further information about this revision of the IFS model uncertainty scheme, including the scientific motivations and implementation details, are provided by Leutbecher et al. in this Newsletter.

Cycle 49r1 includes several scientific and technical changes to the ocean wave model, including a revision of the horizontal grid to match the atmosphere resolution in all forecasts. This corresponds to a reduction of the wave model grid spacing to ~9 km (TCo1279) in medium-range forecasts and ~36 km (TCo319) in sub-seasonal forecasts. To make this increased resolution affordable, the number of frequencies in wave spectra output (but not online computations) is reduced from 36 to 29 frequencies. The most significant scientific changes are to the wind input parametrizations, including a new gravity–capillary model and non-linear growth rates. These updates modulate variations in the drag coefficient with wind speed, which addresses a known underestimation of extreme ocean wind speeds. There are also new sea-state-dependent heat and moisture fluxes. These changes have a strong impact on air temperatures over the oceans throughout the troposphere, primarily in the tropics and winter hemisphere (Bidlot & Janssen, 2024).

The land surface benefits from several IFS model updates introduced in Cycle 49r1. These include a single-layer urban canopy model implemented as a new surface tile, which improves 2 m temperature and 10 m wind speed forecasts over urban areas (McNorton & Balsamo, 2023). Vegetation and leaf-area index maps are updated with new versions from the European Space Agency (ESA) Climate Change Initiative and the new Copernicus Global Land Services-based leaf area index (LAI), respectively (Boussetta & Balsamo, 2021). The overall effect of these modifications is to increase low vegetation cover, reduce high vegetation cover, and improve the representation of seasonal variations in leaf area index. To allow for better near-surface and atmospheric forecast skill, these changes were combined with additional land surface updates including a new soil moisture stress function, a new interpolation method to diagnose 2 m temperatures, and an improved

doi: 10.21957/crx2bn4is8

representation of snow shadowing under high vegetation. Additional information about these changes to the land-surface model and their impact on 2 m temperature forecasts are provided by Ingleby et al. (2024).

Other updates to the IFS forecast model in Cycle 49r1 are summarised below:

- Improvements to short-wave radiation biases by assuming a liquid phase for mid-level convection if temperatures exceed –20°C.
- Improvements to convection–surface coupling and downdraught scaling, which improves nighttime convective organisation.
- A revision of the diagnostic 10 m wind calculation, which removes a limiter and modifies the blending height, leading to reduced 10 m wind biases.
- Introduction of a time-varying source of stratospheric water vapour from methane oxidation.
- A new flexible treatment of aerosols in the IFS radiation code, which includes the potential for future 'hybrid' configurations that combine climatological and prognostic aerosol species.
- Many updates to the Copernicus Atmosphere Monitoring Service (CAMS) modelling systems (IFS-COMPO and IFS-GHG), including more up-to-date anthropogenic emissions, a new wetland emission model for methane, and major revisions to the aerosol model, including updated optical properties and a simple representation of stratospheric aerosols.

Data assimilation and observation usage

Cycle 49r1 includes substantial changes to the IFS data assimilation system, including several updates that improve near-surface weather forecasts. Daytime and nighttime 2 m temperature and 2 m humidity observations from SYNOP weather stations are now assimilated within the atmospheric 4D-Var system. Previous IFS versions assimilated only 2 m humidity during the daytime. This change has a strong positive impact on short-range 2 m temperature forecasts, particularly during the northern hemisphere winter. Several changes to the land data assimilation system also improve near-surface weather forecasts substantially. These include:

- a lapse-rate correction for 2 m temperature that accounts for differences between real-world and simulated orography
- updated background errors in the soil moisture analysis, and
- improved thinning of Interactive Multisensor Snow and Ice Mapping System (IMS) snow cover data combined with an updated diagnostic model for snow cover, which enables the activation of snow cover assimilation over mountainous areas.

These changes constitute a first step towards a more unified land data assimilation system, and they pave the way for further land–atmosphere coupled data assimilation developments (de Rosnay et al., 2022). Their impacts on 2 m temperature forecasts are described in more detail by Ingleby et al. (2024).

The Ensemble of Data Assimilations (EDA) receives a

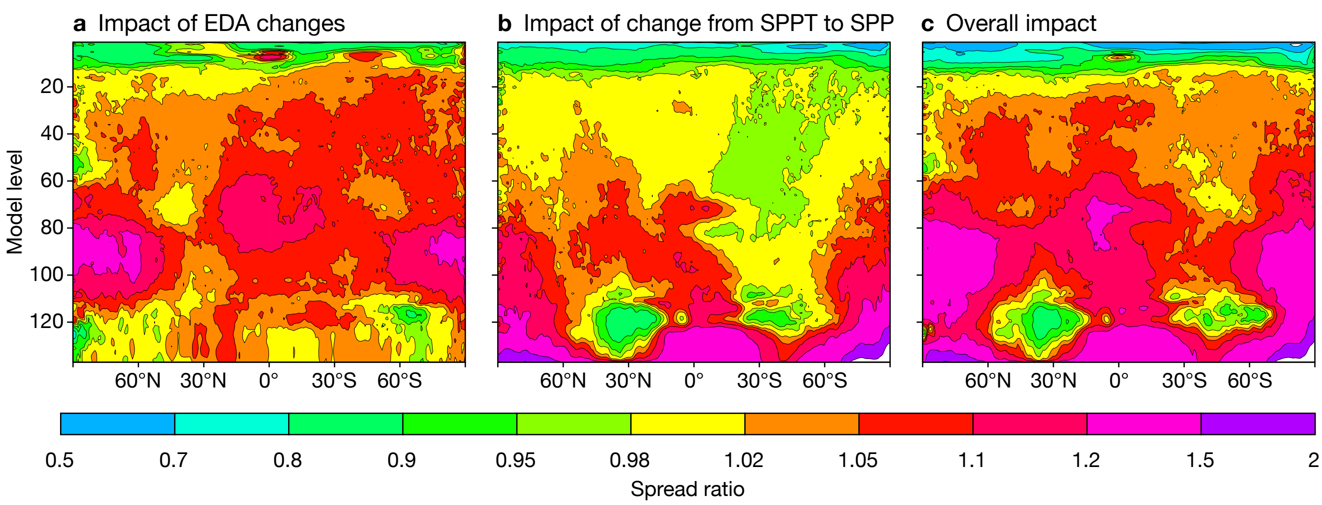

FIGURE 1 Contribution of EDA and model uncertainty changes to spread in Cycle 49r1 for zonal average temperature, showing (a) the impact of the soft re-centred EDA formulation, combined with a reduction of the horizontal grid-spacing of the EDA outer/inner-loop resolution from ~18 km/~100 km to ~9 km/~50 km – this is mostly above the boundary layer; (b) the impact of changing from SPPT to SPP at a grid spacing of ~9 km – this is mostly in the boundary layer; and (c) the total effect of soft re-centring, resolution change, and the switch to SPP. The plots are for the period of 17 December 2021 to 6 January 2022.

major upgrade in Cycle 49r1. Although the ensemble forecast resolution was increased from 18 km (TCo639) to 9 km (TCo1279) in Cycle 48r1, the resolution of the EDA was not changed. In Cycle 49r1 the horizontal grid spacing of the EDA outer loop resolution is reduced to ~9 km and the inner-loop grid spacing is reduced from ~100 km to ~40 km for the control member and to ~50 km for perturbed members. This substantial increase in horizontal resolution is made affordable by soft re-centring of each member (1 outer loop) around a more accurate control member (3 outer loops). Hólm et al. (2022) provide further background on the motivation for this soft re-centring approach. The combination of increased resolution and activation of SPP results in a general increase of EDA spread, though the magnitude of this effect varies regionally and with height. Figure 1 shows the changes in the spread of temperature, where the resolution increase and the soft re-centring contribute most above the boundary layer, and activation of SPP predominantly increases spread in the boundary layer. The combined effect is a 15–20% increase in spread in the extratropics, which leads to a significant improvement in the reliability of the EDA. These changes in spread impact the ensemble forecasts via initial perturbations but also deterministic analyses and forecasts, through the influence of the EDA on 4D-Var background error covariances. The impacts on ensemble variance as a function of horizontal scale are shown as power spectra in Figure 2. The increased EDA resolution in Cycle 49r1 removes the drop in ensemble variance previously evident at scales smaller than 200 km (wavenumber around 100). The variance spectrum at synoptic and planetary scales is also much smoother.

An important step towards a fully integrated Earth system assimilation system in Cycle 49r1 is the activation of microwave imaging radiances over sea-ice surfaces within the atmospheric 4D-Var component. This was not previously possible due to inadequate knowledge of either the sea-ice concentration (SIC) or the surface radiative properties of the sea ice. The problem has been solved using an innovative combination of machine learning and data assimilation to train a new sea-ice surface emissivity model for microwave radiances. A crucial aspect of the model is that it takes empirical input parameters that characterise the radiative properties of the sea ice. The empirical parameters summarise important aspects of the sea ice relating to things like the aging

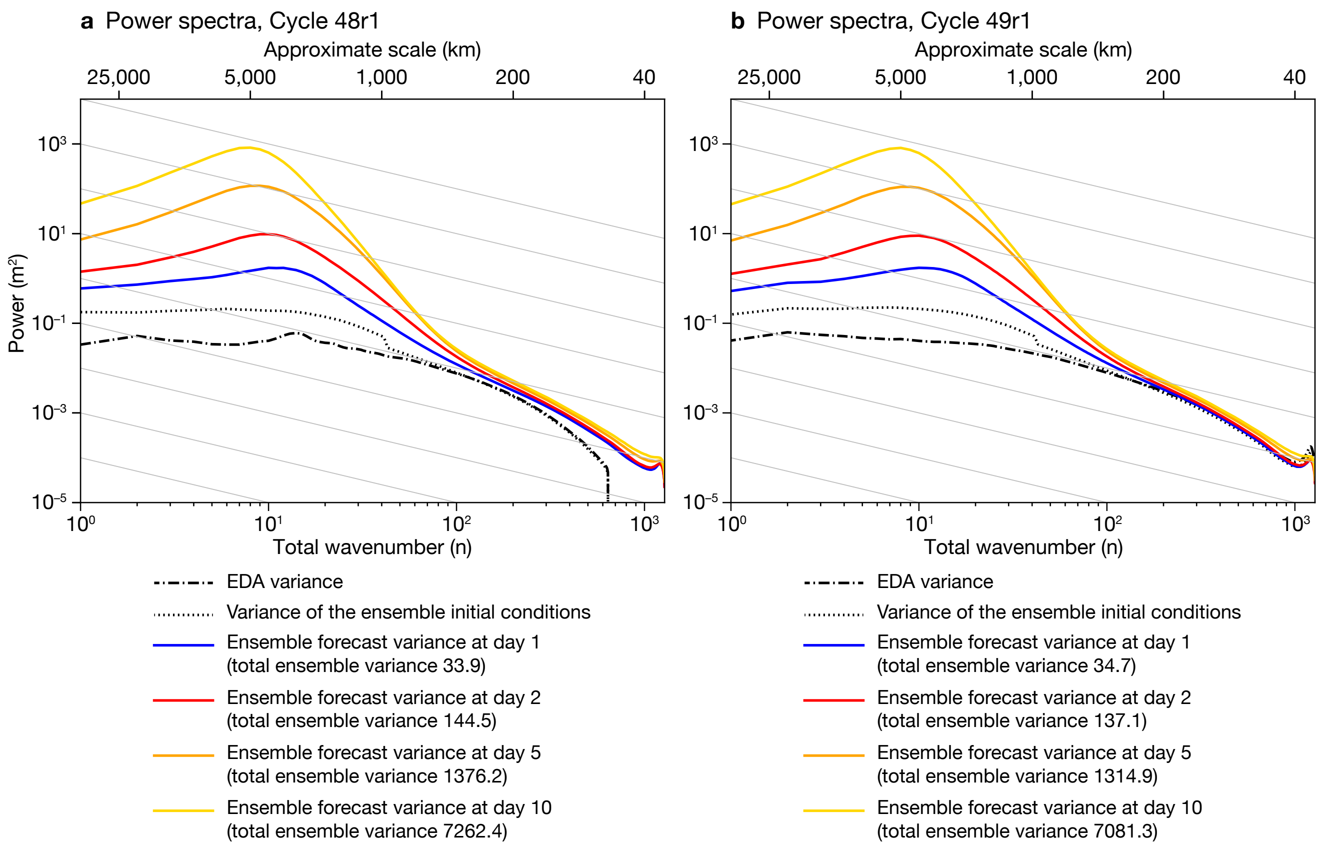

FIGURE 2 Global power spectra of the (ensemble) variance of the EDA and of the ENS for 250 hPa geopotential height against total wavenumber n, valid at 00 UTC during December 2023 – February 2024 for (a) Cycle 48r1 and (b) Cycle 49r1. The total ensemble variance (i.e. the sum over all wavenumbers) is provided in the legend. The use of a logarithmic *x*-axis means that wavenumbers become more tightly packed from left to right. To help discern the impact of this, diagonal lines represent contours of variance per linear unit length on the *x*-axis (with value proportional to the *y*-intercept). Approximate length scales are indicated on the top axis.

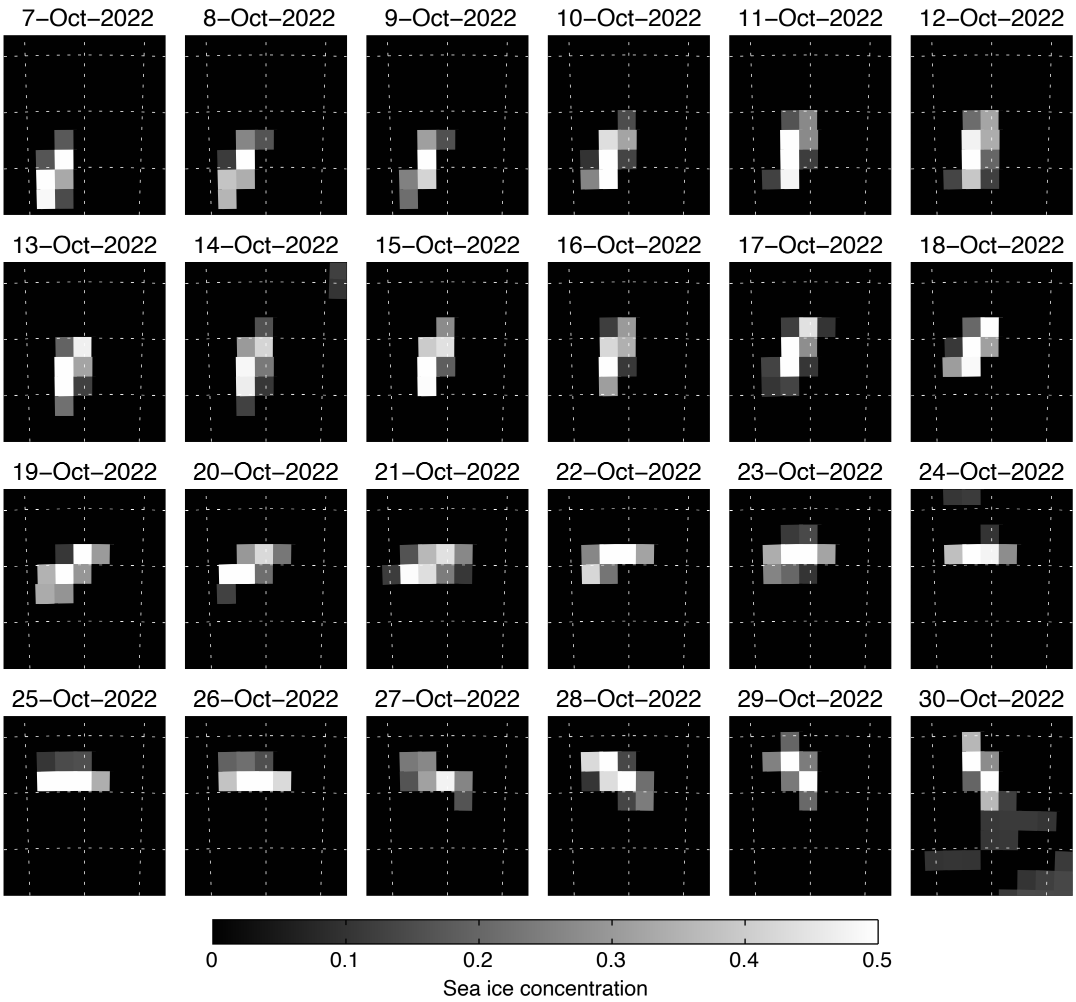

FIGURE 3 Daily maps of the extent of long-lived giant iceberg A76A during October 2022, based on the new AMSR2 SIC retrievals within atmospheric 4D-Var. Iceberg A76A was at this time around 135 km by 26 km in size and rotating in the currents of the Southern Ocean near the Antarctic Peninsula. The pixel size of the AMSR2 retrievals is 40 km by 40 km. Grid lines are shown every 1° latitude and 2° longitude, and the projection is centred on 59.25°S, 50°W. The figure is re-used under creative commons attribution licence from Geer (2024, *https://doi.org/10.1002/qj.4797*).

process of the ice as well as the snow cover on top. These aspects are poorly known, but they have a very strong control over the radiative signature of sea ice observed by microwave imagers. Using an observation space control variable (often referred to as a sink variable), the 4D-Var control vector is extended to estimate SIC and the empirical sea-ice properties at each observation location. The addition of data from Advanced Microwave Scanning Radiometer 2 (AMSR2) and Global Precipitation Measurement Microwave Imager (GMI) in sea ice and neighbouring areas improves forecast scores in the vicinity of Antarctica by around 0.5% out to day 4. The SIC retrievals are of good quality and could be provided as an input to the ocean/sea-ice data assimilation system in a future IFS cycle. An example of the quality of the SIC observations is given in Figure 3, which shows the evolution of a giant iceberg (note that microwave imagers sense any ice coverage on the ocean surface, which differs from stricter definitions of sea ice).

Further details can be found in Geer (2023, 2024).

Other data assimilation and observation usage contributions to Cycle 49r1 include the following:

- A 140% increase in the use of satellite microwave humidity data.
- Activation of additional channels from the Advanced Microwave Sounding Unit-A and the Special Sensor Microwave Imager/Sounder.
- Several updates to non-microwave observations, including reduced thinning of Spinning Enhanced Visible and InfraRed Imager data, assimilation of ground-based Global Navigation Satellite System data, and scene-dependent observation errors for Cross-track Infrared Sounders.
- A vertical extension of Global Navigation Satellite System Radio Occultation assimilation from 50 km to 60 km altitude.

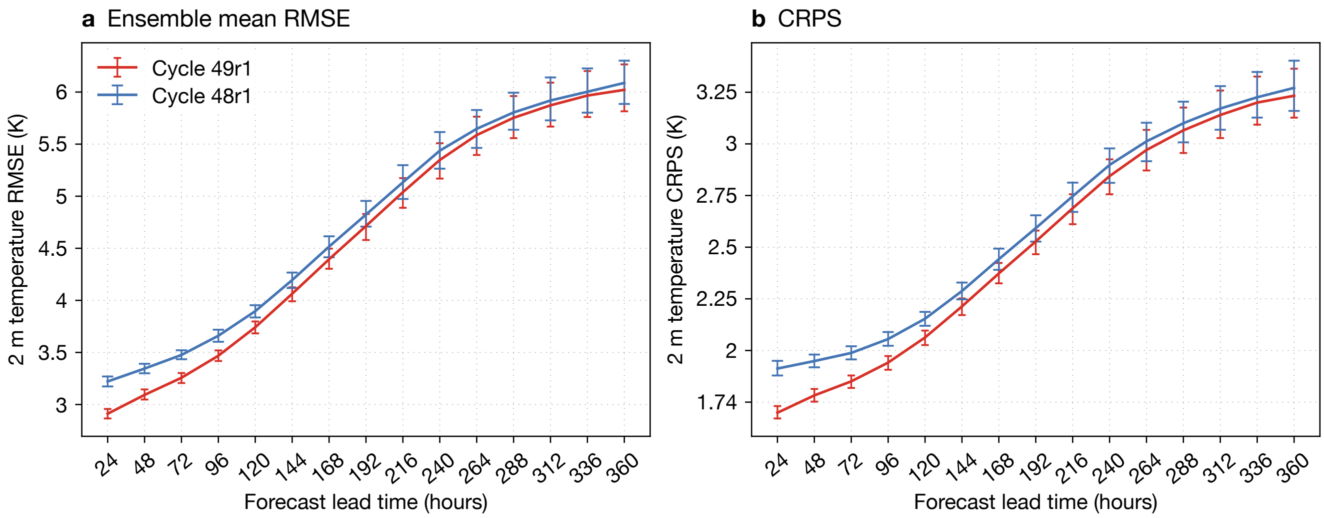

FIGURE 4 The charts show (a) the root-mean-square error (RMSE) of the ensemble mean and (b) the continuous ranked probability score (CRPS) for 2 m temperature from Cycle 48r1 (blue) and Cycle 49r1 (red), verified against observations and averaged over the northern hemisphere. All scores are calculated using 50 perturbed members from 75 forecasts initialised daily between 1 December 2022 and 13 February 2023.

- An upgrade to version 13.2 of the Radiative Transfer for TOVS (RTTOV) model and other updates in the radiance observation operator.
- Activation of Variational Quality Control (VarQC) in the first 4D-Var minimisation.
- Activation of balance constraints above 20 km, which allows stratospheric sounding instruments to generate geostrophically balanced increments from the start of the assimilation window.

Impact on medium-range and sub-seasonal forecasts

Cycle 49r1 substantially improves short- and medium-range forecasts of 2 m temperature and 10 m wind speeds. The largest impacts on 2 m temperature forecasts are for the winter months in the northern hemisphere (Figure 4), where the Continuous Ranked Probability Score (CRPS) is improved by 11% at day 1 and 2% at day 10. In short-range forecasts, the biggest improvements are over Asia and Canada (Figure 5). These improvements reflect the combined impact of many contributions, including the assimilation of 2 m temperature observation data, upgrades to 4D-Var and land-surface data assimilation methodology, and improvements to the IFS land surface model (Ingleby et al., 2024).

Forecasts of 10 m wind speed are improved throughout the year. The largest impacts are for the winter months of the northern hemisphere (Figure 6), where the CRPS is improved by 12% at day 1 and 6% at day 10. The major contributors to improved ensemble forecasts of 10 m wind speed are the increased spread in the boundary layer associated with the switch to the SPP scheme for model uncertainty; improvements to the diagnostic 10 m wind calculation; and the combination of updates to the land-surface model.

Figures 7 and 8 summarise the impact of Cycle 49r1 relative to Cycle 48r1. They use several metrics of deterministic and probabilistic forecast skill for a range of variables and levels across medium-range lead times. Cycle 49r1 has an overall positive impact on both deterministic and ensemble forecast skill scores. For example, CRPS and anomaly correlations are both improved in 73% of comparisons shown in Figure 8. The impact in the tropics is more mixed, especially for verification against analyses. Some of the negative impacts are a consequence of interactions between the new EDA configuration and other contributions to Cycle 49r1. In particular, the increased EDA resolution and the introduction of SPP both have a non-uniform impact on the EDA forecast spread (see Figure 1). This spread informs background errors of the 4D-Var system and provides perturbations for the initialisation of the ensemble forecasts. The final configuration of Cycle 49r1 reflects a compromise that balances evidence from observation-based verification, analysis-based verification, and observation–background/observation–analysis statistics. This involved a tuning process that uniformly reduced the background errors provided to the 4D-Var system by 16% and separately a 15% reduction of the amplitude of the singular vector perturbations added to ensemble initial conditions.

Despite significant changes to the EDA and model uncertainty representations, the scale-dependent growth of forecast uncertainty is, to first order, very similar in Cycle 49r1 and Cycle 48r1 (Figure 2). Nevertheless, there is a systematic reduction in synoptic-scale ensemble variance at medium-range lead times that provides a

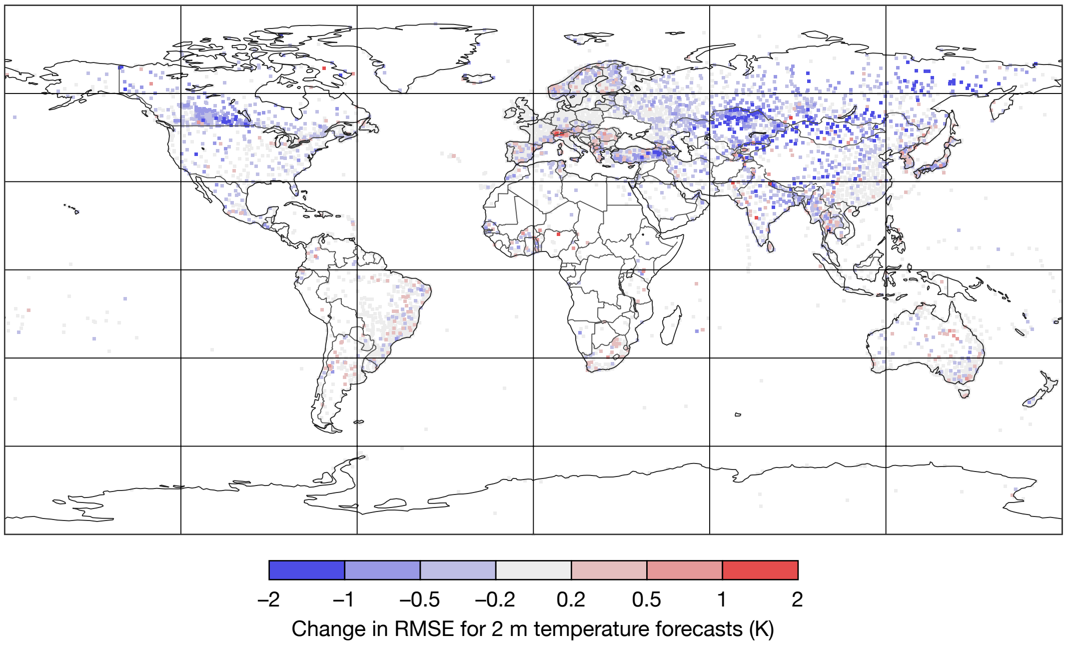

FIGURE 5 Change in root-mean-square error (RMSE) for 24-hour deterministic forecasts of 2 m temperature at 00 and 12 UTC between December 2022 and February 2023 inclusive. Blue colours indicate improvements in IFS Cycle 49r1 with respect to IFS Cycle 48r1.

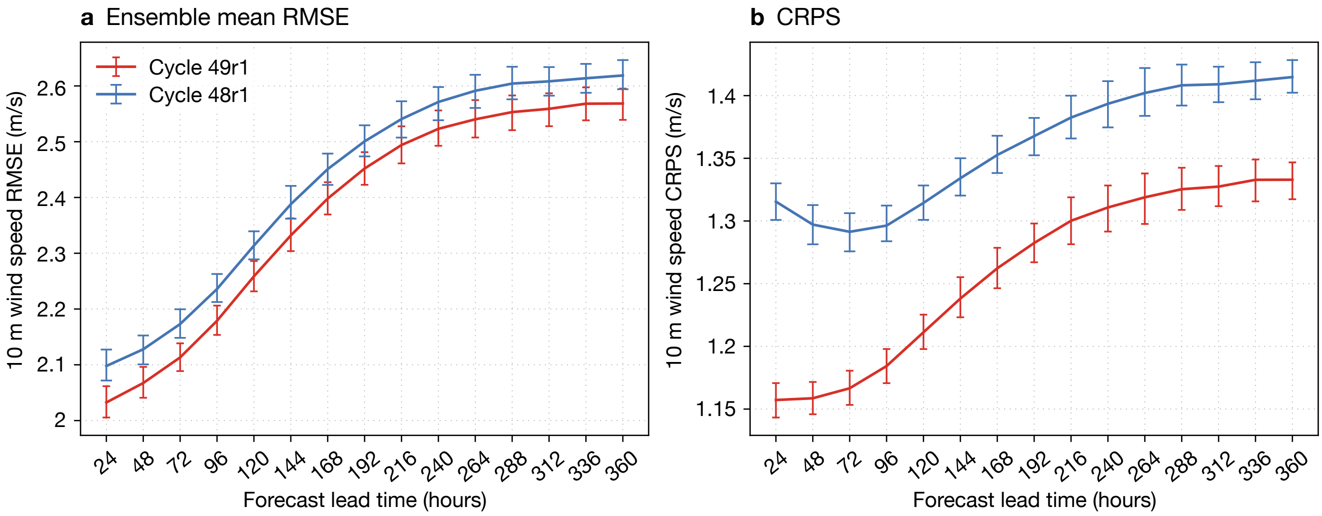

FIGURE 6 As Figure 4, but for 10 m wind speeds.

10% improvement to the over-dispersion in the storm tracks in Cycle 48r1. This improves both ensemble reliability and sharpness (not shown).

Tropical cyclone (TC) intensity and position errors are generally similar in Cycle 49r1 and Cycle 48r1 (Figure 9). The most important change is the significant increase in the ensemble spread of TC intensity, especially at earlier lead times, which is mainly a consequence of the new EDA configuration. This change improves the reliability of TC intensity forecasts, as seen in ensemble spread and ensemble mean root-mean-square error (RMSE) moving closer to each other. In addition, the switch from SPPT to SPP results in a systematic increase in the frequency of tropical cyclones. This increase is likely to be detrimental for weaker systems, which were already overestimated compared to observations, but beneficial for deeper systems, which remain underestimated in Cycle 49r1.

An undesirable feature of TC forecasts in Cycle 49r1 is the increased frequency of unrealistic TC structures that are not axisymmetric and associated very strong winds at initialisation time in perturbed ensemble members. These artefacts are a property of the ensemble initial conditions rather than the forecast model and dissipate within the first 12 hours of the forecast without significantly changing medium-range TC forecast quality. This issue is a consequence of the way ensemble initial conditions are derived from perturbations taken from a more realistic EDA and then re-centred around a reference analysis. It was previously described by Lang et al. (2015) and will be addressed with a revision of the ensemble initial perturbations methodology in a forthcoming IFS upgrade.

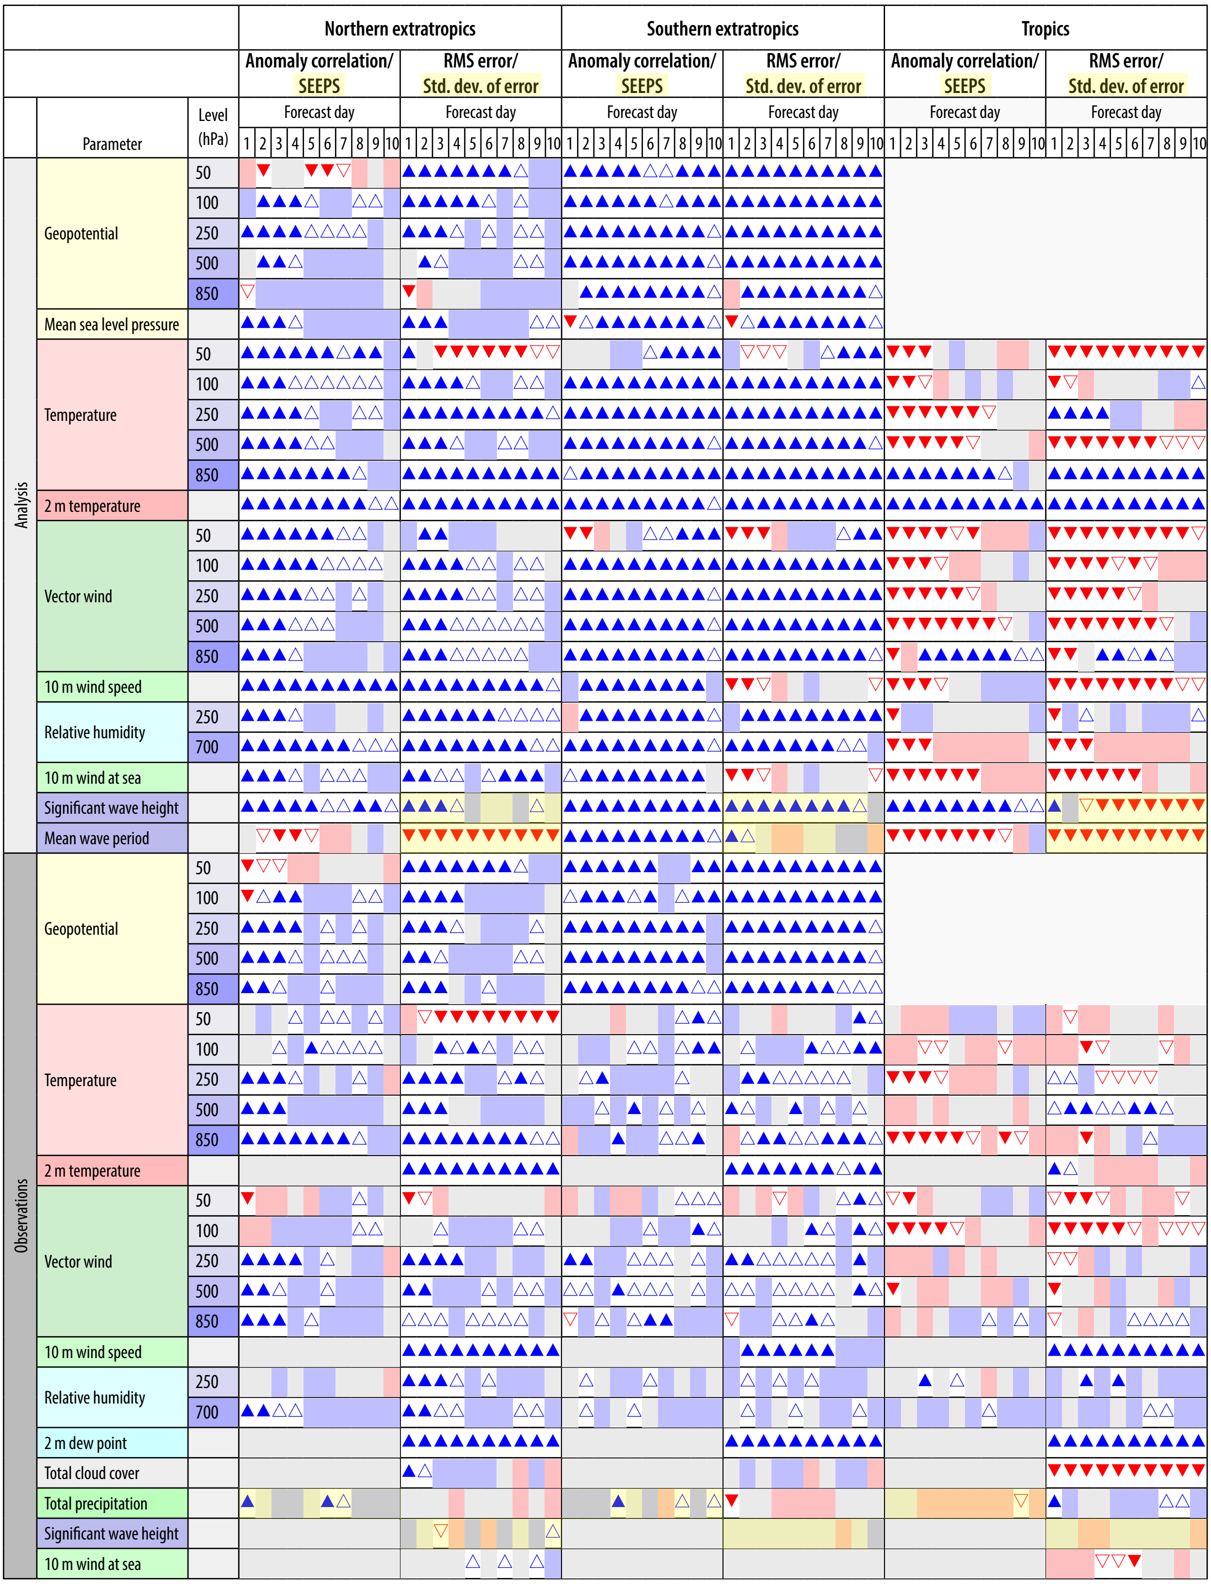

Symbol legend: for a given forecast step...

▲ 49r1 better than 48r1 statistically significant with 99.7% confidence

△ 49r1 better than 48r1 statistically significant with 95% confidence

49r1 better than 48r1 statistically significant with 68% confidence

no significant difference between 48r1 and 49r1

49r1 worse than 48r1 statistically significant with 68% confidence

▽ 49r1 worse than 48r1 statistically significant with 95% confidence

▼ 49r1 worse than 48r1 statistically significant with 99.7% confidence

FIGURE 7 Summary scorecard comparing the difference between control forecasts from IFS Cycle 49r1 and IFS Cycle 48r1 using anomaly correlation coefficients and the root-mean-square error (RMSE). Note that total precipitation is evaluated using the Stable Equitable Error in Probability Space (SEEPS) score rather than correlation, and wave parameters are evaluated using the standard deviation of errors rather than RMSE. Blue colours indicate improvements in IFS Cycle 49r1 with respect to IFS Cycle 48r1. Scores are calculated from more than 1,000 forecasts initialised at 00 and 12 UTC between 1 June 2022 and 8 August 2024.

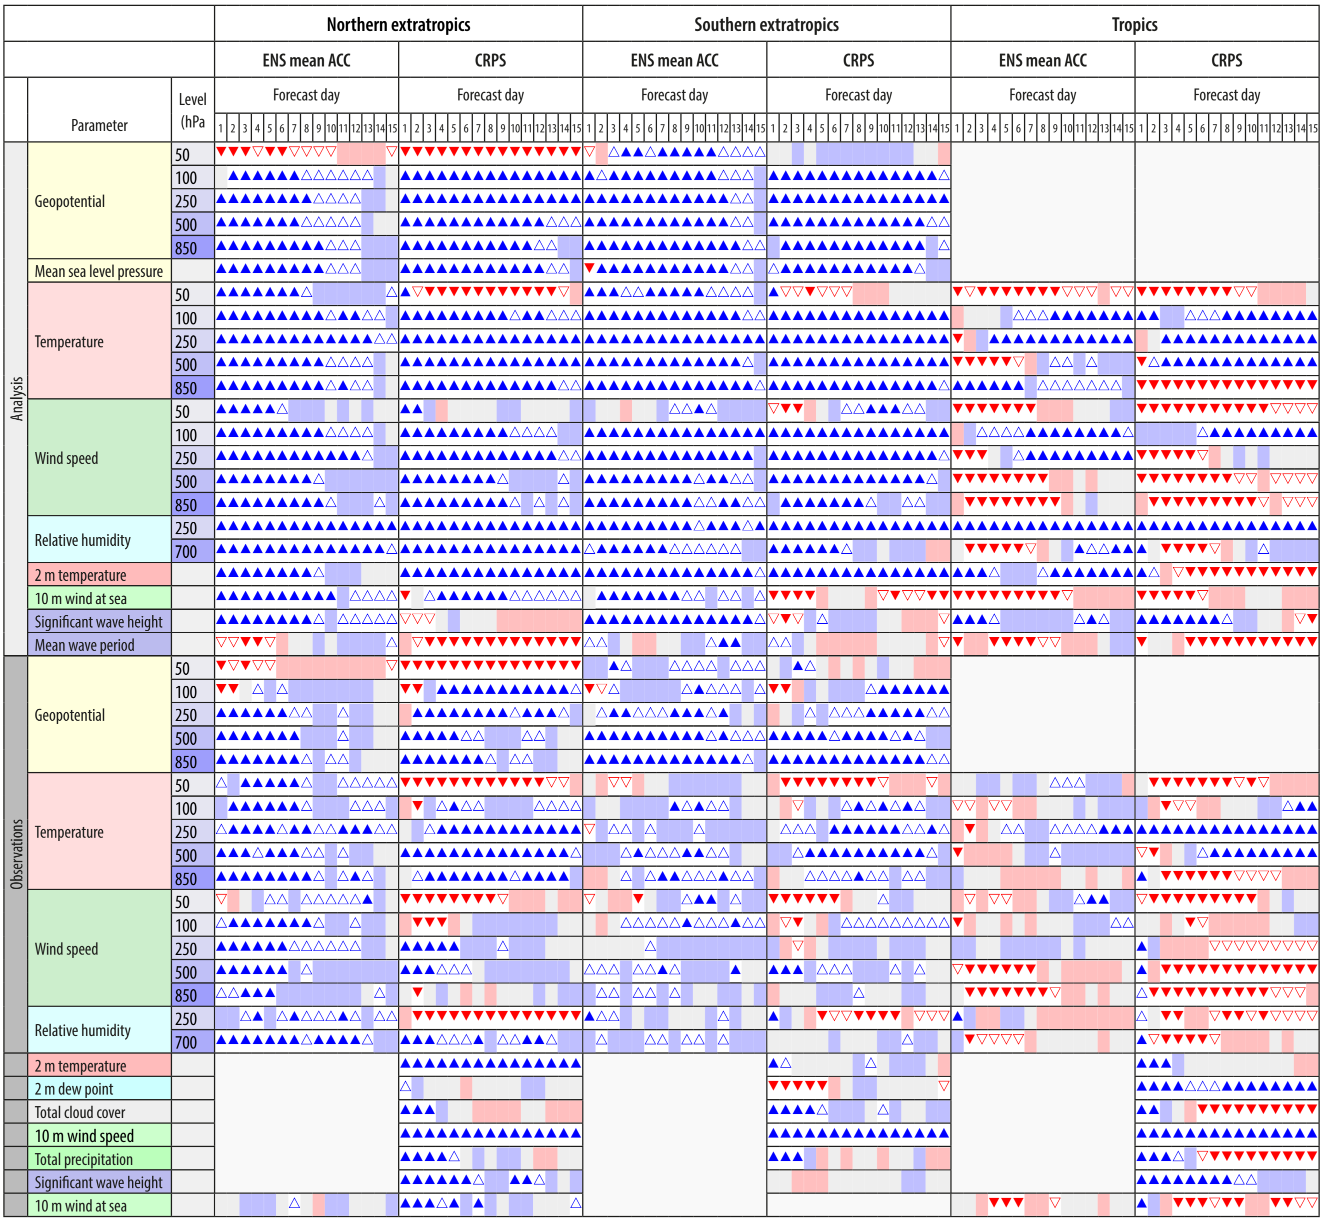

Symbol legend: for a given forecast step...

- ▲ 49r1 better than 48r1 statistically significant with 99.7% confidence
- △ 49r1 better than 48r1 statistically significant with 95% confidence
- 49r1 better than 48r1 statistically significant with 68% confidence
- no significant difference between 48r1 and 49r1
- 49r1 worse than 48r1 statistically significant with 68% confidence
- ▽ 49r1 worse than 48r1 statistically significant with 95% confidence
- ▼ 49r1 worse than 48r1 statistically significant with 99.7% confidence

FIGURE 8 Summary scorecard comparing the difference between ensemble forecasts with IFS Cycle 49r1 and IFS Cycle 48r1, using the anomaly correlation coefficient (ACC) of the ensemble mean and the continuous ranked probability score (CRPS). Blue colours indicate improvements in IFS Cycle 49r1 with respect to IFS Cycle 48r1. Scores are calculated using 50 perturbed ensemble members from more than 400 forecasts initialised at 00 UTC between 1 June 2022 and 8 August 2024.

The ECMWF sub-seasonal forecasting system, which provides an overview of potential weather conditions up to 46 days ahead, is also updated in Cycle 49r1. The most robust impacts on weekly mean forecast anomalies are small but statistically robust changes in ensemble spread, which are driven by the switch from SPPT to SPP. These changes are most evident in the tropics, where ensemble spread in the free atmosphere is reduced by several per cent, which represents a slight improvement in ensemble reliability relative to Cycle 48r1. Despite these changes to ensemble spread, deterministic and probabilistic weekly mean anomaly scores are generally very similar in Cycle 49r1 and Cycle 48r1.

Despite the limited impact on weekly mean scores aggregated over large regions, Cycle 49r1 improves the bivariate correlation skill of Madden–Julian Oscillation (MJO) forecasts at lead times greater than 15 days. These changes are associated with improved MJO

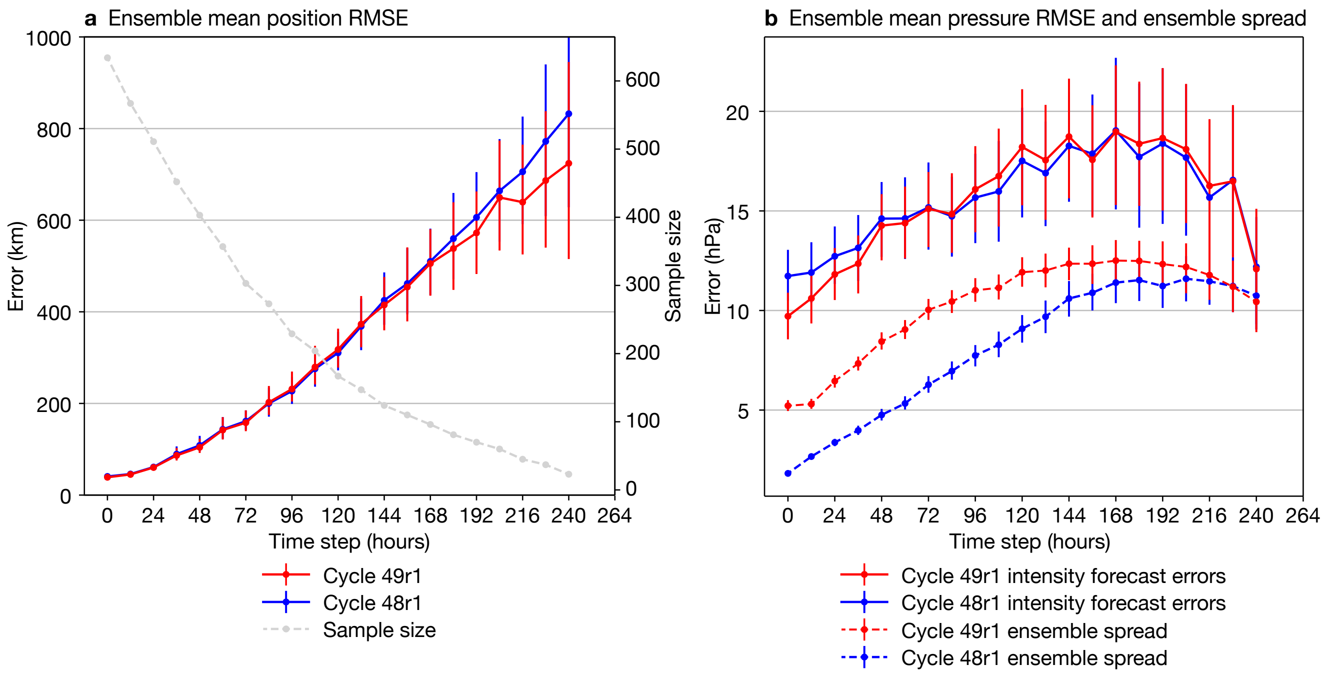

FIGURE 9 Tropical cyclone position and intensity forecast errors in IFS Cycle 49r1 compared to IFS Cycle 48r1, showing (a) ensemble mean position forecast errors of Cycle 49r1 and Cycle 48r1 and (b) intensity (central pressure) forecast errors (solid lines) and ensemble spread (dashed lines) of Cycle 49r1 and Cycle 48r1. The vertical error bars represent 95% confidence intervals. Tropical cyclone scores are calculated using 50 perturbed members from forecasts initialised daily between 3 June 2022 and 29 March 2024.

reliability at sub-seasonal lead times, which is a consequence of the switch from SPPT to SPP.

System configuration changes and updated products

Cycle 49r1 introduces new and revised diagnostic products, including a harmonisation of convection-related parameters. All relevant applications and products in Cycle 49r1 use consistent definitions of Most Unstable Convective Available Potential Energy (MUCAPE) and Most Unstable Convective Inhibition (MUCIN). The original parameters for Convective Available Potential Energy (CAPE) and Convective Inhibition (CIN) are discontinued. Other product changes include new graphical products for true-colour simulated satellite imagery, new parameters related to heat stress (Wet Bulb Globe Temperature, Heat Index, Humidex, Wind Chill Temperature, Universal Thermal Climate Index), and an extension of probabilistic clear-air turbulence products to include contributions from non-orographic gravity-wave dissipation outside of convection areas.

Cycle 49r1 also introduces several major changes to the configuration of operational systems:

- The wave model is now run with the same native grid as the atmospheric model in all forecast systems, and the number of frequencies in wave spectra output is reduced from 36 to 29. For example, for medium-range ensemble forecasts (ENS), the native grid of the wave model will change from a 14 km reduced latitude–longitude grid to use the same TCo1279 (~9 km) grid as the atmosphere.
- Cycle 48r1 increased the horizontal resolution of ENS from TCo639 to TCo1279. One consequence of this upgrade was to make the unperturbed ENS control forecast scientifically equivalent (though not computationally identical) to the high-resolution deterministic forecast (HRES). Although nearly identical, HRES and the ENS control forecast retained some differences in lead time (10 days vs 15 days), output frequencies, and dissemination times. In Cycle 49r1, HRES and the ENS control forecast become scientifically and computationally identical and both are run for 15 days at 00 UTC and 12 UTC. The superfluous ENS control forecast will be stopped in a future IFS upgrade, and the data stream currently known as HRES will become known as the 'control' forecast. The retitled control forecast will be available on the same schedule as the current HRES (i.e. earlier than the perturbed ENS forecast members).
- Cycle 49r1 introduces new configurations for medium-range and sub-seasonal retrospective ensemble forecasts (also known as 're-forecasts' or 'hindcasts'). The ensemble size of re-forecasts (ten perturbed members and one control member) is unchanged from Cycle 48r1. However, the frequency of re-forecasts is changed for both systems. In Cycle 48r1, both medium-range and

sub-seasonal re-forecasts were run every Monday and Thursday for the previous 20 years. In Cycle 49r1, sub-seasonal re-forecasts are run every odd day of the month over the previous 20 years (i.e. 1st, 3rd, 5th…, excluding 29th February) and medium-range re-forecasts are run every other odd day of the month over the previous 20 years (i.e. 1st, 5th, 9th…, excluding 29th February). There are several advantages to these new re-forecast configurations: (i) The increased frequency of sub-seasonal re-forecasts benefits skill assessment and the calibration of real-time forecasts. (ii) The use of fixed re-forecast dates enables direct comparisons with seasonal re-forecasts and between sub-seasonal re-forecasts produced in different years. (iii) The common dates for medium-range and sub-seasonal re-forecasts facilitate resolution sensitivity studies and provide opportunities for the generation of calibrated dual-resolution ensemble products.

Summary and outlook

IFS Cycle 49r1 brings substantial changes to data assimilation methodology, the use of observations, and the underlying IFS forecast model. These changes have an overall positive impact on both deterministic and ensemble forecasts and significantly improve 2 m temperature and 10 m wind forecasts, particularly for the winter months in the northern hemisphere.

The improvements in near-surface weather conditions reflect the combined impact of many contributions, but especially the assimilation of 2 m temperature observation data; upgrades to 4D-Var and land-surface data assimilation methodology; improvements to the IFS land surface model; the switch to the SPP scheme for model uncertainty; and improvements to the diagnostic 10 m wind calculation. The assimilation of microwave imaging radiances over sea-ice surfaces is an important step towards Earth system data assimilation, in which the different components are integrated in a coupled data assimilation framework (de Rosnay et al., 2022). Other major contributions to Cycle 49r1 include the increased resolution and soft re-centring of the EDA; the increased usage of satellite observations; the increased resolution of the ocean wave model; and changes to the frequency of medium-range and sub-seasonal re-forecasts. Looking ahead, the next update of the IFS (Cycle 49r2) will build upon Cycle 49r1 and serve as the foundation for the next ECMWF reanalysis (ERA6) and seasonal forecast system (SEAS6). Cycle 49r2 will not be implemented operationally for medium-range or sub-seasonal forecasts, but all developments for ERA6 and SEAS6 will be included as part of the next operational IFS upgrade (Cycle 50r1).

Further reading

Bidlot, **J.-R.** & **P. Janssen**, 2024: Ocean-wave-related changes in the next model upgrade, *ECMWF Newsletter No. 179*, **18–25**. *https://doi.org/10.21957/jb9py41a6f*

Boussetta, **S.** & **G. Balsamo**, 2021: CONFESS Vegetation dataset of Land Use/Land Cover and Leaf Area Index. CONFESS D1.1. *https://confess-h2020.eu/wp-content/uploads/2021/08/confess-d1-1-v1-0-.pdf*

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. *QJRMS*, **148**(747), 2672–2702. *https://doi.org/10.1002/qj.4330*

Geer, **A.J.**, 2023: Combining machine learning and data assimilation to estimate sea ice concentration, *ECMWF Newsletter* **No. 177**, 14–21. *https://doi.org/10.21957/agh93vs26*

Geer, **A.J.**, 2024: Joint estimation of sea ice and atmospheric state from microwave imagers in operational weather forecasting. *QJRMS*. *https://doi.org/10.1002/qj.4797*

Hólm, **E.**, **M. Bonavita**, & **S. Lang**, 2022: Soft re-centring Ensemble of Data Assimilations. *ECMWF Newsletter* **No. 171**. *https://www.ecmwf.int/en/newsletter/171/news/soft-re-centring-ensemble-data-assimilations*

Ingleby, **B.**, **G. Arduini**, **G. Balsamo**, **S. Boussetta**, **K. Ochi**, **E. Pinnington et al.**, 2024: Improved two-metre temperature forecasts in the 2024 upgrade. *ECMWF Newsletter* **No. 178**, 24–29. *https://doi.org/10.21957/bi49s20qa8*

Lang, **S.T.**, **M. Bonavita** & **M. Leutbecher**, 2015: On the impact of re-centring initial conditions for ensemble forecasts. *QJRMS*, **141**(692), 2571–2581. *https://doi.org/10.1002/qj.2543*

McNorton, **J.** & **G. Balsamo**, 2023: An urban scheme for the IFS, *ECMWF Newsletter* **No. 175**, 18–19. *https://www.ecmwf.int/en/newsletter/175/news/urban-scheme-ifs*

Improving the physical consistency of ensemble forecasts by using SPP in the IFS

Martin Leutbecher, Simon Lang, Sarah-Jane Lock, Christopher D. Roberts, Aristofanis Tsiringakis

Ensemble forecasts need to account for uncertainties in both initial conditions and the forecast model. Since 1998, the latter uncertainties have been represented in ECMWF's Integrated Forecasting System (IFS) via the Stochastically Perturbed Parametrization Tendency scheme (SPPT; Buizza et al., 1999). This scheme is also referred to as 'stochastic physics'. It has been revised several times. SPPT has played an important role through increasing the ensemble spread and boosting the probabilistic skill of ECMWF ensemble forecasts over the past 25 years (see Lock et al., 2019, for details of the operational SPPT configuration). In IFS Cycle 49r1, which will be implemented in November 2024, SPPT will be replaced by the Stochastically Perturbed Parametrizations (SPP) scheme in all ensemble applications. SPP has been developed over several years (Ollinaho et al., 2017; Lang et al., 2021). It represents model uncertainties closer to the sources of errors. The remainder of the article explains the motivation for this revision and how the new scheme works, and it sets out the impacts expected from the revision of the model uncertainty representation.

Motivation

While SPPT has been a success story in terms of its impact on the skill of ensemble forecasts, awareness of some drawbacks has increased over the years. This fuelled interest in developing a representation of uncertainties that has a comparable positive impact on ensemble skill to SPPT, but which can bring additional benefits. In particular, we seek a perturbation method that maintains the conservation properties of the unperturbed model, by ensuring that fluxes at the surface and top of the atmosphere respond consistently to the perturbations within the atmospheric column. Furthermore, by applying the perturbations to individual physical processes, we can introduce a representation of errors in, for example, the shape of a heating profile, and remove the need to taper perturbations near the surface.

Methodology

SPP is a stochastic representation of uncertainties which targets uncertain elements within the parametrizations of individual physical processes. The elements are identified by scientists who develop the parametrizations and have in-depth knowledge of the sensitivities of their schemes to specific choices that need to be made to constrain the parametrization. The version of SPP in Cycle 49r1 has 27 perturbed elements in order to represent the dominating uncertainties in the parametrizations of convection, large-scale cloud, radiation, surface exchanges, vertical turbulent mixing, and orographic gravity wave drag (Table 1). Each perturbation element has its own probability distribution, which is anchored to the unperturbed values used in the deterministic model. Figure 1 shows the distributions sampled by SPP for the elements perturbing the standard deviation of the subgrid orography (HSDT) and the ocean cold skin temperature parametrization (COLDSKIN). The distributions of these elements are limited to a finite range to ensure numerical stability. The distributions are sampled via evolving random fields with prescribed time and space scales (see Lang et al., 2021, for details). Each perturbation element has its own independent random field. Figure 2 displays

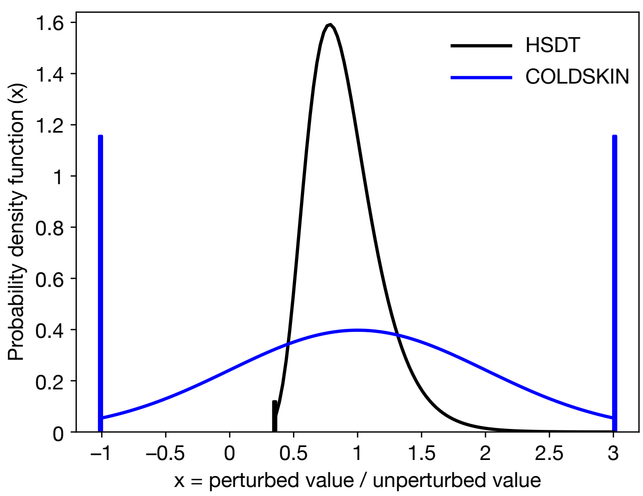

FIGURE 1 Distributions sampled by SPP for the perturbation elements HSDT and COLDSKIN. The continuous distributions of these elements are censored, on the left at 0.3619 (HSDT) and −1.0 (COLDSKIN) and on the right at 3.0, in order to guarantee numerical stability. For visualisation, the finite probabilities of x being at an endpoint of the range of the probability density function are converted into densities by assuming a constant density over an interval of $\Delta x = 0.02$. The mean of both distributions is 1.

doi: 10.21957/mlz238dk1p

Surface fluxes, turbulent mixing and subgrid orography	
CFM	Transfer coefficient for momentum
RKAP	Surface flux uncertainties via von Kármán constant
TOFDC	Turbulent orographic form drag
HSDT	Standard deviation of subgrid orography
VDEXC_LEN	Mixing length-scale stable boundary layer
VDSST	Sea-surface temperature (SST) used in calculation of surface fluxes
COLDSKIN	Cold skin temperature parametrization used for surface fluxes
Convection	
ENTRORG	Entrainment rate
ENTSHALP	Shallow entrainment rate
DETRPEN	Detrainment rate for penetrative convection
RPRCON	Conversion coefficient cloud to rain
CUDU/CUDV	Deep convective momentum transport
CUDUS/CUDVS	Shallow convective momentum transport
RTAU	Adjustment timescale in Convective Available Potential Energy (CAPE) closure
ENTSTPC1	Shallow convection test parcel entrainment
Cloud and large-scale precipitation	
RAMID	Relative humidity threshold stratiform condensation
RCLDIFF	Diffusion for evaporation of cloud at subgrid cloud edges
RLCRITSNOW	Cloud ice threshold for autoconversion to snow
RAINEVAP	Rain evaporation rate
SNOWSUBLIM	Snow sublimation rate
QSATVERVEL	Vertical velocity for adiabatic temperature change in saturation adjustment
FALLSPEED	Hydrometeor terminal fall speeds
Radiation	
ZDECORR	Cloud vertical decorrelation height
ZSIGQCW	Fractional standard deviation of horizontal distribution of water content
ZRADEFF	Effective radius of cloud water and ice
ZHS_VDAERO	Scale height of aerosol normal vertical distribution
DELTA_AERO	Optical thickness of aerosol

TABLE 1 Overview of the active perturbation elements in SPP.

the evolution of the HSDT and COLDSKIN random fields over three days for one ensemble member. The gridpoint values of the random fields are transformed such that they sample the specified probability distributions. The random fields have two components with spatial decorrelation scales of 1,000 km and 3,000 km and corresponding decorrelation time scales of 3 days and 30 days, respectively. The development of the scheme required a number of iterations to constrain the distributions of all elements in order to generate the right level of variance in ensembles at all lead times, using one unified configuration.

SPP inherits the local conservation properties from the deterministic parametrizations. The fluxes at the surface and at the top of the atmosphere are computed using the perturbed elements, resulting in perturbed fluxes that remain consistent with the tendencies in the atmospheric column. Due to the use of multiple perturbation elements, the structure of the tendencies can be altered by the perturbations and can represent, for instance, uncertainties in the shape of heating profiles.

Impacts

Now, we provide an overview of the impacts of SPP on local conservation and on the different ranges of

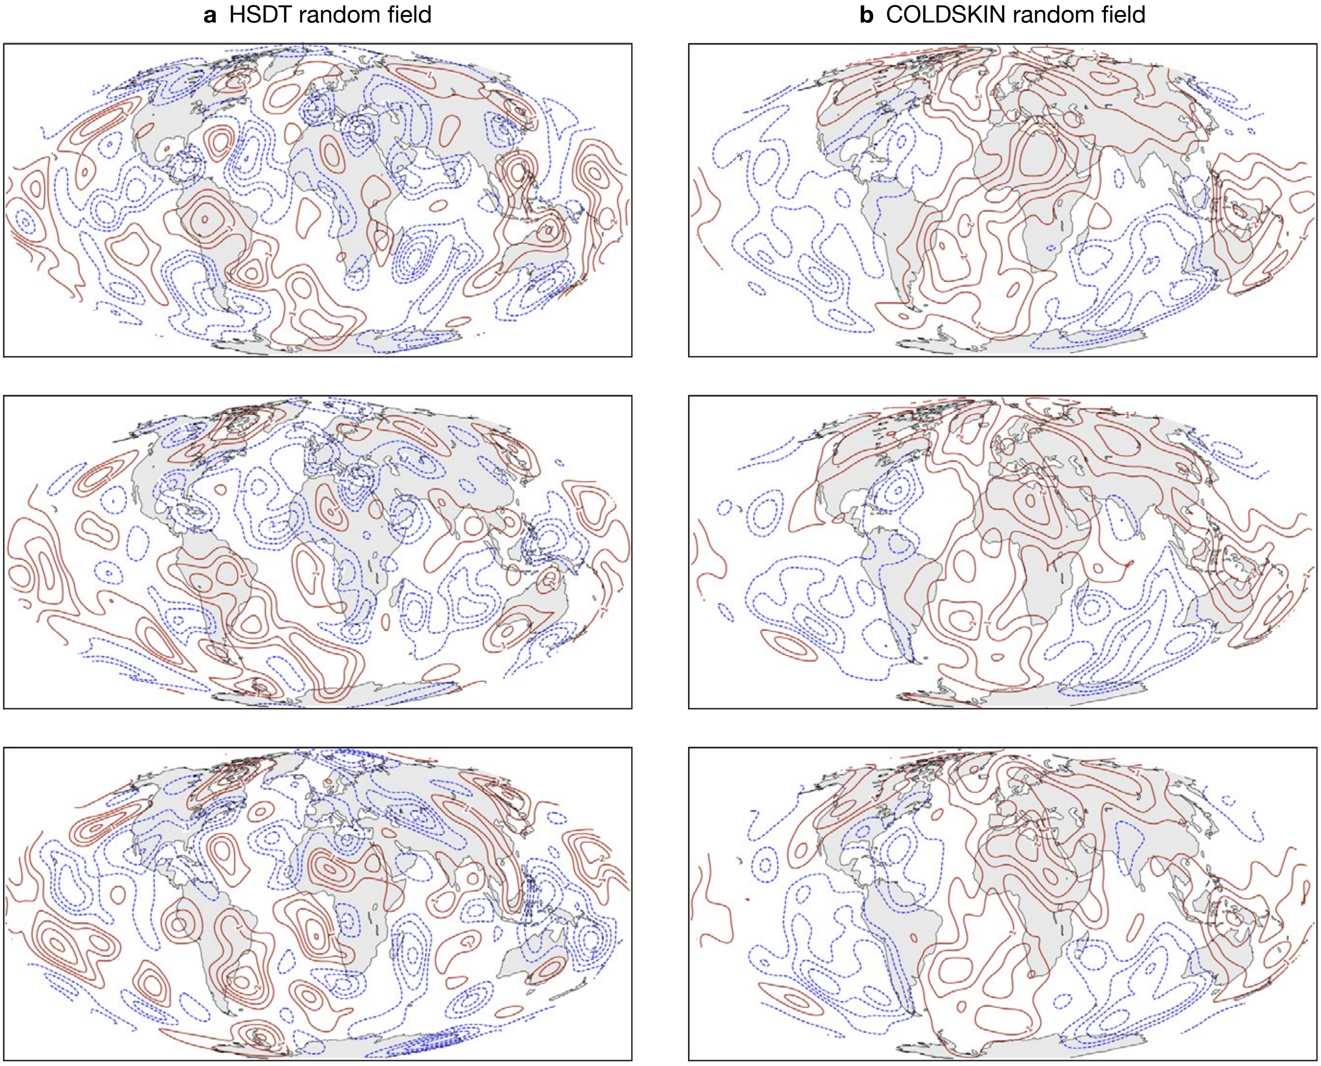

FIGURE 2 Random fields used by SPP in member 1 on 1, 2 and 3 January 2024 (from top to bottom) at 00 UTC for the perturbation elements (a) HSDT and (b) COLDSKIN. The latter has larger spatial and temporal decorrelation scales than the former. The contour interval is 0.5, with values ≥ 0.5 in solid red contours and values ≤ –0.5 in blue dashed contours.

ECMWF's forecasts, from the medium range (up to 15 days) via the sub-seasonal range (up to 46 days) to the seasonal range (up to 13 months), as well as in km-scale ensemble forecasts explored in the EU's Destination Earth (DestinE) initiative.

Local conservation

The parametrizations of physical processes aim to conserve moisture and enthalpy (moist static energy). Column integrals of the tendencies together with the fluxes at the surface and the top of the atmosphere form a nearly closed budget when the semi-Lagrangian averaging of the physics tendencies is deactivated (see Part IV of the IFS documentation for details: *https://www.ecmwf.int/en/publications/ifs-documentation*). Figure 3 (a,b) shows that the residual term in the budget for moisture is very small compared to the precipitation flux in the unperturbed control forecast. The forecast perturbed with SPP achieves a similar level of conservation as the control forecast, while the forecast perturbed with SPPT has a residual that is locally nearly of the same magnitude as the precipitation flux itself (Figure 3 c,d). The level of conservation or lack of conservation shown in the figure for forecast lead times of 45 to 48 hours is representative for all forecast lead times, and it is also representative for other ensemble members or start dates.

For the local budget of enthalpy, the level of conservation of forecasts perturbed with SPP is again similar to the level of conservation in the unperturbed control forecast, while the forecasts using SPPT show large residuals.

Medium-range forecasts

Tropospheric scores in the extratropics are overall improved with SPP, e.g. geopotential (see Figure 4a). Upper-air scores in the tropics are more mixed, with improvements at some levels and degradations at others. For example, there is a strong positive impact for temperature in the tropics at 500 hPa and 250 hPa. Here, SPP shows improvements of around 10%.

Conversely, wind scores in the tropics are degraded, between 1 and 2% for some levels. In general, upper-air spread (ensemble standard deviation) is somewhat decreased, around 1 to 2% in the extratropics. Forecast skill of surface variables like 2-metre temperature and 10-metre wind speed is improved (Figure 4b and c). While 2-metre temperature spread is quite similar, spread of 10-metre wind is increased, by around 25% at a lead time of 48 hours.

SPP results in a noticeable increase in the occurrence frequency of tropical cyclones in the ensemble. The impact on tropical cyclone frequencies originates mainly from the SPP perturbations in the convection parametrization, specifically the perturbations to the deep convective momentum transport (CUDUDV). We have compared the occurrence frequencies to tropical cyclone data from the international best track archive for climate stewardship (IBTrACS). For wind speeds exceeding 20 m/s, the increase in tropical cyclone frequency appears to bring the forecasts closer to the observed frequencies. We believe that for weaker systems, with wind speeds less than 20 m/s, the observed frequencies in IBTraCS may be underestimated.

Sub-seasonal forecasts

In broad terms, the response in sub-seasonal forecasts to changing from SPPT to SPP is similar to that in the medium range. There are small changes in spread (typically of a few percent) throughout forecast lead times. The changes are characterised by reduced spread in upper-air fields and increased spread in surface fields. They are most pronounced in the tropics. The reduced spread with SPP tends to translate into improved spread–error ratios, since using SPPT in the latest ensemble configuration generates somewhat over-dispersive forecasts of some upper-air fields.

A key component of predictability in longer-range forecasts derives from the Madden–Julian Oscillation (MJO, see e.g. Lin et al., 2009). The MJO is a large-scale weather pattern that occurs over the Indian and Pacific Oceans and is associated with global teleconnection patterns. Figure 5 illustrates ensemble spread and root-mean-square error (RMSE) of the ensemble mean

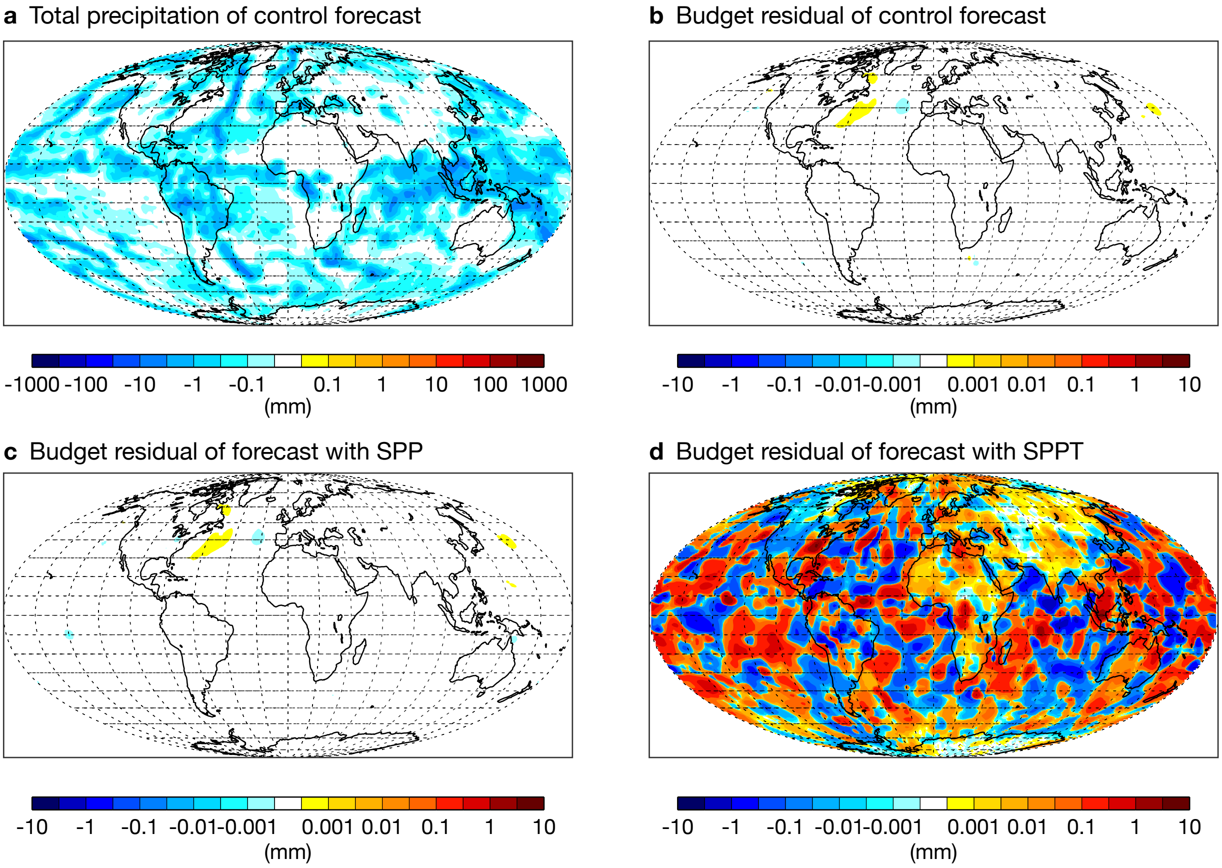

FIGURE 3 Moisture budget terms accumulated during a forecast lead time of 45–48 hours, showing (a) the precipitation flux in the control forecast, and budget residuals for (b) the control forecast, (c) an example SPP perturbed forecast, (d) an example SPPT perturbed forecast. The contour interval values in plots (b)–(d) are 100 times smaller than in (a). The data is from an ensemble forecast at a resolution of 25 km (TCo399) and with 137 vertical levels with Cycle 49r1, starting on 1 December 2022, 00 UTC, with only either SPP or SPPT perturbations in (c) and (d) and no initial condition perturbations.

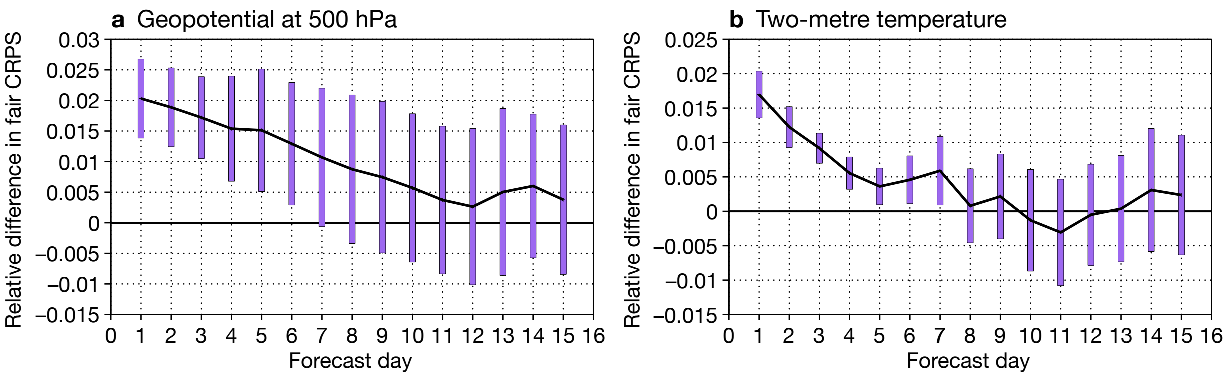

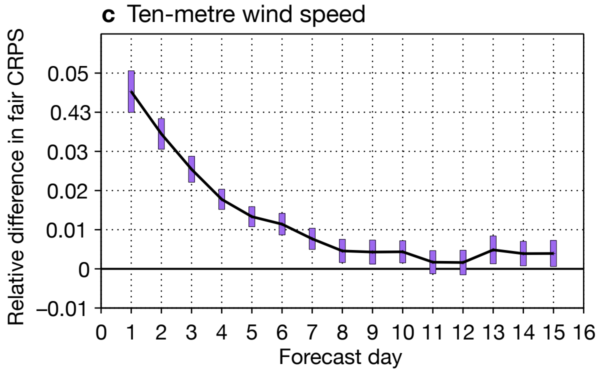

FIGURE 4 Relative differences in fair CRPS (continuous ranked probability score) of an ensemble using SPP and an ensemble using SPPT in the northern extratropics, for (a) geopotential at 500 hPa verified against analyses, (b) 2-metre temperature verified against SYNOP weather station observations, and (c) 10-metre wind speed verified against SYNOP weather station observations. Positive values indicate higher forecast skill for the ensemble using SPP. Shown are combined scores for northern winter 2021/2022 and summer 2022 (282 start dates). The experiments use a resolution of 9 km (TCo1279) and eight perturbed members initialised from operational initial conditions.
The vertical bars show 95% confidence intervals for the score differences.

for the real-time multivariate MJO (RMM) index (following Wheeler and Hendon, 2004) from two sets of re-forecasts: one using SPPT and one SPP. The RMM index is constructed from a combination of tropical upper-air fields: zonal (east–west) winds at 850 hPa and 200 hPa, and outgoing longwave radiation. The goal is a good match between ensemble spread and RMSE. While the forecast skill (represented by the RMSE of the ensemble mean) for the two model uncertainty representations remains very similar, the ensemble spread with SPP is significantly reduced from that with SPPT, and it matches the error much more closely.

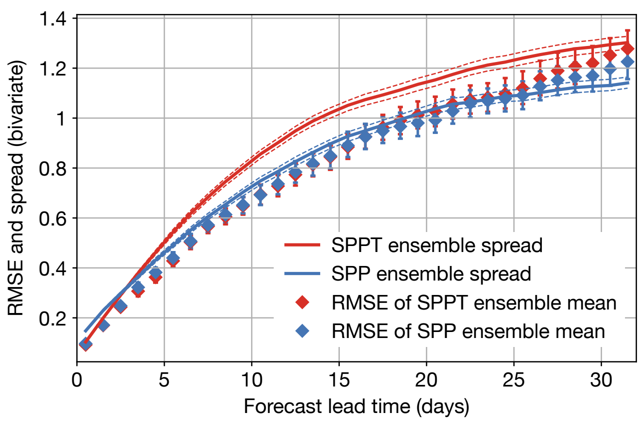

FIGURE 5 Ensemble spread (solid lines) and root-mean-square error (RMSE) of the ensemble mean (diamonds) for the bivariate RMM index for experiments with model uncertainty represented with SPPT (red) and SPP (blue). Confidence intervals (95%) are indicated by the dashed lines (for spread) and bars (for error). Forecasts are started on the first of each month from 1 January 1989 to 1 December 2016.

Seasonal forecasts

Although the change from SPPT to SPP will not be applied to operational seasonal forecasts until the next upgrade of the seasonal system (to SEAS6), the impact of the change has been tested in seasonal configurations.

Switching from SPPT to SPP tends to reduce ensemble spread. For some key forecast variables, that reduction in spread addresses an over-dispersive signal from SPPT. This translates into improved spread–error ratios. For example, for equatorial Pacific regions that are important for El Niño–Southern Oscillation (ENSO) predictions, the spread–error ratios for sea-surface temperatures (SSTs) are mostly improved by the reduced spread with SPP.

The change of the model uncertainty scheme also impacts some forecast biases over these longer lead times. For example, for SSTs in the eastern equatorial Pacific, SPP leads to some warming, while SPPT tends to cool the region. The consequence of the bias changes is complicated, since for both schemes there are regions and seasons for which the change helps (or hinders) by counteracting (or strengthening) the prevailing forecast bias. However, the differences are small relative to the underlying biases.

Kilometre-scale ensemble forecasts

Replacing SPPT with SPP brings benefits in km-scale ensemble forecasts (at a grid spacing of 4.4 km) similar to those in the medium-range. In the tropics, fair CRPS scores for wind speed are slightly degraded for 500 and 850 hPa levels, and for temperature at 850 hPa.

However, surface fair CRPS scores (evaluated against observations) improve between 2 and 5% for 2-metre temperature and 10-metre wind speed. For several extreme weather events (e.g. tropical cyclones, extreme precipitation) evaluated with the km-scale ensemble explored within the EU's Destination Earth (DestinE) initiative, the SPP scheme seems to generally provide better probabilistic predictions and introduces more ensemble spread than SPPT. Furthermore, forecasts using SPP have 24-hour precipitation distributions over the tropics that better match observations (based on IMERG satellite precipitation) than those produced with SPPT.

Conclusion and outlook

Adopting SPP restores the physical consistency of the ensemble members to the level of the unperturbed forecast. With SPP, precipitation, evaporation, sensible and latent heat fluxes as well as perturbed radiative fluxes are consistent with the perturbations in the atmospheric column. This leads to a degree of local conservation of moisture and enthalpy comparable to that achieved in the unperturbed forecast. With SPPT, the forecasts exhibit residuals which are on average at least an order of magnitude larger than the residuals in the control forecast. Moreover, the combination of many independent perturbed elements across the different parametrizations permits introducing a representation of uncertainties that goes beyond amplitude errors of the total physics tendencies. It permits, for instance, the representation of uncertainties in the shape of a heating profile.

The replacement of SPPT with SPP in Cycle 49r1 is the result of a major complex development effort, which had to balance requirements across lead times of hours in the Ensemble of Data Assimilations (EDA) to months in seasonal forecasts. Both SPPT and SPP have some effect on mean errors. The bias changes with SPPT can compensate the mean errors of the unperturbed model for some variables and regions of the atmosphere. In general, the mean errors of forecasts perturbed with SPP are closer to the mean errors of the unperturbed forecast than when SPPT is used. This is a desirable feature that will support IFS development in the future. However, it has the inevitable effect that biases increase for some variables and some regions when SPP replaces SPPT.

Noticeable improvements in skill with the introduction of SPP are achieved for 2-metre temperature and 10-metre wind. Further improvements will appear via the use of SPP in the EDA in Cycle 49r1. SPP has helped to address the underdispersion in the boundary layer, which has been seen as a limiting factor for the assimilation of 2-metre temperature observations. The introduction of SPP together with other changes in the EDA (higher resolution and soft re-centring) has permitted reducing the amplitude of the singular vector initial perturbations. This will partly address the overdispersion in the winter storm tracks in the early medium range.

In the coming years, it is planned to extend SPP into the parametrizations of the land surface. An initial investigation has explored the beneficial impact of increased ensemble spread from stochastic perturbations to land surface parameters in the context of land data assimilation.

SPP has gained popularity as a representation of model uncertainties and has been introduced already in the Canadian Global Ensemble Prediction system and in the regional HARMONIE ensemble prediction system (McTaggart-Cowan et al., 2022; Tsiringakis et al., 2024).

Further reading

Buizza, **R.**, **M. Miller** & **T.N. Palmer**, 1999: Stochastic representation of model uncertainties in the ECMWF ensemble prediction system. *Q. J. R. Meteorol. Soc.*, **125**, 2887–2908.

Lang, **S.T.**, **S.-J. Lock**, **M. Leutbecher**, **P. Bechtold** & **R.M. Forbes**, 2021: Revision of the Stochastically Perturbed Parametrisations model uncertainty scheme in the Integrated Forecasting System. *Q. J. R. Meteorol. Soc.*, **147**, 1364–1381.

Lin, **H.**, **G. Brunet** & **J. Derome**, 2009: An Observed Connection between the North Atlantic Oscillation and the Madden–Julian Oscillation. *J. Climate*, **22**, 364–380.

Lock, **S.-J.**, **S.T.K. Lang**, **M. Leutbecher**, **R.J. Hogan** & **F. Vitart**, 2019: Treatment of model uncertainty from radiation by the Stochastically Perturbed Parametrization Tendencies (SPPT) scheme and associated revisions in the ECMWF ensembles. *Q. J. R. Meteorol. Soc.*, **145**, 75–89.

McTaggart-Cowan, **R.**, **L. Separovic**, **R. Aider**, **M. Charron**, **M. Desgagné**, **P.L. Houtekamer et al.**, 2022: Using stochastically perturbed parameterizations to represent model uncertainty. Part I: Implementation and parameter sensitivity. *Mon. Wea. Rev.*, **150**, 2829–2858.

Ollinaho, **P.**, **S.-J. Lock**, **M. Leutbecher**, **P. Bechtold**, **A. Beljaars**, **A. Bozzo et al.**, 2017: Towards process-level representation of model uncertainties: Stochastically Perturbed Parametrizations in the ECMWF ensemble. *Q. J. R. Meteorol. Soc.*, **143**, 408–422.

Tsiringakis, **A.**, **I.-L. Frogner**, **W. de Rooy**, **U. Andrae**, **A. Hally**, **S. Contreras Osorio et al.**, 2024: An Update to the Stochastically Perturbed Parametrizations Scheme of HarmonEPS. *Mon. Wea. Rev.*, **152**, 1923–1943. *https://doi.org/10.1175/MWR-D-23-0212.1*

Wheeler, **M. C.** & **H.H. Hendon**, 2004: An All-Season Real-Time Multivariate MJO Index: Development of an Index for Monitoring and Prediction. *Mon. Wea. Rev.*, **132**, 1917–1932.

Data-driven ensemble forecasting with the AIFS

Mihai Alexe, Simon Lang, Mariana Clare, Martin Leutbecher, Christopher Roberts, Linus Magnusson, Matthew Chantry, Rilwan Adewoyin, Ana Prieto-Nemesio, Jesper Dramsch, Florian Pinault, Baudouin Raoult

Data-driven weather forecast models are a promising addition to physics-based numerical weather prediction (NWP) models. ECMWF now runs the Artificial Intelligence Forecasting System (AIFS) in an experimental real-time mode. It is run four times daily and is open to the public under ECMWF's open data policy. This AIFS version (henceforth referred to as 'deterministic AIFS') is trained to produce forecasts that minimise mean squared error (MSE) up to 72 h into the forecast. The MSE optimisation leads to excessive smoothing and reduced forecast activity (Lang et al., 2024(a)). This is detrimental to ensemble forecasts, which rely on a realistic representation of the intrinsic variability of the atmosphere.

In this article, we describe two training approaches for data-driven forecast models to produce skilful ensemble forecasts: *diffusion training* (Karras et al., 2022, and Price et al., 2024), where the forecast is the result of a denoising task, and *probabilistic training with a proper score objective adjusted for the finite ensemble size*, such as the fair continuous ranked probability score (fair CRPS; Leutbecher, 2019, and Kochkov et al., 2024).

Model and data

The forecast model for both methods, diffusion-based training and CRPS-based training, is the AIFS (Lang et al., 2024(a)). The AIFS is built around an 'encoder–processor–decoder' architecture. The encoder and decoder are attention-based graph neural networks (GNNs), and the processor is a sliding-window transformer. The latest version of the AIFS at the time of writing (0.2.1) was trained on approximately 40 years of Copernicus ERA5 reanalysis data and 'fine-tuned' on the ECMWF operational analysis from 2019 to 2020 to improve the skill of real-time forecasts.

CRPS-based training

In the AIFS–CRPS configuration, multiple model states (i.e. ensemble members) are propagated forward in time, as shown in Figure 1 (see for example Kochkov et al., 2024). For each ensemble member, a different realisation of random Gaussian noise is injected into the transformer processor. White noise is fed into the model, during both training and inference (forecasting), to be used by AIFS–CRPS to learn a representation of forecast model uncertainty. The CRPS training objective is calculated against the ERA5 deterministic reanalysis at the forecast target time. Perturbed initial conditions are generated by re-centring the ERA5 Ensemble of Data Assimilations (EDA) on the ERA5 deterministic reanalysis. This is consistent with the use of the EDA for ECMWF's Integrated Forecasting System (IFS), as described in Lang et al., 2015. In addition to model parallelism (sharding, see Lang et al., 2024(a)), the AIFS code can distribute ensemble members across several

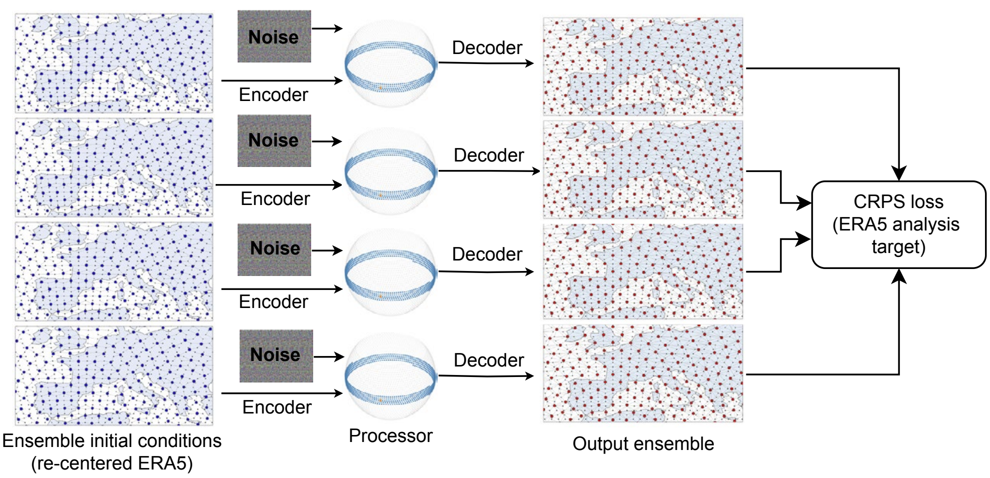

FIGURE 1 Probabilistic training with CRPS optimisation (AIFS–CRPS): the AIFS propagates four ensemble members that are then optimised jointly through the CRPS loss. The ensemble members can reside on separate GPU devices; in this case, a differentiable all-gather operation happens before the loss computation. Ensemble member trajectories start from different initial conditions (re-centered ERA5 EDA, see text) and receive different noise inputs.

doi: 10.21957/ma3p95hxe2

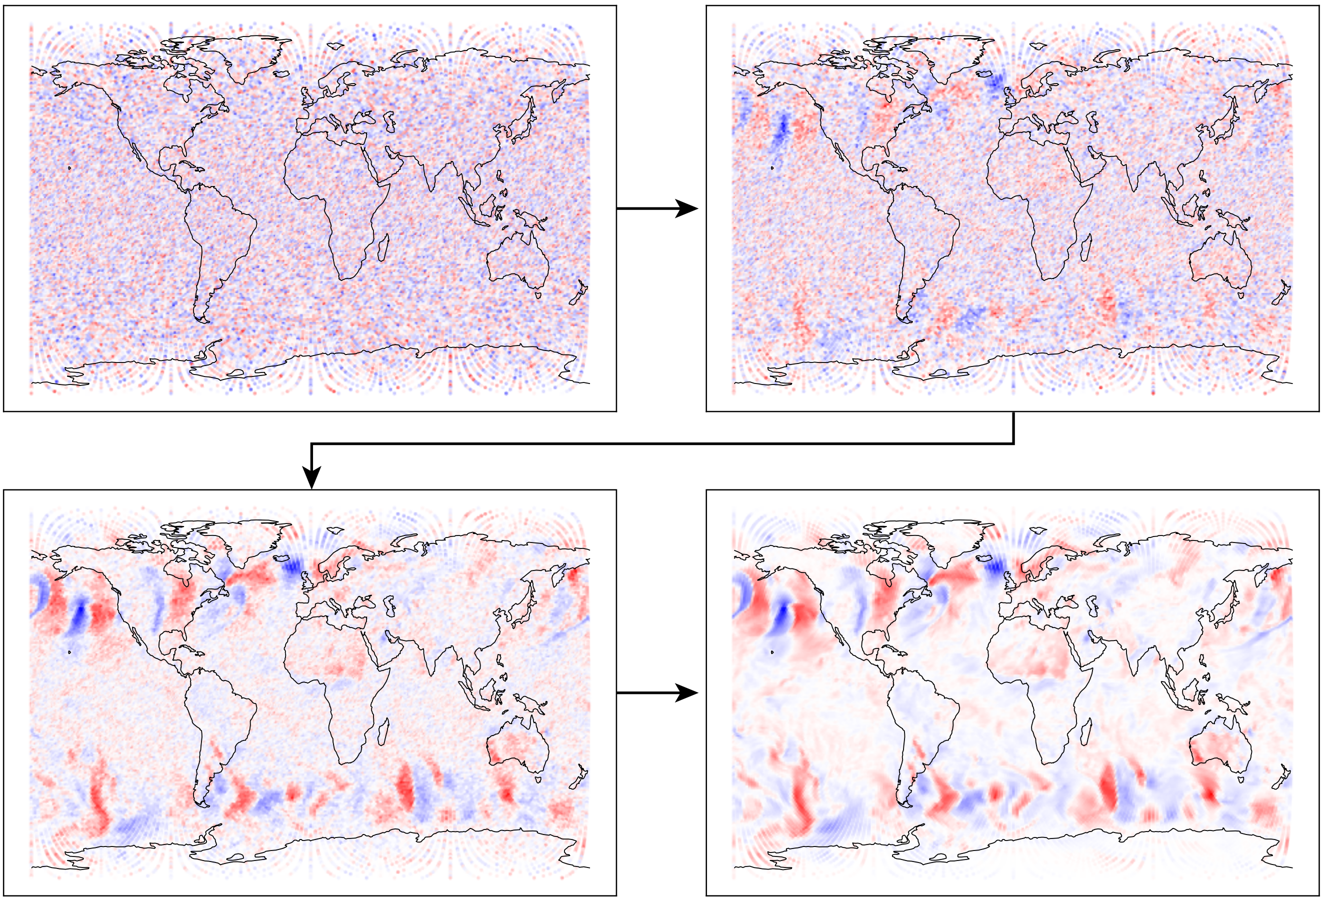

FIGURE 2 Diffusion training (AIFS–Diffusion): we show four (non-consecutive) steps from a denoising diffusion chain. Starting from pure Gaussian noise (top left), the model generates a 12-hour tendency (bottom right). The model has many variables and levels; for illustration purposes, the meridional wind component at 850 hPa has been selected.

graphics processing units (GPUs) to enable the training of larger ensembles at higher spatial resolution. AIFS–CRPS also implements autoregressive rollout during training, with 6-hour time steps; this makes it possible to optimise CRPS up to several days into the forecast. We found that a four-member ensemble was sufficient during training to arrive at a model that shows good probabilistic skill in both training and inference. Larger ensemble sizes are used during inference.

Diffusion-based training

In the diffusion approach (AIFS–Diffusion), the AIFS learns to remove noise from a forecast state, conditioned on the initial conditions and a noise schedule (Price et al., 2024; Karras et al., 2022; and Figure 2). During training, the model 'sees' different noise levels, i.e. increasingly noisy forecast states, all the way up to 'pure' noise. The model iterates on the same state using a sampling process, arriving at a 12-hour forecast *tendency* after 20 denoising steps. This increases the computational cost of a single forecast trajectory. Diffusion-based training usually requires a significantly larger number of training steps than deterministic training. On the other hand, AIFS–Diffusion does not incur the overhead of propagating multiple ensemble members as in AIFS–CRPS.
We have found that both ensemble configurations have comparable training costs at a horizontal grid spacing of approximately one degree (111 km).

Inference

During inference, AIFS–CRPS and AIFS–Diffusion start from the initial conditions of the operational IFS ensemble. The initial conditions include the singular vector component of the initial perturbations. Both AIFS–CRPS and AIFS–Diffusion are then run autoregressively to generate 15-day forecasts.
The ensembles are configured with a forecast step of 12 hours (AIFS–Diffusion, cf. Price et al., 2024) and 6 hours (AIFS–CRPS). Each ensemble member is independent, and thus the forecast generation is fully parallel. The cost of an AIFS–Diffusion forecast is significantly higher than that of an AIFS–CRPS forecast because the diffusion model is called multiple times per forecast step. That said, both data-driven approaches are very cheap when compared to the computational cost of an IFS ensemble member trajectory: e.g., when run on a single NVIDIA A100 GPU device, AIFS–

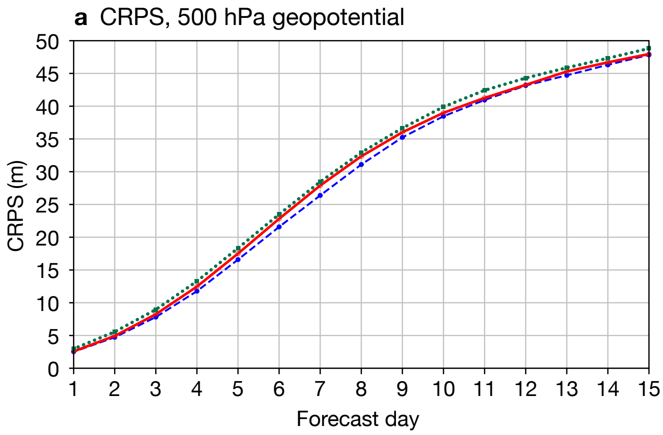

FIGURE 3 Fair CRPS scores of (a) 500 hPa geopotential height and (b) 850 hPa temperature, comparing the 50-member operational IFS ensemble with 8-member ensembles initialised from the ERA5 EDA for a 3-month period in 2019, using models trained with the proper score optimisation (AIFS–CRPS) and diffusion (AIFS–Diffusion) techniques.

Diffusion needs only about 2.5 minutes to produce a 15-day forecast ensemble member. For reference, one IFS Cycle 48r1 ensemble member takes about one hour to produce (excluding I/O), on 96 AMD Epyc Rome central processing units (CPUs). The operational IFS ensemble runs at a spatial resolution of approximately 9 km (Lang et al., 2023).

Forecast evaluation

To enable rapid testing at a small computational cost, we have thus far only trained models at a horizontal grid spacing of one degree, which is consistent with the configuration used for the development of the first deterministic (v0.1) AIFS system.

We found that both approaches produce skilful ensemble forecasts. In Figure 3 we compare AIFS–Diffusion and AIFS–CRPS initialised from perturbed, re-centred ERA5 analyses at O96 horizontal grid spacing (ca. one degree) to the 2019 IFS operational ensemble (ca. 18 km horizontal grid spacing). AIFS–Diffusion and AIFS–CRPS produce well-calibrated forecasts and generate realistic forecast variability. In contrast to deterministic AIFS

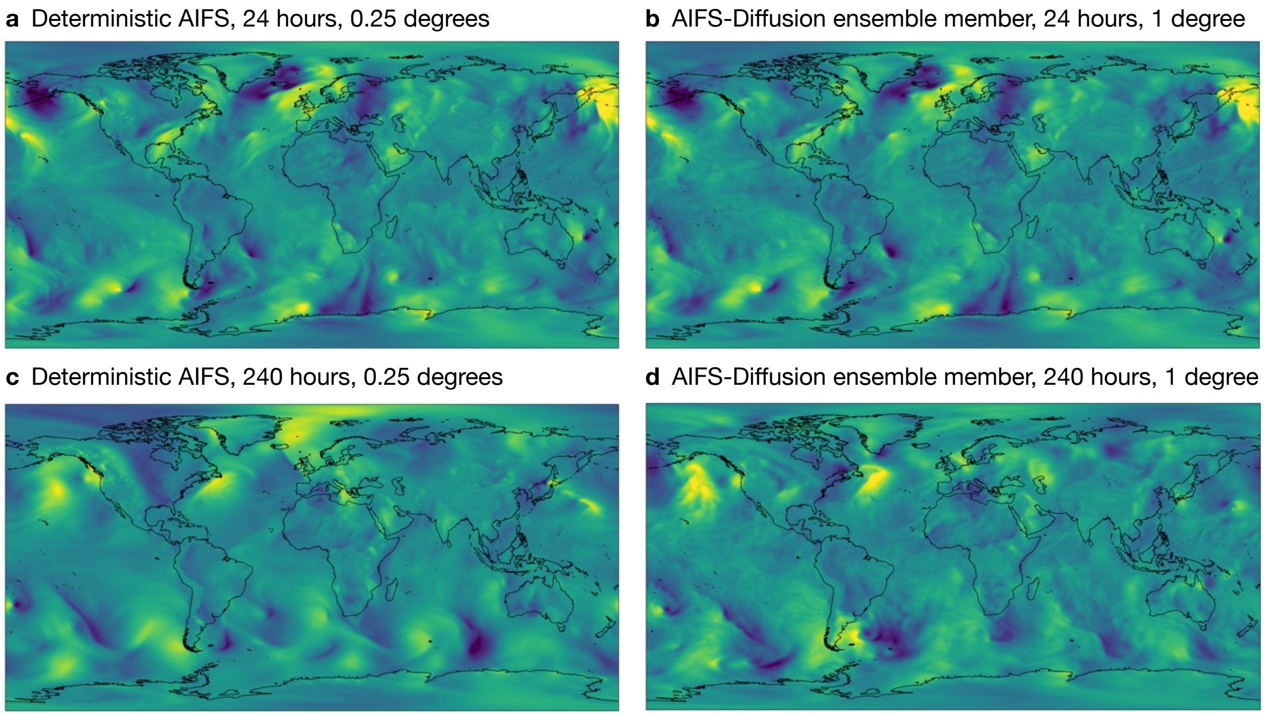

FIGURE 4 Depicted are (a) a 24-hour forecast of the deterministically (MSE) trained AIFS at N320 (a forecast with a horizontal grid spacing of ca. 0.25 degrees), (b) a 24-hour AIFS–Diffusion ensemble member at O96 (a forecast with a horizontal grid spacing of ca. 1 degree), (c) the same as (a) but showing a 240 h forecast, and (d) the same as (b) but showing a 240 h forecast. The forecasts the AIFS produces after probabilistic training (diffusion or fair CRPS) show a similar level of detail at short- and medium-range lead times.

forecasts, probabilistically trained AIFS ensemble members retain a similar level of detail at short- and medium-range lead times, as evidenced in Figure 4.

Implementation

AIFS–Diffusion was chosen as the first candidate for experimental real-time implementation. The operational IFS ensemble provides perturbed initial conditions for the data-driven ensemble forecast. After fine-tuning on operational IFS analyses, the resulting model is competitive with the 9 km IFS ensemble for upper-air scores (see Figure 3 in Lang et al., 2024(b)). It now runs twice daily in a 51-member configuration and produces a similar set of variables to that of the deterministic AIFS (Lang et al., 2024(b)). It is important to note that, while the control member of the AIFS–Diffusion ensemble configuration is started from unperturbed initial conditions, it nonetheless includes a representation of model uncertainty because of the stochastic sampling involved in calculating the forecast.

To better quantify and understand its forecast performance, the real-time AIFS–Diffusion ensemble is periodically evaluated by ECMWF analysts – see, e.g., the recent episode of exceptionally heavy rainfall in the United Arab Emirates described by Magnusson et al., 2024.

A cold snap over western Europe

The forecasting skill of the diffusion-trained ensemble can be illustrated with an example from France. A cold spell was observed over parts of central and western Europe in late April 2024. The cold air caused late-season, potentially damaging frost during the flowering period of fruit trees and grapevines. Figure 5 shows 2-metre temperature ensemble forecasts from the IFS ensemble and the experimental real-time AIFS–Diffusion ensemble. The forecasts are averaged over a 1x1-degree box located near Troyes, France, a winemaking region. Both forecasting systems successfully forecast the 24 April cold anomaly about 8–10 days before the event.

Sub-seasonal forecasts

Early evidence strongly suggests that the ensemble AIFS will also have a role to play in sub-seasonal forecasting. While deterministically trained data-driven forecast models are known to develop large biases over relatively long forecast horizons (Ben-Bouallègue et al., 2023), the systematic errors of the two

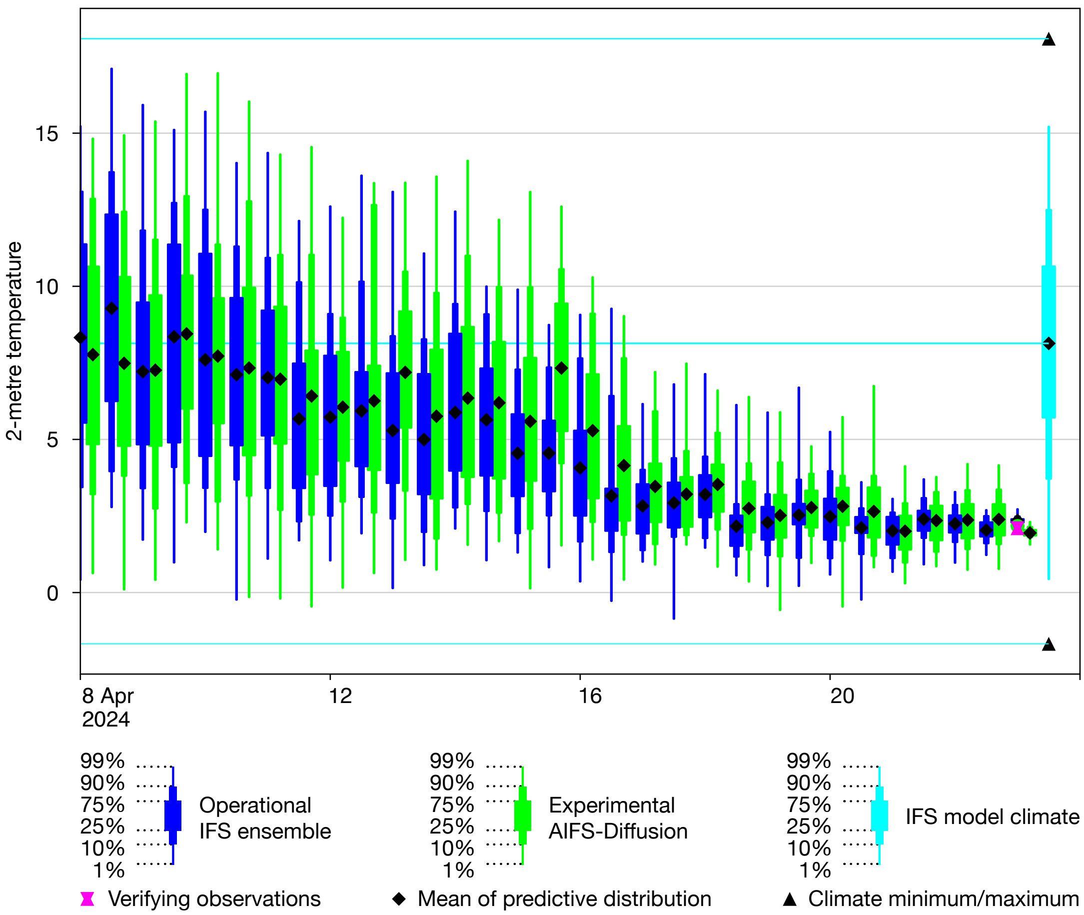

FIGURE 5 Ensemble forecasts of 2-metre temperature ahead of a cold spell over Europe, in late April 2024. The forecasts are averaged over a 1x1-degree box centred around 48.3°N, 4°E (near Troyes, France).

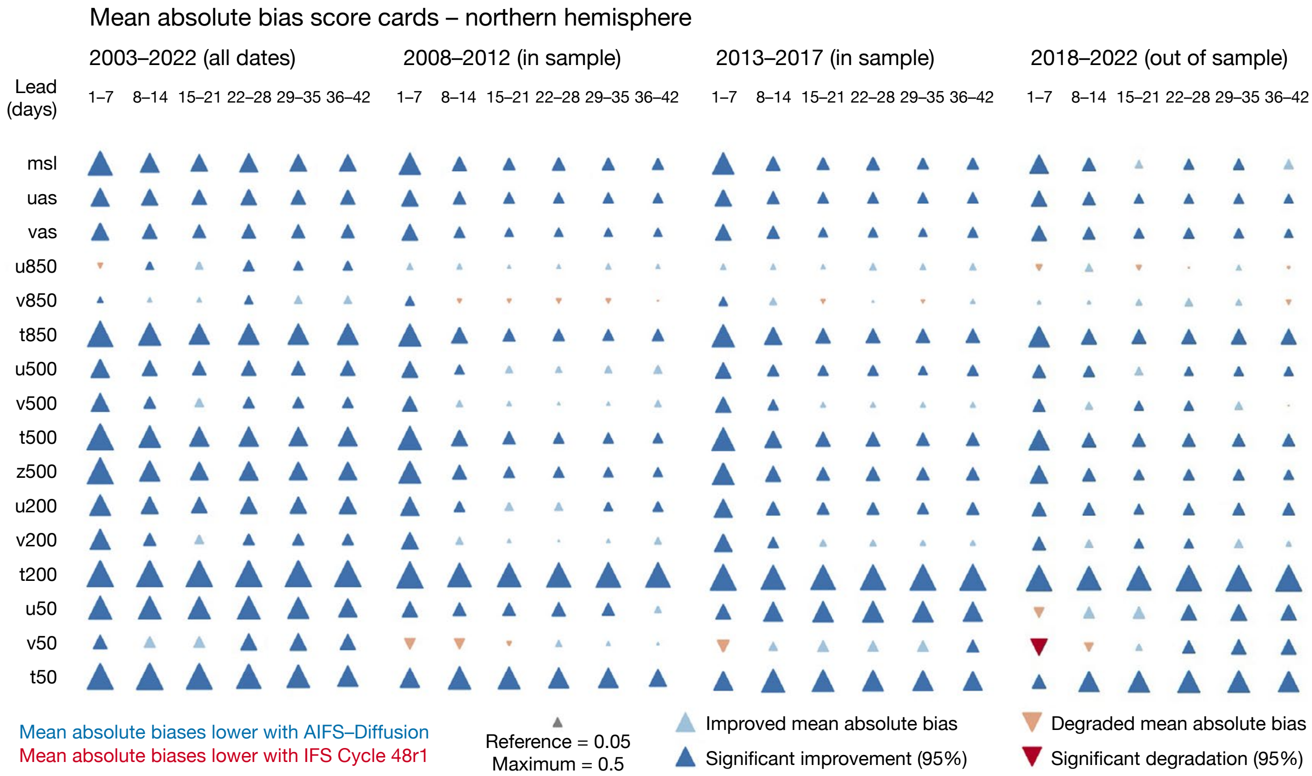

FIGURE 6 Scorecard summarising changes in mean absolute bias (MAB) for the northern hemisphere (30°N–90°N) for AIFS–Diffusion versus operational IFS (Cycle 48r1) sub-seasonal hindcasts, calculated as 1 – MAB_{AIFS}/MAB_{IFS} as described in Roberts et al. (2021). MAB is shown estimated for all available dates (2003–2022; left) and three different 5-year subsets, including data not used for training (2018–2022; right). Upward (blue) triangles indicate that absolute biases aggregated across all locations and start dates in AIFS–Diffusion are reduced compared to IFS Cycle 48r1. The variables shown are mean sea-level pressure (msl) and zonal/meridional wind at 10 m (uas/vas); temperature (t) and zonal/meridional wind (u/v) at different pressure levels (850, 500, 200 and 50 hPa); and geopotential height (z) at 500 hPa. For both systems, MAB is calculated relative to ERA5 using 8-member 46-day ensemble forecasts initialised every Monday and Thursday within the re-forecast period. Symbol areas are proportional to the fractional change in bias score and significance from the distribution created by block-bootstrap resampling of the available start dates.

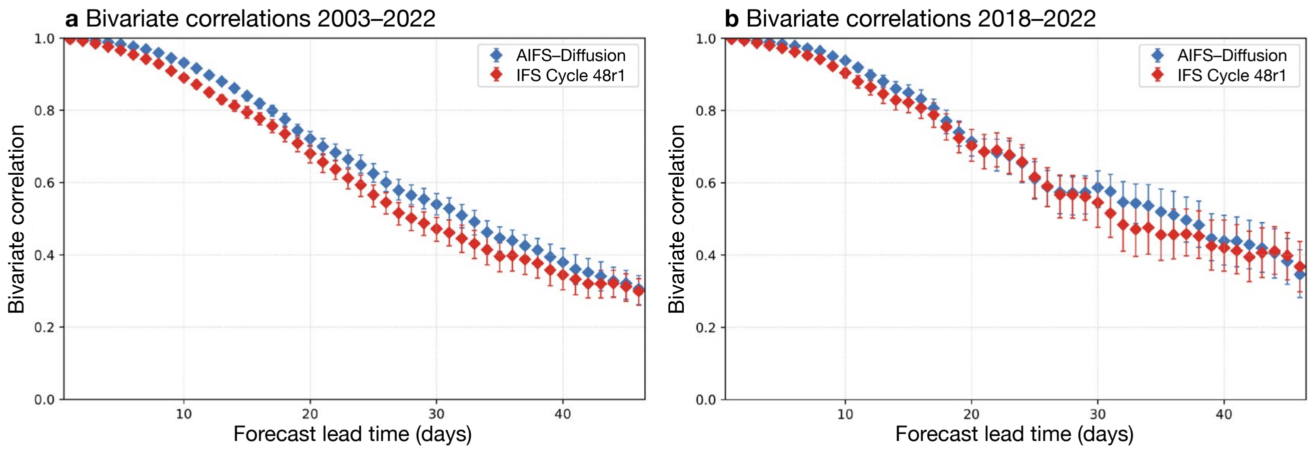

FIGURE 7 Bivariate correlations for an MJO index calculated from 200 hPa and 850 hPa zonal wind anomalies for AIFS–Diffusion (blue) and IFS Cycle 48r1 (red) calculated (a) using all available dates (2003–2022) and (b) data not used for training (2018–2022). Higher correlations mean better forecasts. The MJO index used here is an approximation for the full Wheeler and Hendon (2004) Real-time Multivariate MJO index as it excludes contributions from outgoing longwave radiation that are not available from AIFS–Diffusion. For both systems, correlations are calculated relative to ERA5 using 8-member 46-day ensemble forecasts initialised every Monday and Thursday within the re-forecast period. Error bars represent the 2.5th and 97.5th percentiles of the distribution created by block-bootstrap resampling of the available start dates.

probabilistic models described here are comparable to or smaller than the biases of the physics-based IFS, for a range of forecast parameters (see Figure 6 for AIFS–Diffusion vs the IFS). Notably, preliminary analyses of sub-seasonal AIFS–Diffusion ensembles show significant forecast skill, outperforming (weeks 1 and 2) or matching (week 3 and later) the skill of the IFS when predicting the Madden–Julian Oscillation (MJO), as shown in Figure 7.

Outlook

Probabilistic training of data-driven models results in skilful ensemble forecasts that also overcome one of the main limitations of deterministically trained models: the over-smoothing of forecast fields. Ongoing research aims to further increase forecast skill, to improve the fine-tuning approaches of the ensemble models on operational IFS analyses, to increase the temporal resolution, and to decrease horizontal grid spacing to 0.25 degrees. It is likely that higher-resolution ensembles will improve forecast scores for surface fields such as 2-metre temperature, precipitation, and 10-metre winds, as well as the representation of tropical cyclones.

Because data-driven ensemble forecasts are much cheaper to produce than their physics-based counterparts, it will be possible to add an AIFS–CRPS ensemble configuration to the experimental real-time suite, running alongside the diffusion-based system. This will allow a comprehensive evaluation of the strengths and weaknesses of both approaches.

Meteograms along with mean and spread products from the experimental AIFS real-time ensemble are available as open charts (*https://charts.ecmwf.int*) under ECMWF's open data policy. Further charts and data will be available in the near future.

Further reading

Ben-Bouallègue, **Z.**, **M.C.A. Clare**, **L. Magnusson**, **E. Gascon**, **M. Maier-Gerber**, **M. Janousek et al.**, 2023: The rise of data-driven weather forecasting. *https://arxiv.org/abs/2307.10128*

Karras, **T.**, **M. Aittala**, **T. Aila** & **S. Laine**, 2022: Elucidating the Design Space of Diffusion-Based Generative Models. *https://arxiv.org/abs/2206.00364*

Kochkov, **D.**, **J. Yuval**, **I. Langmore**, **P. Norgaard**, **J. Smith**, **G. Mooers et al.**, 2024: Neural general circulation models for weather and climate, *Nature*, **632**, 1060–1066. *https://doi.org/10.1038/s41586-024-07744-y*

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

Lang, **S.** & **M. Chantry**, 2024(b): Enter the ensembles. ECMWF website. *https://doi.org/10.21957/a791daf964*

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

Lang, **S.T.K.**, **D. Schepers** & **M. Rodwell**, 2023: IFS upgrade brings many improvements and unifies medium-range resolutions, *ECMWF Newsletter* **No. 176**, 21-28. *https://doi.org/10.21957/slk503fs2i*

Leutbecher, **M.**, 2019: Ensemble size: How suboptimal is less than infinity? *Q.J.R.M.S.*, **145** (Suppl. 1): 107–128. *https://doi.org/10.1002/qj.3387*

Magnusson, **L.**, **T. Hewson** & **M. Chevallier**, 2024: Unprecedented rainfall in the United Arab Emirates, *ECMWF Newsletter* **No. 180**, 2–3. *https://doi.org/10.21957/hd38w19ab6*

Price, **I.**, **A. Sanchez-Gonzalez**, **F. Alet**, **T.R. Andersson**, **A. El-Kadi**, **D. Masters et al.**, 2024: GenCast: Diffusion-based ensemble forecasting for medium-range weather. *https://arxiv.org/abs/2312.15796*

Roberts, **C.D.**, **F. Vitart** & **M.A. Balmaseda**, 2021: Hemispheric Impact of North Atlantic SSTs in Subseasonal Forecasts, *Geophysical Research Letters*, **48**(4). *https://doi.org/10.1029/2020GL091446*

Wheeler, **M.C.**, & **H.H. Hendon**, 2004: An All-Season Real-Time Multivariate MJO Index: Development of an Index for Monitoring and Prediction, *Monthly Weather Review*, **132**(8), 1917–1932. *https://doi.org/10.1175/1520-0493(2004)132<1917:AARMMI>2.0.CO;2*

The Copernicus Interactive Climate Atlas: a tool to explore regional climate change

José Manuel Gutiérrez (Instituto de Física de Cantabria, Spain), Daniel San Martín (Predictia, Spain), András Horányi, Anca Brookshaw, Carlo Buontempo (all ECMWF), Javier Díez-Sierra (Instituto de Física de Cantabria, Spain), Christopher Goddard (ECMWF), Max Tuni (Predictia, Spain)

The Copernicus Interactive Climate Atlas (*http://atlas.climate.copernicus.eu*, C3S Atlas in short) was released by the Copernicus Climate Change Service (C3S) in early 2024. It is a new C3S application (*https://cds.climate.copernicus.eu/applications*) which enables an interactive exploration of the Earth's climate, from recent changes and trends to possible climate futures under different emission scenarios. It uses key datasets that are available in the C3S Climate Data Store (CDS), including observation-based datasets (E-OBS), reanalyses (ECMWF's ERA5, ERA5-Land and ORAS5), and comprehensive global (CMIP5/6) and regional (CORDEX) climate projections. The goal is to produce authoritative climate change information for a wide range of physical variables. The variables characterise various types of climatic impact-driving conditions (heat and cold, wet and dry, wind and radiation, snow and ice, ocean, circulation) relevant for climate change risk assessments. The C3S Atlas is a new resource for policy makers wishing to formulate effective climate policy and for other users who need to visualise and analyse climate change information, particularly at the regional scale. This C3S tool is an evolution of the Intergovernmental Panel on Climate Change (IPCC) Interactive Atlas (IPCC-IA), which was frozen in 2021 with the publication of the Sixth Assessment Report's (AR6) WGI (Working Group I) section (see https://www.ipcc.ch/report/ar6/wg1).

The AR7 cycle has just started, and so the scope of the new Atlas has not been defined yet. In that sense, there is no formal agreement with the IPCC regarding the future. At the same time, the C3S Atlas has consolidated its role to become the natural evolution of the AR6 Atlas, and we believe the IPCC (WGI/II) might be interested in building on the C3S Atlas for AR7, in combination with other tools. We look forward to continuing the dialogue with the IPCC to ensure C3S can support the evolution of the IPCC's report in the most useful and efficient way.

The C3S Atlas has additional datasets and more variables than the IPCC-IA, and it will have further enhancements in the future. One key additional feature is the possibility to select a country or multiple countries as predefined regions (e.g. a combination of Portugal and Spain is shown in Figure 1), or to compute regional products on-the-fly for any customised regions. This makes it possible, for instance, to select any transboundary regions, where the information provided by the C3S Atlas can support coordinated actions of multiple countries. In any of the predefined and user-defined regions, the user can visualise regional information products, such as time series, climate stripe plots or annual cycles. The C3S Atlas provides different options to download the different products and the underlying data, including the monthly gridded dataset of the full set of variables across the different datasets underpinning the C3S Atlas. See *https://doi.org/10.24381/cds.h35hb680* for CDS access to the dataset. Figure 1 shows a screenshot of the C3S Atlas, illustrating its main components and some of the different products provided. Box A provides a summary of new features of the C3S Atlas.

Key variables for risk assessment

The C3S Atlas expands beyond the list of 21 variables/indices included in the IPCC-IA in order to better characterise key hazards. It thus supports and expands the global and regional climate change assessment made in the AR6 WGI report. Besides near-surface air and sea temperatures, precipitation and wind, the C3S Atlas includes additional essential climate variables, such as near-surface humidity, surface radiation fluxes, soil moisture and runoff, as well as various indices characterising different types of climatic impact-drivers. The selection of such new variables is based on the IPCC experience and particularly on user feedbacks on the IPCC-IA. For ease of navigation and to improve the user experience, all indices have been grouped around common themes: heat and cold, wet and dry, wind and radiation, snow and ice, ocean, and circulation. The first version of the C3S Atlas included a total of 30 variables and indices (Table 1). Periodic updates are planned to incorporate new variables and indices aligned with C3S needs and with key international initiatives, such as the IPCC. The full description of variables and indices is available in the user guidance for the underlying dataset

 doi: 10.21957/ah52ufc369

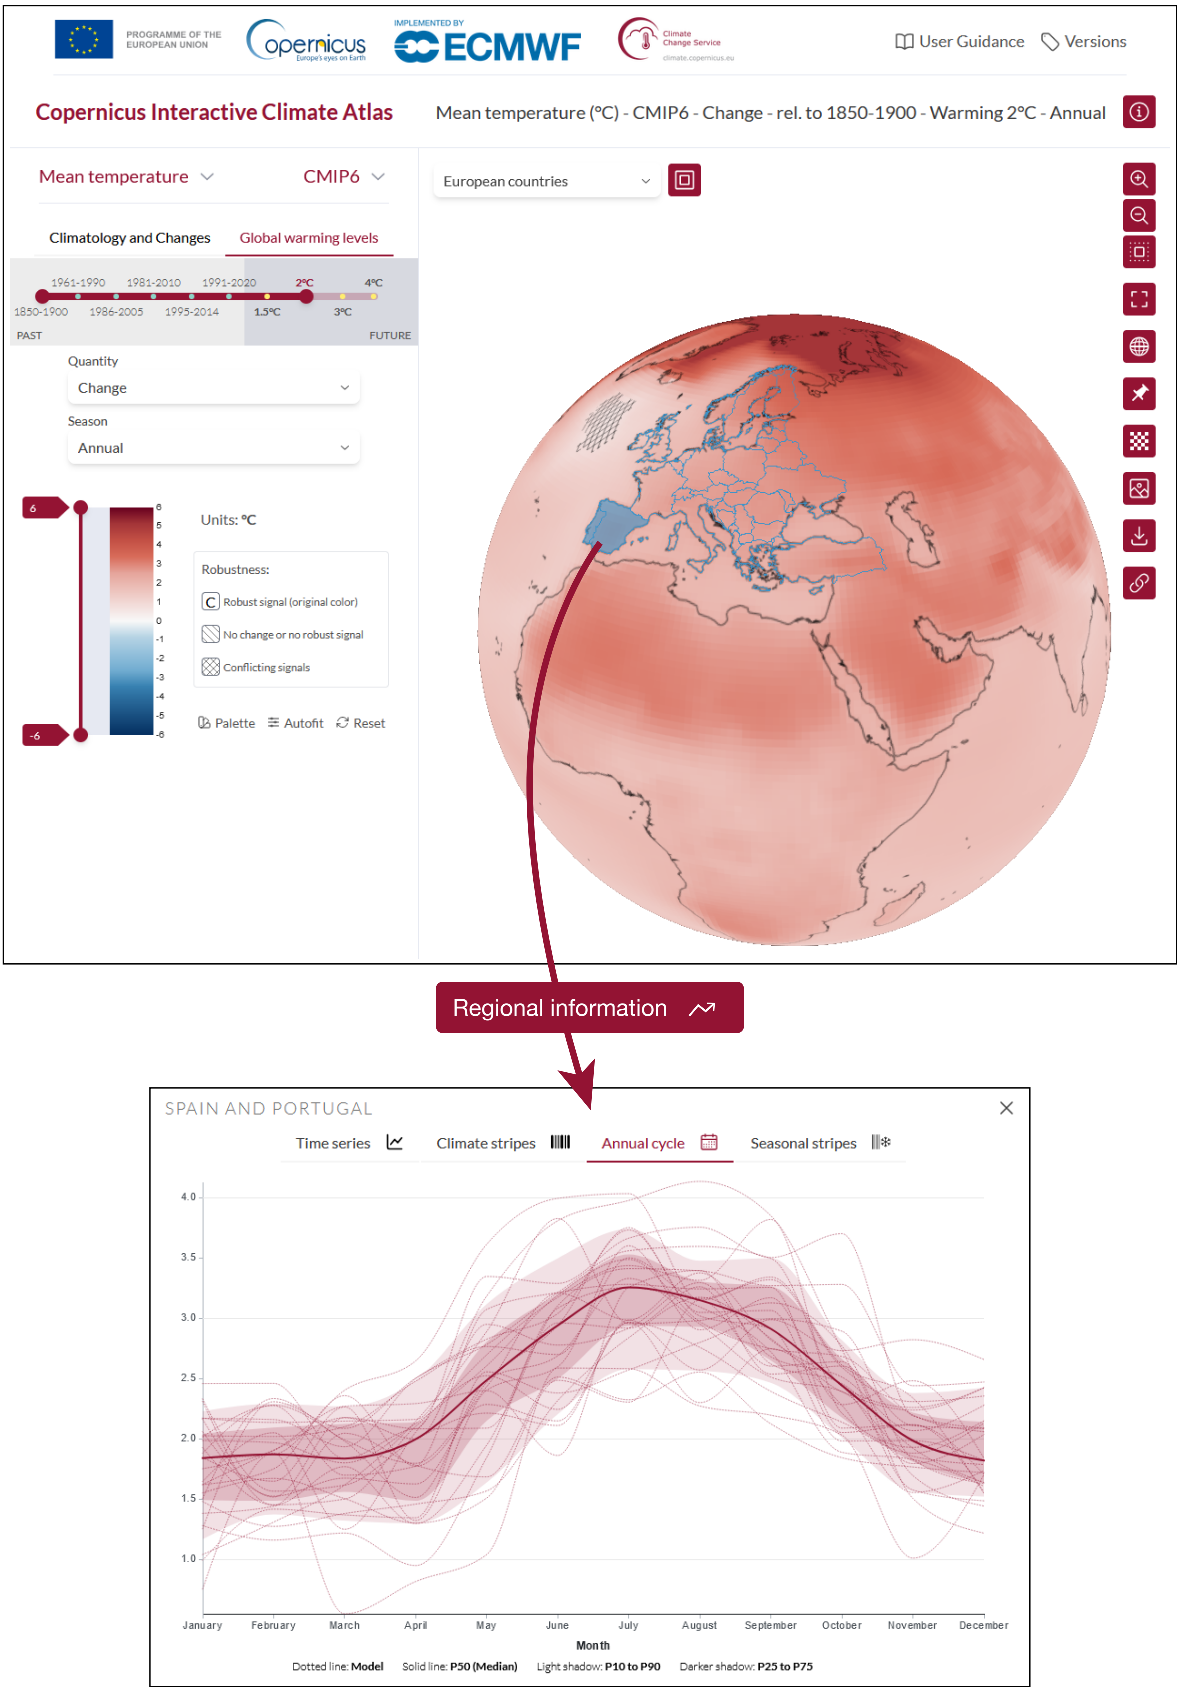

FIGURE 1 The top panel shows the landing page of the Copernicus Interactive Climate Atlas (C3S Atlas). It shows the mean temperature increase for global warming of 2°C, relative to the pre-industrial baseline 1850–1900. The left control panel enables users to configure the analysis dimensions (variable and dataset, season, periods, magnitudes: changes, trends, warming levels, etc.), and the right toolbar with buttons enables interactive exploration and figure/data exporting. Interactivity includes the selection of predefined (or customised) regions and the possibility to visualise regional information products, as shown in the bottom panel (selecting Spain and Portugal in the 'European countries' as predefined regions in the regional selector).

a **Summary information about the new features of the C3S Atlas**

The new Copernicus Interactive Climate Atlas (C3S Atlas) is a C3S application which can be used to access and show authoritative datasets for monitoring and assessing the evolution of key climate variables and indices. This harmonised and processed information (common grids, calendars, etc.) is also published as an additional C3S dataset in the CDS. It ensures reproducibility and reusability and provides a generic dataset for climate change risk assessment. This data harmonisation, pre-processing and quality control can help with the use of the data and can prevent any possible errors made in these steps by the users. It also facilitates downstream applications. Finally, it provides several new features with respect to any previously used tools in general and the IPCC-IA in particular. Such features are:

- the possibility of selecting national and trans-national regions, either predefined or user-defined, for detailed regional climate information
- periodic updates, including additional variables and indices (e.g. nine additional variables with respect to the IPCC-IA in the first version: SPEI6, huss, evspsbl, mrsos, mrro, clt, rsds, rlds and psl – see Table 1 for the variable codes)
- using quality-assured datasets from the C3S CDS, such as ERA5-Land (in addition to ERA5 originally included in the IPCC-IA) as a global reanalysis, the ORAS5 ocean reanalysis, CORDEX-CORE Regional Climate Model (RCM) simulations as well as regional EURO-CORDEX data, and Global Climate Models (GCMs from CMIP5 and CMIP6)
- the monthly dataset behind the C3S Atlas is fully available in the C3S Climate Data Store (CDS): https://doi.org/10.24381/cds.h35hb680.

under '2.2 Variables and indices' at: *https://confluence.ecmwf.int/display/CKB/Gridded+data+underpinning+the+Copernicus+Interactive+Climate+Atlas%3A+Description+of+the+datasets+and+variables*

This set of 30 indices has been systematically computed (where data are available) for a number of authoritative and quality-assured C3S datasets (Table 2). They provide information about the past and present climate (observations and reanalyses) and our possible climate futures (global and regional climate projections), with key complementary lines of evidence for climate change risk assessment.

The C3S Atlas includes observation-based products, such as the E-OBS gridded observational database, which tracks temperature in Europe since 1950, as well as reanalyses like ERA5 and ERA5-Land. These reanalyses represent the state of our climate since 1940 (1950 for ERA5-Land). The application shows the data until 2022 and will be updated regularly, for instance every March, when all the data for the previous year are consolidated. Meanwhile, the near-real time ERA5 dataset ERA5T – up to five days before the current day – is available in the Climate Data Store. The C3S Atlas also includes ocean data (ORAS5 ocean reanalysis) from 1958 to 2014 to analyse the climatic conditions of different oceanic variables (sea-surface temperature and sea-ice extent).

The Atlas includes state-of-the-art climate model projections to explore a variety of possible climate futures. They describe the evolution of the climate system based on simulations produced with Global and Regional Climate Models (GCMs and RCMs). They are forced with historical conditions and different future emission scenarios over the period from 1850 to 2100. The Coupled Model Intercomparison Project Phase 6 (CMIP6) is the reference framework for global climate projections and was essential to the work of the IPCC AR6. The application also provides access to the fifth phase of CMIP, CMIP5, which was the basis for the IPCC Fifth Assessment Report (AR5) and the special report for 1.5° global warming. This is still widely used as an alternative line of evidence.

Besides the global projections, the C3S Atlas also gives access to the regional projections created by the Coordinated Regional Climate Downscaling Experiment (CORDEX) driven by CMIP5 global forcings; in particular, CORDEX-EUR-11 high-resolution projections and CORDEX-CORE, complementing the CORDEX domain-by-domain information included in the IPCC Interactive Atlas. These two datasets provide the projections with the highest resolution available for the European domain (12.5 km) and for a global mosaic covering all land inhabited regions (25 km), respectively.

Downloading the visual products and the underlying data

A variety of options exist to download the visual outputs and the underlying data in user-friendly formats. Maps can be downloaded as PNG files and the underlying data as NetCDF, commonly used by climate practitioners, or GeoTIFF, commonly used by the

Group category	Code	Index	Units
heat and cold	T	Monthly mean of daily mean temperature	°C
	TN	Monthly mean of daily minimum temperature	°C
	TX	Monthly mean of daily maximum temperature	°C
	TNn	Monthly minimum of daily minimum temperature	°C
	TXx	Monthly maximum of daily maximum temperature	°C
	TX35	Monthly count of days with maximum temperature above 35°C	days
	TX35ba	Monthly bias adjusted TX35	days
	TX40	Monthly count of days with maximum temperature above 40°C	days
	TX40ba	Monthly bias adjusted TX40	days
	FD	Monthly count of frost days	days
	HD	Annual heating degree-days	°C day
	CD	Annual cooling degree-days	°C day
wet and dry	PR	Monthly mean of daily accumulated precipitation	mm
	RX1day	Monthly maximum of 1-day accumulated precipitation	mm
	RX5day	Monthly maximum of 5-day accumulated precipitation	mm
	CDD	Annual consecutive dry days	days
	SPI6	Standardized Precipitation Index (SPI) for 6 months cumulation period	1
	SPEI6	Standardized Precipitation Evaporation Index (SPEI) for 6 months cumulation period	1
	huss	Monthly near surface specific humidity	1
	evspsbl	Monthly evaporation including sublimation and transpiration	mm
	mrsos	Monthly soil moisture in upper soil portion	kg m^{-2}
	mrro	Monthly total runoff	kg m^{-2}
snow and ice	prsn	Monthly mean of daily accumulated snowfall precipitation	mm
	siconc	Monthly mean of sea-ice area percentage	%
wind and radiation	sfcwind	Monthly mean of daily mean wind speed	m s^{-1}
	clt	Monthly fraction of cloud cover	1
	rsds	Monthly surface solar radiation downwards	W m^{-2}
	rlds	Monthly surface thermal radiation downwards	W m^{-2}
ocean	sst	Monthly mean of sea surface temperature	°C
circulation	psl	Monthly sea level pressure	Pa

TABLE 1 Description of the 30 climate variables and indices included in the C3S Atlas.

Project	C3S Atlas Subset	Resolution	CDS-catalogues
CMIP6	CMIP6	1° (*)	Projections-cmip6
CMIP5	CMIP5	2° (*)	Projections-cmip5-daily-single-levels [daily] Projections-cmip5-monthly-single-levels [monthly]
CORDEX	CORDEX-CORE	0.25° (*)	Projections-cordex-domains-single-levels [daily/monthly]
	CORDEX-EUR-11	0.125° (*)	
ERA5	ERA5	0.25°	Reanalysis-era5-single-levels [hourly] Reanalysis-era5-single-levels-monthly-means [monthly]
ERA5-Land	ERA5-Land	0.1°	Reanalysis-era5-land [hourly] Reanalysis-era5-land-monthly-means [monthly]
E-OBS	E-OBS	0.125° (*)	Insitu-gridded-observations-europe [daily]
ORAS5	ORAS5	0.25° (*)	Reanalysis-oras5 (consolidated, single levels) [daily]

TABLE 2 Descriptions of the datasets and CDS catalogues used for the C3S Atlas dataset. The column ‘resolution’ indicates the horizontal resolution and an asterisk indicates that the original dataset from the CDS has been regridded to the common nested sub-grids: 2°, 1°, 0.5°, 0.25°, 0.125°.

Geographical Information Systems (GIS) community. All regional information products can be exported in PDF and PNG formats, and the underlying data (numbers) as a CSV file. All these products are distributed with an open licence, facilitating reusability.

Moreover, the entire dataset underpinning the C3S Atlas has been published in the C3S CDS catalogue ('Gridded dataset underpinning the Copernicus Interactive Climate Atlas', *https://doi.org/10.24381/cds.h35hb680*).

Discovering the C3S Atlas: an illustrative example

A selection panel with the main choices (see left in the top panel of Figure 1) is available to explore recent and future climate, including the selection of the variable, dataset, period of analysis, and part of the year (referred to as 'season', offering a choice between annual, seasonal or monthly). The selection of the dataset determines the details of analysis, which are different for observational and reanalysis datasets and for climate projections. The different options are described below using runoff as an illustrative example. Note that it is part of the set of variables (runoff, evaporation and soil moisture) included in the C3S Atlas to further characterise the hydrological cycle and hydrological droughts, thus extending and complementing the rainfall information originally included in the IPCC-IA. For observational and

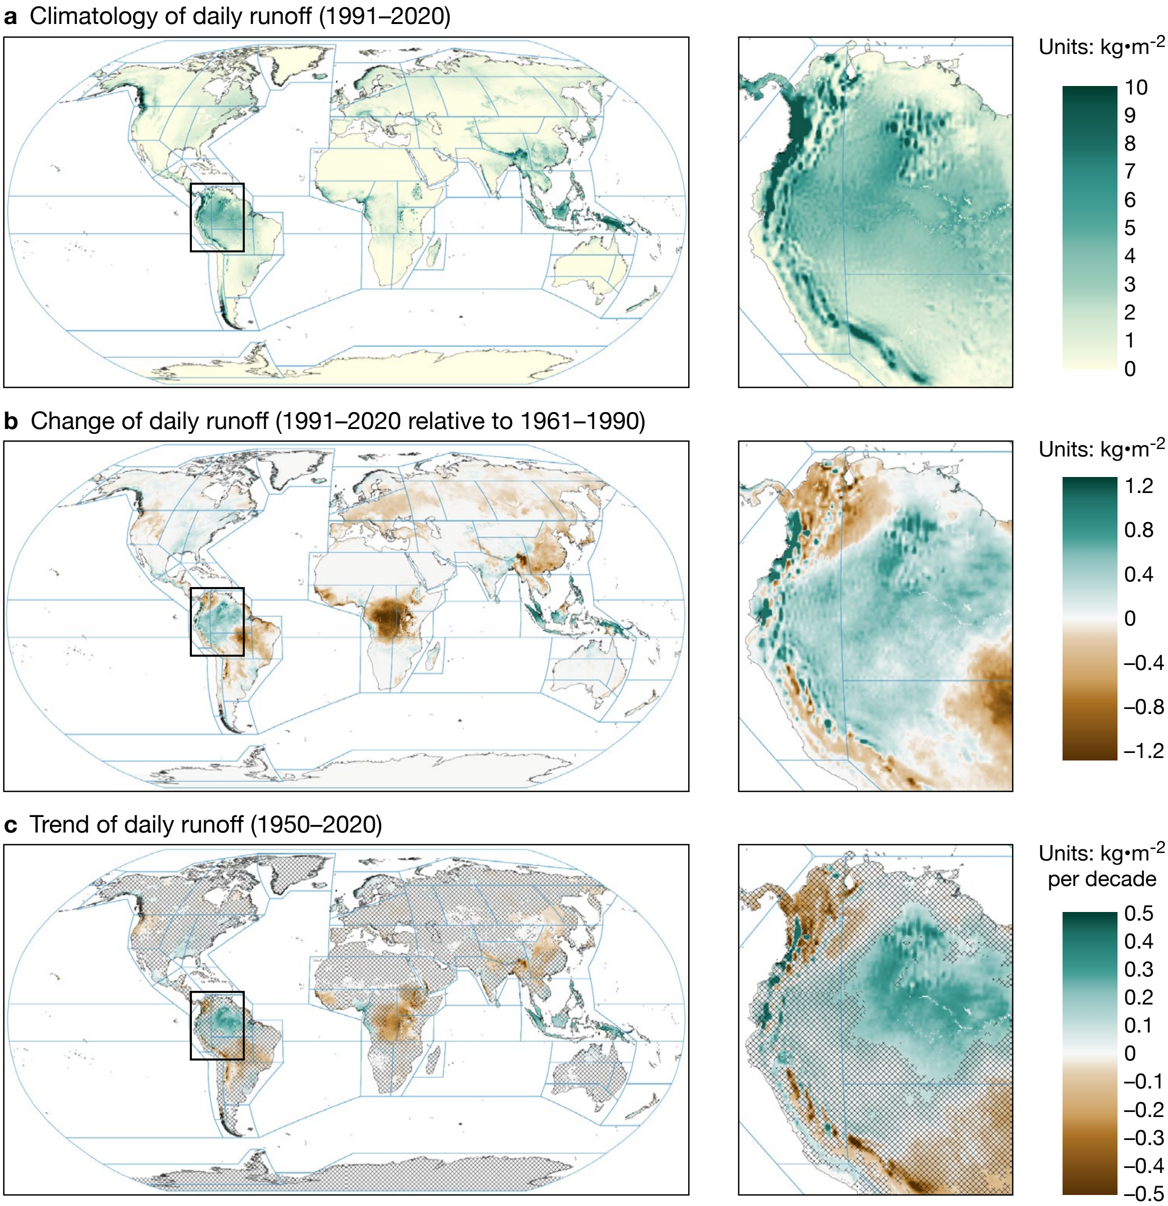

FIGURE 2 Different dimensions of analysis for the annual mean daily runoff variable (kg·m^{-2}) for the ERA5-Land dataset, showing (a) the climatology of the reference period 1991–2020, and two alternative magnitudes of change: (b) change for 1991–2020 relative to the 1961–1990 period and (c) the trend for the 1950–2020 period (showing non-significant trend regions with crosses). Besides the global map, the C3S Atlas makes it possible to zoom in to get full regional details, e.g. in the northern part of South America, as shown in the figures.

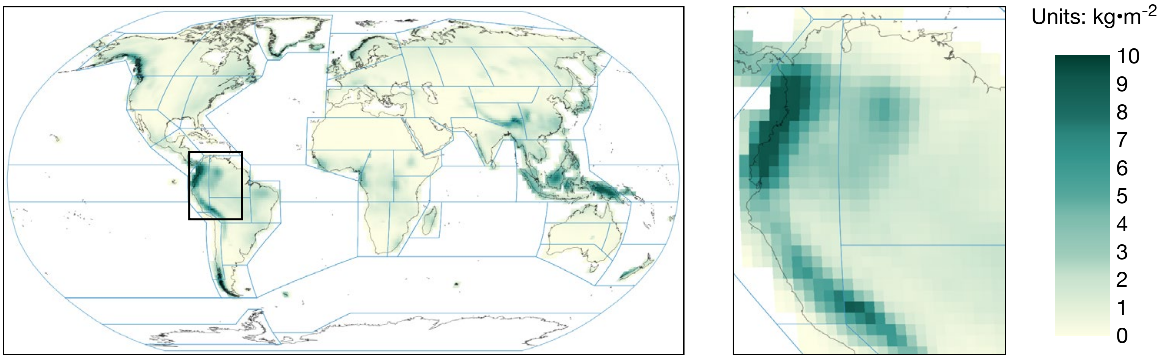

FIGURE 3 Climatology of mean daily runoff from the CMIP6 ensemble mean for the reference period 1991–2020. Besides the global map, the C3S Atlas makes it possible to zoom in to get full regional details, e.g. in the northern part of South America, as shown in the right panel.

reanalysis datasets, the C3S Atlas offers climatologies, changes relative to a baseline period, and trends, for a number of predefined historical periods. The resulting information is graphically represented in the form of a map showing gridded information for the selected analysis period. Each dataset is shown in its full spatial extent (global or regional), with its corresponding spatial resolution. For instance, Figure 2 illustrates the possibilities of the C3S Atlas to analyse present climate conditions and recent changes for the annual mean daily runoff variable ($kg{\cdot}m^{-2}$, corresponding to mm for water) from ERA5-Land, showing consistent change and trend patterns.

For the climate projection datasets, the C3S Atlas enables the display of climatologies and changes for historical periods, as in the previous case (see Figure 3).

Besides the historical periods, which can be used with observations and reanalysis, the ‘climatology and changes’ dimension makes it possible to explore future periods (near-, medium- and long-term, defined as 2021–40, 2041–60 and 2081–2100, respectively) across different emission scenarios: Representative Concentration Pathways (RCPs) for CMIP5 and CORDEX or the Shared-Socio-economic Pathways (SSPs) introduced for CMIP6. An additional dimension of analysis is the policy-relevant Global Warming Levels (GWL) used extensively in the IPCC AR6 report. In particular, the C3S Atlas enables the selection of 1.5°C, 2°C, 3°C and 4°C. Global warming levels have been computed, following the methodology used in the IPCC AR6 WGI Atlas, using the 20-year periods when models first reach a particular global warming level relative to the pre-industrial 1850–1900 period. These periods are shown in the Atlas, as illustrated in Figure 4.

Figure 5 illustrates the use of the C3S Atlas to contrast different lines of evidence for future projections.
It shows runoff changes according to the global CMIP6 ensemble and the regional CORDEX-EUR-11 one. They provide consistent information for the future changes in runoff under global warming of 2°C, with respect to conditions in 1991–2020.

Climate change maps include information on the robustness or uncertainty of the displayed signal. Following the AR6 WGI method, the C3S Atlas uses model agreement (with an 80% threshold) and signal emergence (relative to internal variability) to provide three robustness categories: (a) robust signal, (b) no change or no robust signal, and (c) conflicting signal.

Detailed regional information with a variety of visual products

Beyond spatial map information, the C3S Atlas makes it possible to explore regionally aggregated information for a number of predefined regions, shown in the ‘region set’ selector. Single or multiple regions can be selected by clicking directly on the map. Predefined regions include:

- the IPCC AR6 reference regions, which were used in the AR6 WGI report for regional climate change assessment
- the EUCRA regions, which are used in the European Climate Risk Assessment, and
- European countries, including those countries covered by the regional European datasets: E-OBS and CORDEX-EUR-11.

The regionally averaged information shown by the C3S Atlas for these predefined regions is pre-computed and can be explored interactively by clicking the ‘regional information’ button, which becomes visible when a region is selected. The user can also select customised regions, which can be defined using the ‘user defined’ option in the region selector. A custom new region can be drawn directly on the map after clicking the ‘pencil’ button. This action creates an offline job which enters a queue system, passing through different states until completion, typically in a few seconds.

Complementary aspects of regional information are

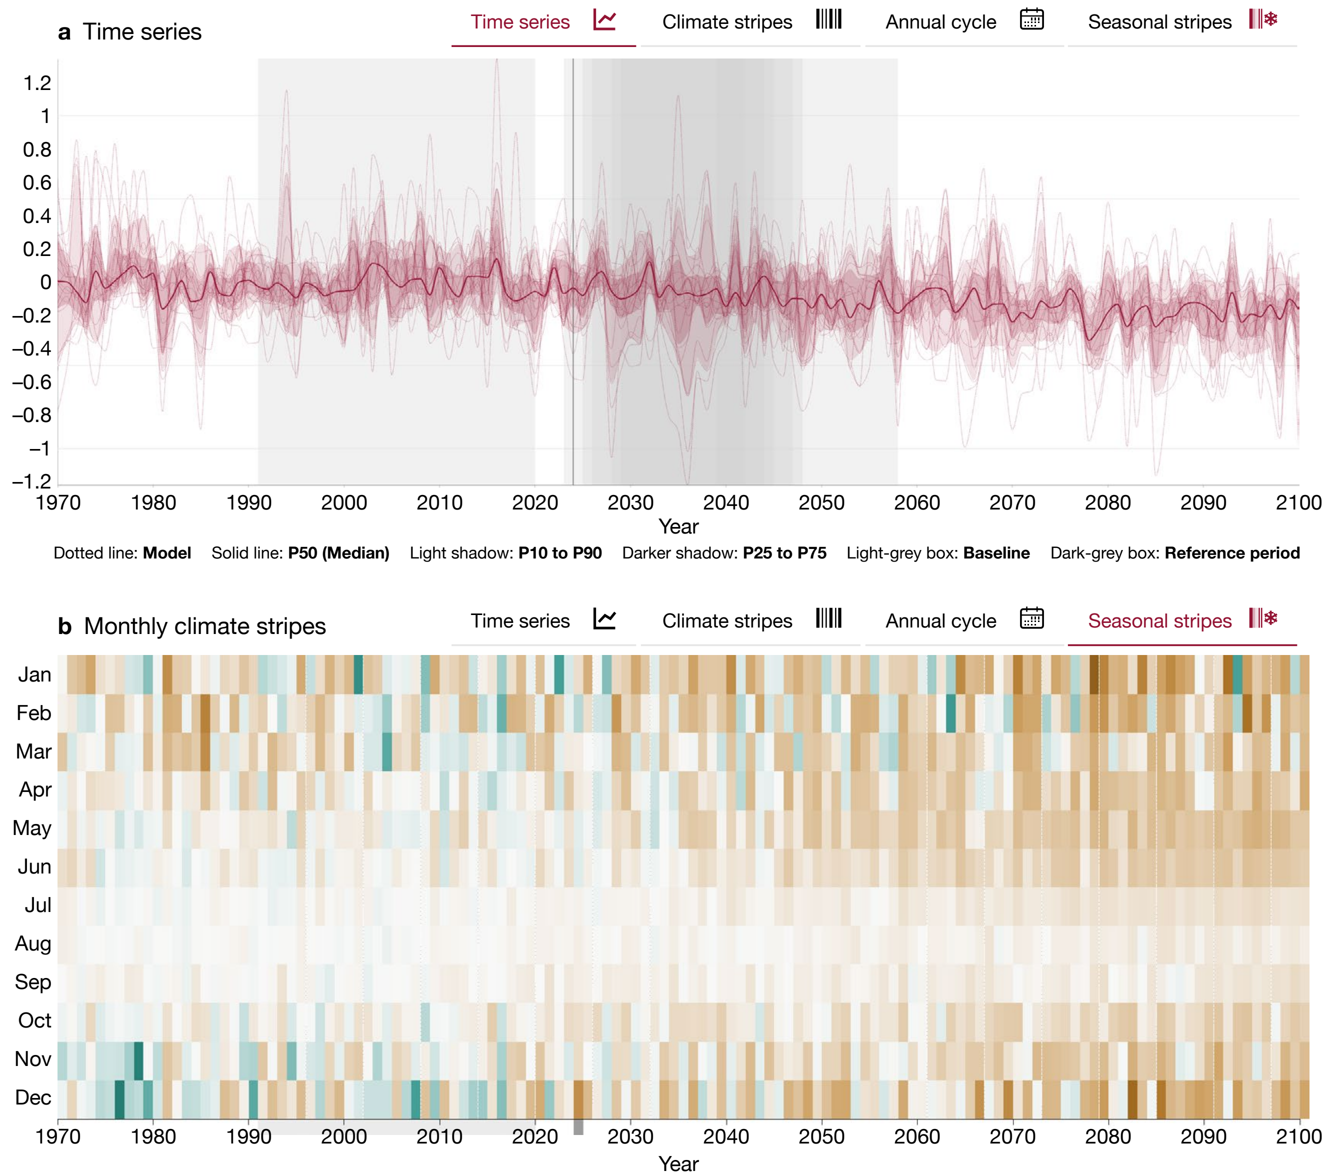

FIGURE 4 Different graphical products for regional information on runoff over a predefined region (selecting Spain and Portugal among the predefined European countries), including (a) a time series showing all ensemble members and (b) monthly climate stripes, showing the ensemble median monthly changes. Note that the time series plot includes light grey shading indicating the selected baseline (1991–2020) and darker grey shading for reference periods (2ºC global warming level). The latter are shown as 20-year periods from the ensemble members, when the given global warming level is first reached. All regional graphical information products make it possible to export the results in PDF and PNG formats, and also to export the underlying data (numbers) as a CSV file.

shown using different graphical products (Figure 4). All these graphical elements are dynamically updated when changing the choices in the selection panel.

FAIR principles for reproducibility and reusability

The development of the IPCC Interactive Atlas embodied a pioneering effort to integrate FAIR data principles (for Findability, Accessibility, Interoperability and Reusability) into climate policy reports (such as IPCC reports), thereby significantly enhancing their transparency and reproducibility (Iturbide et al., 2022). These principles have been adopted and expanded in the C3S Atlas by publishing the underpinning dataset in the C3S catalogue as described above, thus facilitating findability, accessibility, and reusability, and thoroughly documenting the data sources and processes used to produce the climatic products shown by the C3S Atlas. These activities will be complemented in the future by providing reusable code in Jupyter notebooks. These notebooks will illustrate the workflow followed to produce the C3S Atlas dataset and graphical products. Additionally, machine-readable standard provenance information will be provided for reproducibility. This information comprises a comprehensive description of the main climate data sources (primarily CMIP5/6 and CORDEX subsets), post-processing methods (such as

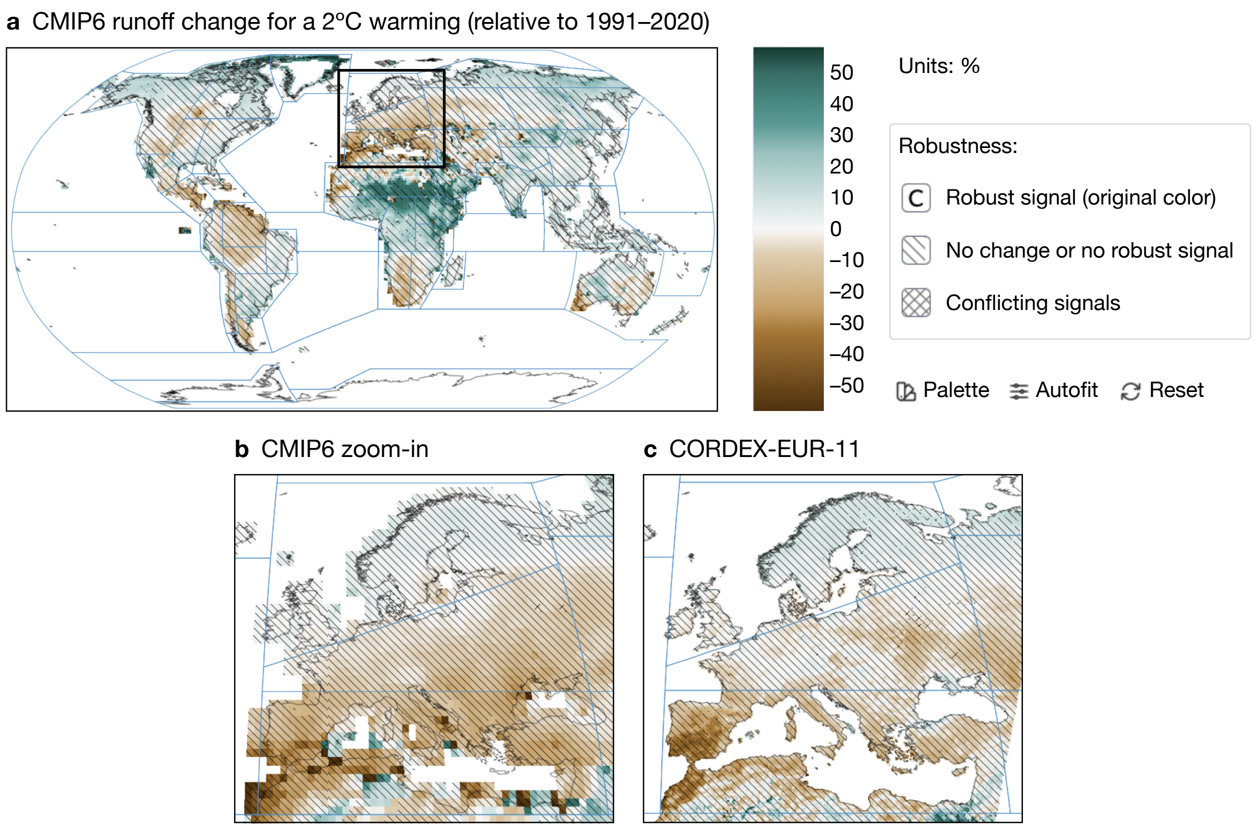

FIGURE 5 Climatology of mean daily runoff from the CMIP6 ensemble mean for the reference period 1991–2020. Besides (a) the global map, the C3S Atlas makes it possible to zoom in to get full regional details, e.g. in (b) Europe, and show (c) the higher-resolution CORDEX-EUR-11 simulations.

temporal aggregation and regridding), calibration techniques (including bias adjustment), and graphical outputs (such as geographical extent, colour bars, and displayed entities and layers).

You can find all the practical navigation details of the C3S Atlas in the comprehensive User Guide: *https://confluence.ecmwf.int/display/CKB/Copernicus+Interactive+Climate+Atlas%3A+User+Guide*

Acknowledgement

This work has been funded by the Copernicus Climate Change Service (C3S) through the contract (C3S2_381) led by the Spanish National Research Council (Instituto de Física de Cantabria, IFCA) with the participation of the data management company Predictia. ECMWF implements C3S on behalf of the European Commission.

Further reading

Iturbide, **M.**, **J. Fernández**, **J.M. Gutiérrez**, **A. Pirani**, **D. Huard**, **A. Khourdajie et al.**, 2022: Implementation of FAIR principles in the IPCC: the WGI AR6 Atlas repository, *Scientific Data*, **9**(629). *https://doi.org/10.1038/s41597-022-01739-y*

ECMWF publications

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

Technical Memoranda

919 **Hewson**, **T.** & **M. Chevallier**: Use and verification of ECMWF products. *September 2024*

918 **Haiden**, **T.**, **M. Janousek**, **F. Vitart**, **M. Tanguy**, **F. Prates** & **M. Chevallier**: Evaluation of ECMWF forecasts. *September 2024*

ESA or EUMETSAT Contract Reports

Duncan, **D.** & **N. Bormann**: Assessing RFI flags at passive microwave bands with an NWP model. *October 2024*

Salonen, **K.**, **P. Weston** & **P. de Rosnay**: Quarter 3 2024: Operations Service Report. *October 2024*

EUMETSAT/ECMWF Fellowship Programme Research Reports

64 **Scanlon**, **T.**, **A. Geer**, **N. Bormann** & **P. Browne**: Improving Ocean Surface Temperature for NWP using All-Sky Microwave Imager Observations. *August 2024*

ECMWF Calendar 2024/25

2024	
Oct 21–22	Finance Committee
Oct 22	Policy Advisory Committee
Oct 29	Advisory Committee of Co-operating States (virtual)
Nov 4–8	Training course: Predictability and ensemble forecast systems
Nov 11–15	Training course: Numerical methods for weather prediction
Nov 19–22	NWP SAF Workshop on Satellite Observations of the Earth System
Dec 10–11	Council
2025	
Feb 3–6	Training course: Use and interpretation of ECMWF products
Mar 3–7	Training course: Parametrization of subgrid physical processes
Mar 17–21	Training course: Data assimilation & machine learning
Mar 24–28	Training course: EUMETSAT/ECMWF NWP-SAF satellite data assimilation
Apr 7–11	**ECMWF's 50th anniversary events in Bonn, Germany**
	Apr 7–11 Annual Seminar 2025
	Apr 8–9 The evolution of Copernicus Services at ECMWF: stakes and challenges
	Apr 9–10 Workshop on ancillary data for land surface and Earth system modelling
	Apr 9–10 Workshop on surface process coupling and its interactions with the atmosphere
	Apr 9–10 Workshop on data assimilation: initial conditions and beyond
Apr 29	Policy Advisory Committee (virtual)
Apr 30	Finance Committee (virtual)
July 3–4	Council (virtual)
Oct 20–21	Technical Advisory Committee (virtual)
Oct 27–28	Finance Committee
Oct 28	Policy Advisory Committee
Dec 4–5	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.

Newsletter | **No. 181** | Autumn 2024

European Centre for Medium-Range Weather Forecasts

www.ecmwf.int