

Newsletter

No. 179 | Spring 2024

A cold winter in northern Europe

Forecast performance 2023

Ocean-wave-related update

A time-varying tropospheric aerosol
climatology

Medicane Daniel

Introducing earthkit

© Copyright 2024

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

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

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

Publication policy

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

The ECMWF Newsletter is not peer-reviewed.

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

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

Longer ranges

In addition to providing medium-range weather forecasts up to 15 days ahead, which is the main purpose of ECMWF, we have long produced extended-range forecasts up to 46 days ahead, and seasonal forecasts up to 13 months ahead. It is important to understand the applicability of such forecasts: they cannot produce the details of future weather conditions on a day-to-day basis, but they can provide projected deviations from average conditions over the weeks and months ahead. This makes them particularly useful to sectors such as agriculture, water management, energy and health, among others. Just as our medium-range forecasts are subject to continual improvement, so are longer-range forecasts. For example, in last year's update of our Integrated Forecasting System (IFS) to Cycle 48r1, we substantially reorganised our extended-range forecasts: they now have 101 instead of 51 ensemble members, they run daily rather than twice-weekly, and they have a consistent horizontal resolution of 36 km. The seasonal forecasting system, which is currently SEAS5, is updated less frequently, at intervals of four to six years. The next system, SEAS6, is due to be introduced next year, and active preparations for this upgrade are under way.

An example of the use of our extended-range forecasts can be seen in this Newsletter, in the article on a cold winter in northern Europe. The fact that seasonal forecasts are the subject of continuous research becomes clear in the articles on improving the boundary forcing in reanalyses and seasonal forecasts, and in the feature article on a new time-varying tropospheric aerosol climatology for the IFS. The latter article also makes clear the intimate connections between ECMWF weather forecasts and the work carried out by the EU-funded Copernicus Atmosphere

Monitoring Service (CAMS), implemented by ECMWF. It describes how a time-varying climatology will address two shortcomings: the current system is unable to represent changes in anthropogenic aerosols, and it is incompatible with the representation of aerosols in CAMS. There is also an article on the new Copernicus Interactive Climate Atlas, provided by the EU-funded Copernicus Climate Change Service (C3S) implemented by ECMWF. This includes climate projections from a range of sources.

Our cooperation with the World Meteorological Organization (WMO) is illustrated by an article on supporting the new WMO Incident Management System and by a report on a WMO fellowship recently completed at ECMWF. Additional feature articles describe recent updates related to ocean wind waves in ECMWF's Earth system model, which will be implemented in IFS Cycle 49r1 later this year; they provide a detailed look at our forecasts of extreme rainstorms in Greece, Bulgaria, Türkiye and Libya in September 2023; and they present a new open-source Python development, called earthkit, providing powerful tools to speed up weather and climate science workflows. Overall, we aim for seamless predictions across the different forecasting ranges to maximise the usefulness of our products.

Florence Rabier
Director-General



Contents

Editorial

Longer ranges 1

News

A cold winter in northern Europe 2

Forecast performance 2023 4

Predicting the onset of monsoon winds over the Maldives . . 6

New observations January – March 2024 7

Supporting the WMO Incident Management System. 8

Fuel for the fire 9

A new observation dashboard 10

Improving the support to users with a new forum service. . 12

Improving the boundary forcing in reanalyses and seasonal forecasts 13

EUPPBench: A forecast dataset to benchmark statistical postprocessing methods. 15

Copernicus Interactive Climate Atlas: a new tool to visualise climate variability and change 16

Earth system science

Ocean-wave-related changes in the next model upgrade . . 18

A new time-varying tropospheric aerosol climatology for the IFS. 26

Medicane Daniel: an extraordinary cyclone with devastating impacts 33

Computing

Introducing earthkit 48

General

ECMWF publications 54

ECMWF Calendar 2024 54

Contact information 54

A cold winter in northern Europe

Linus Magnusson

While the world as a whole experienced an unprecedentedly high average temperature in 2023, Northern Europe was one of the few areas that saw below-average temperatures compared to 1991–2020. Especially the autumn and the winter were dominated by cold anomalies, as seen in the composite of weekly anomalies in ECMWF’s ERA5 reanalysis from 2 October 2023 to 25 February 2024 (see the top panel in the first figure). For the evaluation in this article, we will focus on land points in a box covering most parts of Finland and the north-eastern corner of Sweden (see the box in the third figure, 60–67°N and 20–30°E).

Prediction of mean anomalies

The first figure shows composites of extended-range ensemble mean predictions of weekly mean 2-metre temperature anomalies for weeks starting on 2 October to 18 February, based on Monday forecasts. The anomalies over Europe were well captured in week-2 forecasts. A cold anomaly over Northern Europe was present on average in the week-3 forecast (days 15–21), but with lower amplitude (not shown). A longer forecast (week-4) had the warm signal present for continental Europe but was neutral for northern Europe. It still predicted that northern Europe would be colder than central Europe in terms of anomalies.

Predictions of week-to-week variability

The second figure shows week-to-week variations of extended-range ensemble mean predictions and ERA5 analyses for the box over Finland/Sweden. Here we find negative anomalies in ERA5 for 17 out of 21 weeks (3 of the warm weeks appeared after mid-January). For the

Two-metre temperature anomalies.

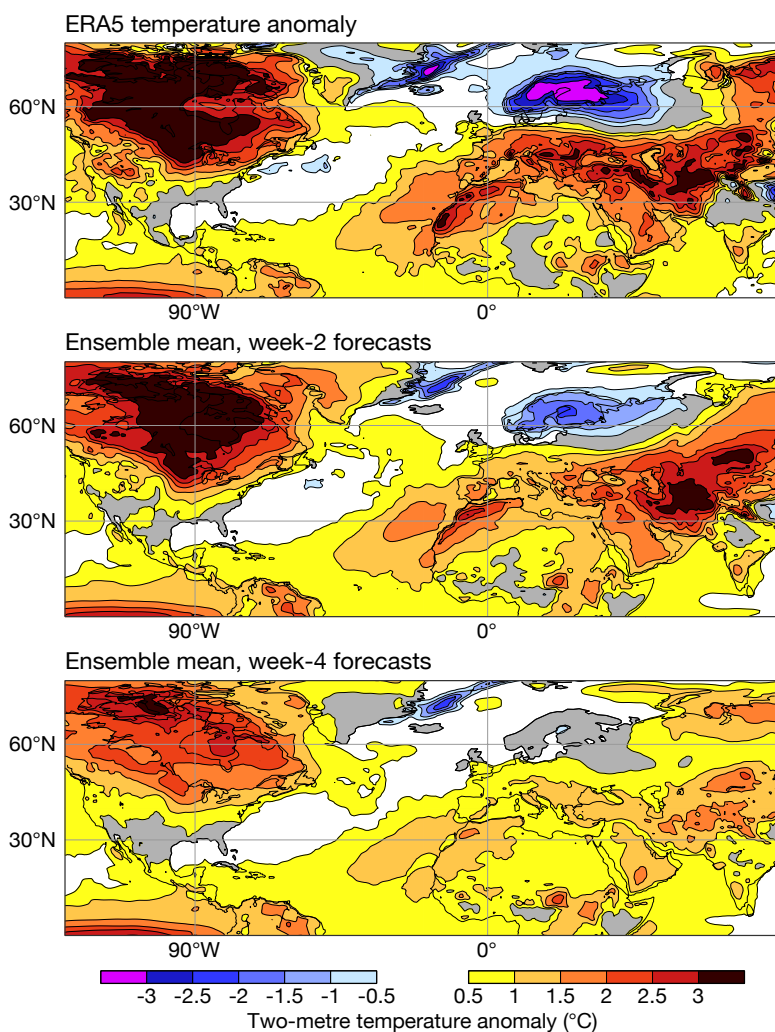
Composites of weekly 2-metre temperature anomalies for the weeks commencing 2 October 2023 to 18 February 2024 from the ERA5 reanalysis (top), the extended-range ensemble mean from week-2 forecasts (middle), and the extended-range ensemble mean from week-4 forecasts (bottom).

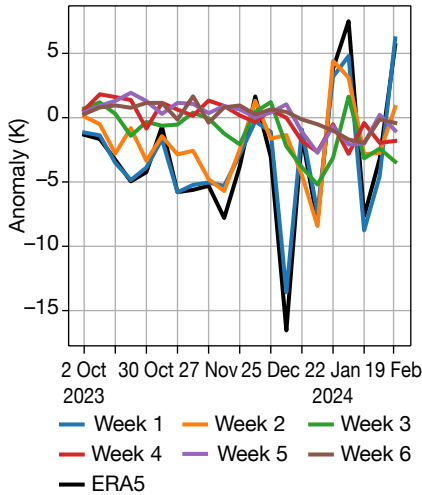
week-4 forecasts over the region, most of the forecasts valid before New Year predicted a slight warm anomaly, and all forecasts afterwards predicted a cold anomaly. This resulted in a negative correlation with the outcome (−0.3). A similar signal is found for week 5 and 6 (not shown). Here we can note that the seasonal forecast from 1 October predicted a warm anomaly (above +1°) for November–February. The week-3 forecasts mainly predicted cold anomalies in January, which could be related to a sudden stratospheric warming at the beginning of January. The week-3 forecast gave a false alarm for cold weather in the week starting on 15 January. However, for the following week the week-3 forecast captured the warm anomaly by persisting the signal from the week-2 forecast. The overall correlation with the outcome for week-3

forecasts was 0.22. The week-2 forecast captured well the week-to-week variability (correlation of 0.65) except for the coldest week of the period, starting on 1 January, which had a weekly anomaly of −16° in the box in the ERA5 analysis.

Problematic forecast for beginning of January

To understand the (lack of) skill for the first week of January further, we will focus on a 3-day period (3–5 January). During these days, the coldest temperatures since 1999 were measured in both Sweden and Finland (−44.6°C and −44.3°C), and two stations in Sweden had a 3-day mean temperature below −40°C (2–4 January). According to the ECMWF analysis of 2-metre temperature, large parts of central Finland had a 3-day mean temperature





Weekly anomalies of 2-metre temperature over parts of Finland and north-western Sweden. The chart shows weekly anomalies in the ERA5 reanalysis and according to ECMWF extended-range ensemble mean forecasts one to four weeks ahead, over land points inside the box in the third figure.

nudging technique pointed to short-range errors on the downstream side of a trough over the central U.S. These affected the subsequent development of storm Geraldine, which hit France. Once the uncertainties in this development were overcome, medium-range predictions captured the extreme cold.

Prediction of local 2-metre temperature

Predicting very cold temperatures during calm conditions has been a long-standing challenge for ECMWF's Integrated Forecasting System (IFS) as well as other models. It is believed that this issue is due to a combination of several model deficiencies that need to be tackled separately. In IFS Cycle 48r1 (implemented in June 2023), a multi-layer snow model was introduced that improved temperatures during clear conditions over snow. Looking at the spatial pattern of the

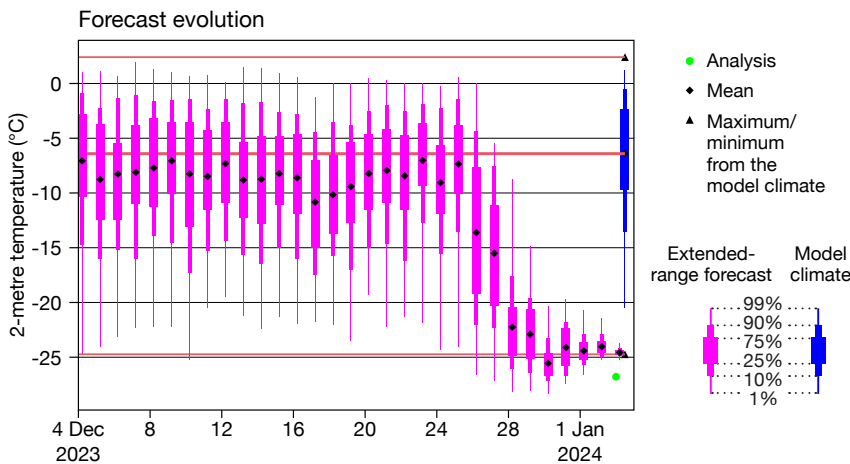
below -30°C (see the left-hand bottom panel of the last figure). Looking at the forecast evolution of 2-metre temperature inside the box (top panel of the last figure), based on the daily extended-range ensemble, the signal for the extreme cold appeared between 26 and 28 December (the same was found for the medium-range ensemble, not shown). The day-6 forecast from 25 December as well as earlier forecasts included large errors over Europe. Predictability experiments using the

9 km control forecast averaged over the first 72 hours for 3 January 00 UTC (bottom right panel of the last figure), we see that the forecast underestimated the coldest extreme, such as the region where temperatures were below -30°C (brown shading). As reported in ECMWF Newsletter No. 178, in the upcoming Cycle 49r1 planned for late 2024, several changes are targeting 2-metre temperature. Two-metre temperature from SYNOP weather station observations will be used in the atmospheric data assimilation, while it is already used for the land data assimilation. This, together with changes related to the representation of the land surface and the interpolation to 2-metre temperature, is expected to improve the 2-metre temperature for the first forecast day(s).

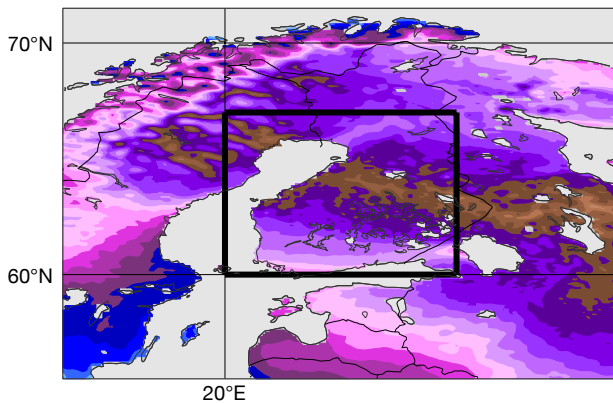
Summary and outlook

In this article, we have discussed the prediction of cold temperatures in northern Europe during the winter 2023/24 on different timescales. During this period, the predictability of week-to-week anomalies did not reach much further than 2–3 weeks. This could be because the seasonal anomaly was due to a 'driver of predictability' that was not captured by the model, or because the cold season was a result of a series of synoptic events that led to a colder than normal season. In any event, more research is warranted on the interface between medium-range and sub-seasonal predictability.

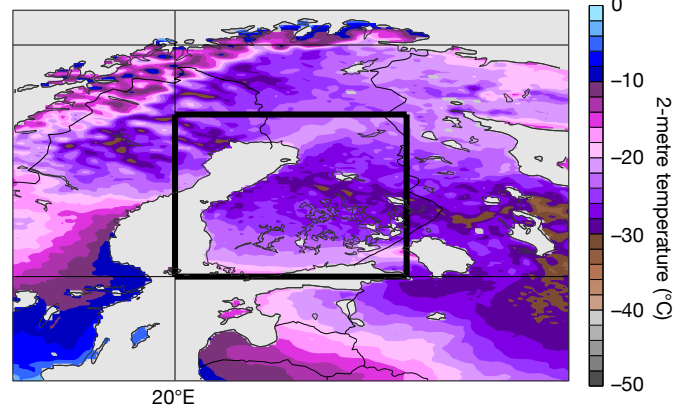
Predictions of 2-metre temperature 3–5 January. Forecast evolution plot for 3-day (3–5 January 2023) 2-metre temperature over parts of Finland and north-eastern Sweden (the box shown in the maps) (top), the analysis (bottom left), and a 0–3 day ensemble (ENS) control forecast (bottom right) of 2-metre temperature for the same dates.



Analysis



0–3 day ensemble control forecast



Forecast performance 2023

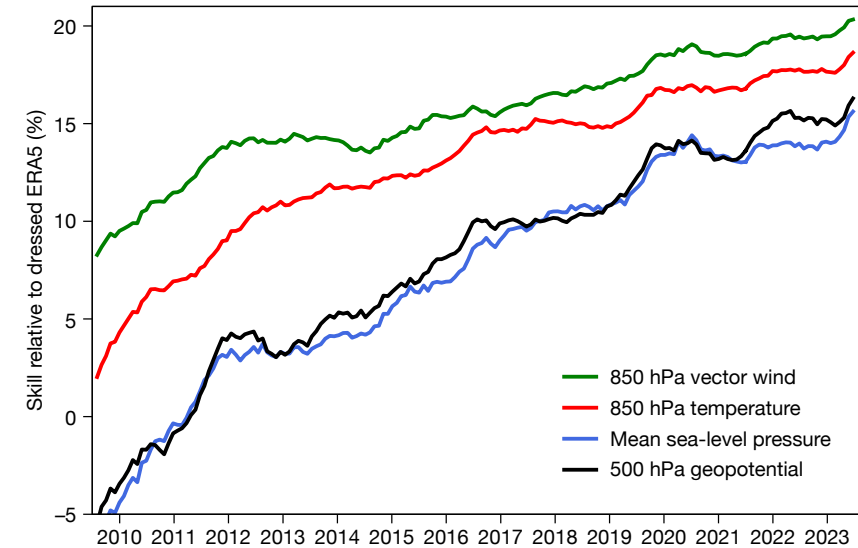
Thomas Haiden, Matthieu Chevallier

ECMWF maintains a comprehensive range of verification statistics to evaluate the accuracy of its forecasts. Each year, a summary of verification results is presented to ECMWF's Technical Advisory Committee (TAC). Their views about the performance of the operational forecasting system in 2023 are given in the box.

With the implementation of Cycle 48r1 of ECMWF's Integrated Forecasting System (IFS) on 27 June 2023, the horizontal resolution of medium-range ensemble forecasts (ENS) has increased from 18 to 9 km. This is the same resolution as the current high-resolution forecast (HRES). Cycle 48r1 also brought changes to extended-range ensemble forecasts, notably an increase in members from 51 to 101, and many improvements to the data assimilation system as well as the forecast model. The largest forecast skill improvements are associated with the ensemble forecast because of the ENS resolution upgrade. ENS scores of surface variables are improved by 2% to 6%, upper-air variables typically by around 1% to 3%.

The Cycle 48r1 resolution increase improves ENS tropical cyclone track and intensity forecasts, with position errors reduced by up to 10% and core pressure errors reduced by around 20%. Tropical cyclone track forecasts are improved because the 9 km model has a reduced translation speed bias. Intensity forecasts are improved because the higher horizontal resolution allows for a better representation of the strong horizontal gradients associated with tropical cyclones. This is also reflected in a strongly increased tropical cyclone intensity spread of the ENS, leading to a better spread-error relationship.

Compared to other centres, ECMWF has been able to consolidate its overall lead for upper-air parameters in the medium range thanks to the new Cycle 48r1. For surface parameters, such as 2 m temperature and precipitation, some other centres have, however, taken the lead in the short range. Nevertheless, for both 2 m temperature and 10 m wind speed, the number of large ENS errors has further decreased. Cycle 49r1, which



Upper-air skill of the ENS relative to forecasts based on the ERA5 reanalysis system.

Skill of the ENS for upper-air parameters at day 5 in the northern extratropics, relative to the skill of a Gaussian-dressed ERA5 forecast. Values are running 12-month averages, and verification is performed against own analysis. Due to the averaging, only about half of the full impact of Cycle 48r1 is visible (data until the end of December 2023 is included).

is scheduled for implementation in autumn 2024, will bring substantial improvements near the surface due to the assimilation of 2 m temperature over land as part of the 4D-Var data assimilation system and the introduction of the SPP scheme for model perturbations, among other changes. In terms of forecast skill for ocean wave parameters, ECMWF is leading compared to other global centres both in peak period and significant wave height.

ECMWF's new experimental in-house machine-learning (ML) model AIFS (Artificial Intelligence/Integrated Forecasting System) improves on the HRES in some aspects, and its forecasts are competitive with the best from external machine-learning models. All of these models (both external and in-house) are trained on ERA5 reanalysis data.

The increase in ensemble size from 51 to 101 of the extended-range forecast has led to improvements in forecast performance across all parameters over the forecast range from week one to week four. The change in forecast frequency from bi-weekly to daily has further increased the effective utility of the forecast. A major challenge for

extended-range prediction is still posed by transitions in the large-scale flow pattern that lead to marked changes in temperature anomalies, such as the switch from warm to cold anomalies in Northern Europe in October 2023.

ECMWF's seasonal forecast provided very good predictions of the transition early in 2023 from La Niña to El Niño conditions, even on the long-range (one year) timescale. The exceptionally warm northern hemisphere autumn season 2023 was well captured. However, the cold anomaly in northern Europe was missed.

Note that the TAC's assessment of forecast performance in the box below is based on evaluation results that were available by September 2023, so it does not reflect the improvements due to Cycle 48r1 which became visible in the scores after that date.

The complete set of annual verification results is available in ECMWF Technical Memorandum No. 911 on 'Evaluation of ECMWF forecasts, including the 2023 upgrade', downloadable from <https://www.ecmwf.int/en/publications/technical-memoranda>.

The following are other sources of

information about verification and forecasting system changes.

- Verification as part of ECMWF's charts page:
<https://charts.ecmwf.int>
- World Meteorological Organization (WMO) inter-comparison of global model forecast skill:

<https://wmo.cdn.ecmwf.int>

- WMO ocean wave model inter-comparison results:
<https://confluence.ecmwf.int/display/WLW/WMO+Lead+Centre+for+Wave+Forecast+Verification+LC-WFV>
- List of 'Known IFS Forecasting

Issues':

- <https://confluence.ecmwf.int/display/FCST/Known+IFS+forecasting+issues>
- IFS cycle changes since 1985:
<http://www.ecmwf.int/en/forecasts/documentation-and-support/changes-ecmwf-model>

Assessment of ECMWF's Technical Advisory Committee, 19–20 October 2023

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

- congratulated ECMWF on the successful implementation of 48r1, the first model cycle to be implemented at the new Bologna facility;
- welcomed early indications showing 48r1 having generally positive impacts on verification scores whilst recognizing that the impact of 48r1 is generally not yet reflected in many of the scores and metrics used in this assessment summary;
- welcomed that some of the largest improvements in the 48r1 scorecards are for the ENS where the change in horizontal resolution to match HRES has brought positive impact, validating the decision to increase resolution;
- congratulated ECMWF's change in the extended forecast running daily at 48r1, noting this provides additional lead time for major events and transitions in the weeks 2 to 4 period;
- recognized that in the absence of a new model upgrade in 2022 due to the move to Bologna, changes in scores often reflected interannual variability;
- noted that in spite of there being no model upgrade in 2022, ECMWF maintained their lead over other centres in many metrics, especially upper air;
- noted that the magnitude of the lead between ECMWF and other centres in the upper air has decreased for some scores, for example 850 hPa temperature CRPSS and verification against radiosondes in the tropics, but increased for others, for example 500 hPa anomaly correlation scores;
- noted that outside of summer, ECMWF now lags behind some other centres for short-range scores in 2 m temperature and precipitation over Europe; this is thought to be linked to challenges in cold, stable airmasses and handling of light precipitation and other centres improving their own outputs;
- acknowledged that regional biases in surface weather parameters are complex, for example in 2m temperature and dewpoint biases around Europe, and welcomed ongoing research here;
- noted an increased bias in cloud cover introduced at 47r3 is likely to be addressed by 48r1 whilst a night-time continental positive bias in 10m wind scores is likely to be addressed at 49r1;
- acknowledged that 47r3 has resulted in improved 850 hPa wind in the tropics but 850 hPa temperature here has degraded compared to analysis. Meanwhile, improvement in net short-wave radiation in the tropics has led to the smallest errors in this parameter ECMWF has ever had;
- acknowledged that ECMWF maintained its long-term lead over other centres for significant wave height and its more recently developed lead for peak period scores;
- noted some of the smallest errors so far in the handling of tropical cyclones, although there is still a slow bias for tropical cyclone propagation. Early indications suggest 48r1 offers some general improvements above 47r3;
- noted that changes in some EFI verification scores, for example a drop in scores for precipitation, relative to last year are considered related to interannual variability rather than deficiencies in the forecast system;
- acknowledged that improvements in CAMS following an update in 2022 have led to a return to previous verification scores. Noted, also, that the meteorological performance of some outputs in the Tropics performs comparably to HRES even though the resolution of CAMS here is 40 km;
- noted that extended-range parameters continue to show improvements over climatology, more so at week 2 than weeks 3 and 4, and the highest ever real-time forecast precipitation skill was registered. Improvements for warm summer anomalies are greater than for cold winter anomalies, reflecting inherent forecast challenges;
- acknowledged that in spring 2022 ECMWF seasonal output was too fast, with a transition from La Niña to neutral conditions, but by autumn 2022 ECMWF seasonal output captured the transition from La Niña to El Niño well. This is important given tropical Pacific SST [sea-surface temperature] anomalies are the biggest driver of global seasonal variation;
- noted that the seasonal forecast for winter 2022–23 showed good skill in capturing the anomalously warm winter in Europe, capturing the idea of the warm anomaly being greater in northern and eastern Europe. Signals elsewhere in the northern hemisphere were also good even though the magnitude was not quite captured, for example over North America;
- noted that the seasonal forecast for summer 2023 captured the warm European summer anomaly but did not capture the blocking transition that led to a shift in the location of the largest anomalies over the course of the summer. Further improvement in identification and magnitude of heat waves in the weeks 2 to 4 period would be welcomed;
- appreciated the continued development of new diagnostics and very good support ECMWF provided to Member and Co-operating States over the last year, with engagement via many mechanisms including online support, the annual UEF [Using ECMWF's Forecasts], online seminars, new format of site visits and meteorological representatives at Member States.

Predicting the onset of monsoon winds over the Maldives

Ahmed Shabin (Maldives Meteorological Service)

Having completed a World Meteorological Organization (WMO) fellowship at ECMWF, it is a pleasure to look back and share my reflections. The fellowship scheme between the WMO and ECMWF is part of a wider WMO Fellowship Programme that provides specialised training placements to develop capacity in least-developed and developing countries.

The one-year fellowship at ECMWF provided firsthand experience of the day-to-day operations and the final stages of the roll-out of the new Integrated Forecasting System (IFS) Cycle 48r1 in June 2023. It also provided the opportunity to take part in training courses at the Centre and in forums with experts from across the globe.

Working closely with the Forecast

Department, a research project was formulated to improve the retrieval and processing of model data and expand computer programming skills. The study focused on assessing the forecast skill of the IFS in predicting the onset of monsoon winds over the Maldives area.

The monsoon onset is declared by the Maldives Meteorological Service (MMS) based on observations against a set of criteria. These criteria include the reversal of the lower-level zonal (east–west) wind direction with a sustained average wind speed of 10 knots or more for two consecutive days over the three regions of the country. The southwest monsoon and the northeast monsoon are triggered by the propagation/withdrawal of the Intertropical Convergence Zone (ITCZ) over the country around the first week of May and in mid-to-late December,


		Onset region		
		South	Central	North
Area Hovmöller	1,000 hPa	78%	57%	48%
	850 hPa	83%	70%	65%
Line Hovmöller	1,000 hPa	74%	65%	52%
	850 hPa	91%	70%	57%

ERA5 hit rate compared to declared dates. Percentage hit rate of ERA5 wind composites for 22 years compared to the monsoon declared dates, based on the declaration criteria.


respectively. The onset of the southwest monsoon over the Maldives is the first indication of the onset of the greater Asian monsoon over the region. Therefore, increasing the forecasting skill of the MMS would have regional benefits as well.

How good is ERA5 as a proxy for observations?

The first question for the study was to analyse how to process data from ECMWF’s ERA5 reanalysis as a proxy for observation data and how close this was to the ground reality. To address this, the ERA5 data was processed and plotted using area and line Hovmöller averaging, each for two levels (1,000 and 850 hPa). As the country is centered along 73°E, a narrowing longitude band starting from 60–70°E was tested. The longitude band 75–70°E was chosen for the Hovmöller averaging as a reasonable representation whilst allowing for enough data to go into the computation. The parameter showing the resultant wind (speed) and the zonal wind sign was computed to reflect the monsoon onset criteria. These composite figures were then analysed against the past 22 years of southwest monsoon-declared dates for three regions of the Maldives by the MMS (see the first figure). The analysis looked at how many of these cases for each region satisfied the onset criteria,



Hovmöller chart. Area-averaged Hovmöller chart of ERA5 resultant wind speed and zonal wind direction for the southwest monsoon onset period of 2020. Red, green and magenta lines show the onset dates declared by the MMS for southern, central and northern regions, respectively. The year 2020 demonstrates an example of when ERA5 and monsoon onset dates corresponded quite well. Across the years analysed, the results were a mixed range of perfect match and erratic.



Predictability of monsoon winds onset.

The predictability of the onset of monsoon winds over the Maldives region is shown by the Brier Score, computed between IFS hindcasts and the ERA5 reanalysis for both monsoon periods from 2003 to 2022. The left-hand plot is for the southwest monsoon and the right-hand plot for the northeast monsoon. The two horizontal lines delimit the three regions in the Maldives (south–central–north).

in other words the hit rate of ERA5 compared to formal monsoon onset dates declared by the MMS.

For area Hovmöller averaging at 1,000 hPa, over 70% of ERA5 hits were in agreement for the southern region, and 57% and 48% for the north and central regions, respectively (see the table). The difference in values across the regions can be attributed to the fact that the southwest onset declaration is strongly focused on the southern region, being the first area of onset over the country (as the ITCZ moves northwards); difficulties in clearly defining the regional boundaries; plus evolving practices over the two decades. As the results for the line Hovmöller averaging were fairly similar, for simplicity, the study continued using ERA5 area-averaged Hovmöller at 1,000 hPa as a good proxy for the observations. A similar study was not possible for the northeast monsoon due to the lack of long-term onset declaration dates for the monsoon.

How far ahead can the IFS predict monsoon winds?

To address this question, medium range ensemble (ENS) reforecast data

of the IFS from 2023 (hindcasts consisting of ten perturbed members and a control member) was retrieved for a two-month period centred on average monsoon onsets for the 20-year period of 2003–2022. The data was processed as area-averaged Hovmöller of the probability of derived parameters consisting of both the resultant wind (speed) and the sign of the zonal wind. A corresponding dataset for each hindcast run was prepared from the reanalysis data. Based on these two datasets, a Brier Score was computed to assess the accuracy of the ENS hindcast for the past 20 years.

While a Brier Score of 0 represents perfect accuracy in probabilistic predictions, a score of 1 indicates that the predicted forecasts are as far from the observation as possible, in this case, the ERA5 values. The analysis for southwest monsoon winds showed an overall Brier Score of less than 0.4 across the 15-day forecast window and a Brier Score of less than 0.25 up to 7 days ahead (see the left-hand panel of the second figure). This indicates that the IFS model has good predictive skill in forecasting the onset of southwest monsoon winds

over the Maldives region.

The analysis for northeast monsoon wind onset showed an overall Brier Score of less than 0.3 across the 15-day forecast window and a Brier Score of less than 0.2 up to 10 days ahead (see the right-hand panel of the second figure). This indicates that the IFS model has even better predictive skill in forecasting the onset of northeast monsoon winds over the Maldives region. With this information, during the final stage of the study, a model product was designed for the Maldives region. The product is intended to inform forecasters of the percentage of ECMWF ensemble members from the real-time forecasts that fulfill the monsoon wind onset criteria across the longitude band of 70–75°E from 7°N to 1°S. In addition to forecasting monsoon wind, the product can predict wind burst episodes related to large-scale variability, such as the Madden-Julian Oscillation.

The fellowship has undoubtedly opened up a new avenue for the MMS to maximise its uptake of information and guidance from ECMWF, whilst also forging closer cooperations through initiatives such as UN Early Warning for All.

New observations January – March 2024

The following new observations have been activated in the operational ECMWF assimilation systems during January – March 2024.

Observations	Main impact	Activation date
NOAA-21 CrIS radiances	Temperature, humidity, dynamics	16 January 2024
SPIRE GNSS-RO bending angles from NOAA	Temperature and winds in upper troposphere/lower stratosphere	24 January 2024
Wave height data from Sentinel 3A/3B/6A, Cryosat-2 and SARAL	Wave analysis/forecasts and atmospheric boundary layer over the ocean	6 February 2024

Supporting the WMO Incident Management System

Daniel Varela Santoalla

In 2020, ECMWF received a request from the World Meteorological Organization (WMO) to support the establishment of an Incident Management System (IMS) for Regional WMO Integrated Global Observing System (WIGOS) Centres (RWCs). This system was required to manage the observational incidents the RWCs could find within the network of stations under their responsibility. Identification of observational issues, which can concern station metadata (for example incorrect coordinates/altitudes), or data communication/coding, or other reasons, is made by RWCs primarily via the WIGOS Data Quality Monitoring System (WDQMS) webtool (<https://wdqms.wmo.int>). This is hosted and operated by ECMWF.


ECMWF's role

ECMWF has been running the Atlassian Jira software as an issue tracking and incident management system since 2011. During this period, we have accumulated expertise in configuration strategies and operational management of the tools provided. Having been granted a so-called 'open-source licence' by Atlassian meant that we could afford to use the tools widely without incurring any licensing costs until February 2024.

After the WMO requested it, ECMWF started supporting the WMO in establishing a prototype for a RWC IMS, leveraging our experience in the field with the requirements drawn up by the WMO. Finally, in 2023, the WMO decided to establish their own Jira system. With technical support from ECMWF staff, the IMS for RWCs was migrated to a WMO cloud-based solution, which became operational on 1 February 2024.

Three phases

The collaboration consisted of three distinct phases. During the first, we worked together to design a system that would meet the project's specific requirements. This involved understanding the needs and the roles



WMO RWC Incident Management Workflow. The chart shows what happens to observational incidents raised by RWCs.

of RWC and National Focal Points and mapping all those into Jira's technical features, particularly user roles, screens, workflows, priority levels, and custom fields. Roles and the permissions based on them needed to be defined precisely so that only certain users could perform some actions, like opening and closing incidents or commenting on them.

The second phase consisted of the implementation of these requirements and features in an actual Jira project hosted in ECMWF's infrastructure. This implementation required several iterations with a few internal WMO test users until the design was validated and deemed appropriate for broader testing. This more comprehensive testing required training sessions with members of some specific National Focal Points and Regional WIGOS Centers of some WMO regions to validate the solution and understand any further user requirements that could be of any value for the continuous improvement of the system. One of these sessions, delivered as a webinar, included the National Focal Points of WMO Region III (South America) and the Regional WIGOS Centres for this region, which are Argentina and Brazil. ECMWF participated with the WMO in the delivery of this webinar.

One of the key aspects demonstrated in the webinar was the

meaning of the workflow statuses and the conditions under which incidents could be transitioned from one status to another. The figure shows a chart displaying statuses and possible transitions.

Once the solution was validated, the third phase of this project involved pre-operational running of the solution on the existing infrastructure at ECMWF. This pre-operational phase was necessary to validate the solution further and to give WMO staff some time to acquire the necessary expertise to build an infrastructure that could be used to deploy the solution as designed and tested in the previous phases.

Migration to the WMO

Once this infrastructure was ready to take over operations, all data collected during the pre-operational phase was migrated from the ECMWF infrastructure to the newly built WMO infrastructure. This migration required a good degree of coordination as the data transfer methodology employed required that both origin and destination systems were running the same version of the Jira software. All the hard work came to fruition on 1 February 2024, after which all Regional WIGOS Centres and National Focal Points were instructed to move to the newly deployed system at <https://jira.wmo.int>.

Fuel for the fire

Joey McNorton, Francesca Di Giuseppe

Landscape fires impact local ecosystems, communities, air quality and global atmospheric conditions, often with devastating consequences. Effective wildfire management and prevention strategies depend on accurate forecasts of fire danger. The backbone of the early warning component of the EU's Copernicus Emergency Management Service (CEMS), and the most used index by forest agencies and meteorological services in Europe, is the fire weather index (FWI), which was developed in the 1970s. ECMWF has been producing FWI forecasts since 2017 as part of its duties as the computational centre for the CEMS fire activity. The FWI uses atmospheric conditions and effectively estimates the moisture content of dead vegetation. A major limiting factor to this approach is the lack of consideration of vegetation abundance or type. Furthermore, as the FWI was trained in specific conditions, Canadian pine forests, it has been shown to have limitations when applied at a global scale. As part of ECMWF's continuing effort to broaden the use of our Integrated Forecasting System (IFS) into the field of environmental forecasting, specifically to provide more accurate fire forecasts, a new global fuel characteristic model and a ten-year dataset have been developed (see <https://doi.org/10.5194/bg-21-279-2024>). These combine meteorological and land surface variables from the IFS with satellite observations. This newly developed mid-complexity model is part of an overall ambitious plan to develop a fire module – SPARKY – for the IFS. In its present form, the fuel component of SPARKY provides real-time daily forecasts of fuel composition and state with a grid spacing of 9 km.

Seeing the wood for the trees

The most notable limitation of existing fire danger metrics used in early warning systems is the lack or simplification of fuel availability. For example, a static fuel map may be used. At the climate scale, this is overcome with the use of dynamic vegetation models, although these

lack consensus at fine spatial resolutions and fail to account for the impact of human-induced disturbances on the landscape.

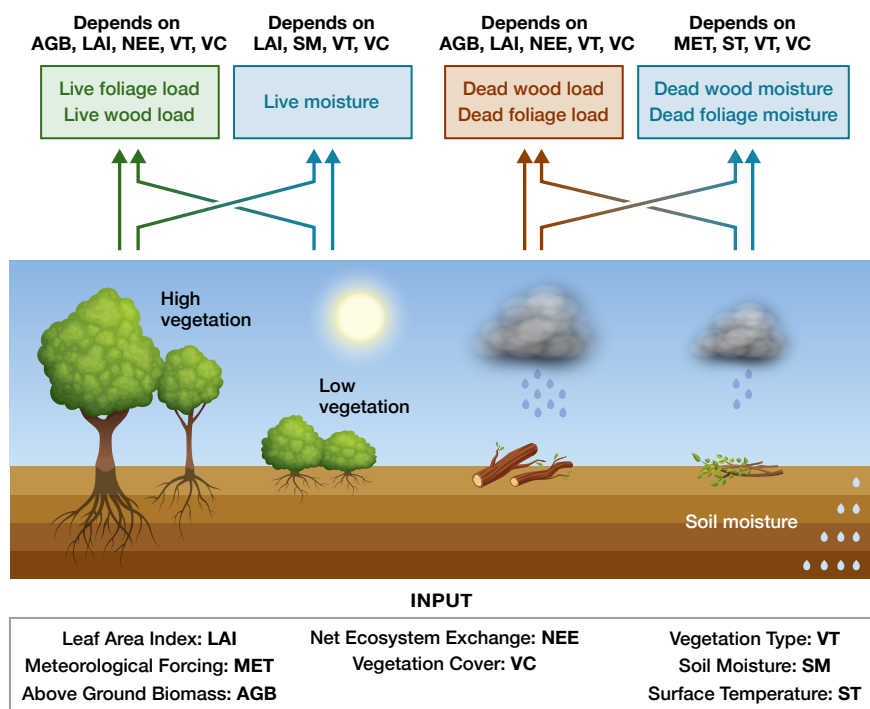
A new fuel model developed at ECMWF is initialised from satellite observations of above-ground fuel, taken from a multi-year European Space Agency – Climate Change Initiative (ESA CCI) product. To provide daily fuel updates, the model incorporates variations in vegetation amount output from the land surface component of the IFS, ECLand. As ECLand has some biases, the model employs a bias correction scheme, which is used by the EU's Copernicus Atmosphere Monitoring Service (CAMS) implemented by ECMWF, to provide more accurate carbon fluxes. To account for changes in total fuel mass, and not just the carbon component, the fuel model is further bias-corrected using the difference in above-ground fuel from ESA CCI from consecutive years. This results in daily estimates of total fuel load.

The next step is to categorise the fuel load into live wood, live leaves, dead

wood, and dead foliage mass. This is done by considering the type of vegetation used in ECLand and using satellite estimates of the vegetation state, using a variable known as the leaf area index (LAI). Using these methods, the model estimates that 78% of the above-ground fuel globally is alive, whilst the remaining 22% is dead, although this varies depending on location and season. Of the living component, 97% comprises woody fuel and 3% is leafy fuel. The dead fuel is approximately equally weighted between wood and foliage, where the foliage includes small twigs and branches.

The moisture state

Although existing fire danger metrics make estimations of the moisture content of dead vegetation, these are often simplified and do not extend to the state of living fuel. The newly developed ECMWF fuel model estimates both live fuel moisture content (LFMC) and dead fuel moisture content (DFMC), further dividing the DFMC between woody and foliage fuel.



Fuel characteristics. Schematic of the fuel characteristic model showing the required input and the output variables.

Estimations of LFMC are made using a combination of observed vegetation state, LAI, modelled soil moisture, and mapped vegetation types. These are combined to give an estimated LFMC, which sums three components of vegetation hydrology: winter moisture content is when dormant vegetation is dependent on soil moisture alone; drought resistant moisture content is dependent on the LAI alone; the combined active season moisture content is estimated by combining soil moisture and LAI. The model was optimised with a theoretical maximum LFMC per vegetation type, using observations taken from vegetation samples in the field.

A well-tested DFMC model, known as

the Nelson Model, has historically been used to estimate the wetting and drying rates of a specific type of fuel. The Nelson approach utilises air temperature, humidity, radiation, and precipitation to derive heat and moisture transfer of the fuel. The new ECMWF fuel model generalises the Nelson equations to make it applicable to multiple fuel types, depending on the drying rate of the fuel. The partitioning between wood and foliage DFMC is dependent on the vegetation type defined by ECLand.

Fanning the flames

The fuel component of the SPARKY model provides a crucial first step toward disentangling the contributing

factors of fuel and weather to understand fire evolution globally. By combining the fuel component with meteorological data and land cover information, the aim is to further improve fire forecasting and attribution. One of our goals, using machine learning, is a forecast at 1 km resolution of the 10-day probability of fire. Further avenues will be explored in collaboration with the CEMS service and the forest fire expert group of the ARISTOTLE-ENHSP (All Risk Integrated System Towards Trans-boundary holistic Early-warning — European Natural Hazards Scientific Partnership). Our ambition is to define new, innovative fire weather indices and move the service forward to the new challenges ahead.

A new observation dashboard

Mohamed Dahoui, Paolo Battino

Observations are essential for numerical weather prediction (NWP) systems. They are used by the data assimilation system to produce the best estimate of the initial conditions (analysis) for subsequent forecasts. The observing system has grown substantially in the past two decades and continues to evolve rapidly. At the time of writing, ECMWF receives over 700 million observation reports per day, of which approximately 40 million are used. The observations originate from approximately 100 data sources, ranging from individual satellites to in-situ observation types.

Checking the health and usage status of the observing system is of interest to a wide range of users, from scientists working directly with observations to data providers and managers. ECMWF has developed a new operational observation dashboard providing near-real-time quick access to the status of the different components of the observing system (<https://obsstatus.ecmwf.int/>). The dashboard design offers a high-level view of the observing system for non-expert users, and a more detailed view for expert users.

Observation dashboard design and features


The observation dashboard is

intended to serve the needs of a wide range of expert and non-expert users thanks to featuring a cascading hierarchy of the observing system structure. It also uses a traffic light system to indicate the absence or presence of availability and quality anomalies detected by ECMWF's automatic data checking system (<https://www.ecmwf.int/en/newsletter/174/earth-system-science/use-machine-learning-detection-and-classification>). The dashboard offers a link to the recently implemented on-demand web monitoring system (<https://www.ecmwf.int/en/newsletter/178/earth-system-science/demand-web-plotting-observation-monitoring-statistics>). The design of the system enables the following features:

- The full observing system monitored by ECMWF is supported.
- New observations will automatically appear in the appropriate location within the observing system structure. Obsolete observations (no reports for more than 30 days) are automatically removed from the dashboard. This procedure enables the system to function with minimal maintenance.
- The status is updated twice a day to reflect the usage status and

anomalies affecting the two main 4D-Var analysis cycles. The status of ocean observations is updated once a day.

- The dashboard provides a history of the last few days to allow users to monitor the evolution of anomalies.
- Non-used observations are supported by the system. The associated icons are, however, greyed out to highlight the passive nature of these observations. Although passive observations are not supported by the automatic anomaly detection system, the dashboard will highlight cases when the data are temporarily not received by ECMWF.
- Users interested only in investigating data affected by anomalies have the option to hide channels not affected by warnings.
- For all components of the observing system, the dashboard offers direct links to time series from ECMWF's on-demand web monitoring system.
- For components of the observing system affected by anomalies, the dashboard provides a short description of the anomaly and a direct link to the time series



Dashboard processing framework. The diagram shows the dashboard system architecture with an image of the initial screen seen by users.

produced by the automatic anomaly detection system.

The implementation of this dashboard will significantly improve the efficiency of monitoring activities, enhance the sharing of information with external users (including data providers), and improve the communication and reporting of aspects related to observations.

System architecture

As a real-time status dashboard, the application needs to be highly available, with as few potential points of failure as possible. For this reason, we opted for a very lean architecture, with a minimum number of 'moving parts'. The visible part of the dashboard is a front-end-only application, that runs in the browser and has no back-end component. This application gets the data from an object store, which is an Amazon Simple Storage Service (S3) bucket underpinned by a MinIO server hosted on the ECMWF data centre. Twice a day, when the suite runs, a script extracts the data from the Event database, formats it in a comma-separated values (CSV) file and sends the file to the object store. Having set up a specific retention policy, a built-in

feature of the MinIO object store, each file is removed after seven days.

See the flow chart for an overview of the system's architecture.

By simply using CSV files and leveraging an existing S3 object store with its built-in features, we have avoided the implementation of a separate backend and relative database. To some extent, we can say that the S3 object store fulfils at the same time the role of transient interchange data storage and that of a backend for the front-end application. Clearly, this architecture is minimalist and does not offer the potential to support more advanced features that would be easy to implement with a conventional backend/front-end architecture. However, it supports current requirements and reduces complexity and maintenance of the web stack (and therefore the likelihood of downtime) to an absolute minimum.

Expected evolution of the system

After its first operational implementation, the dashboard will undergo several enhancements that will enable an even better description of the observing system structure and related events. The planned upgrades include:

- adding support of observations used in atmospheric composition
- extending the support of satellite data from the global domain to other geographical regions
- extending the support of data to all geophysical quantities provided by each data type (currently an aggregated status is provided)
- fetching observation-related events from other sources (one candidate is the EUMETSAT user notification service)
- enhancing the traffic light system to allow a greater breakdown of data events
- adding more description and help texts within the dashboard.

Conclusion

The observations dashboard was initially intended to communicate detected observation anomalies. The project evolved towards a more inclusive and multi-purpose system that can serve the needs of a wide range of expert and non-expert users. The system as it is can be used as a first access point to explore ECMWF's observation monitoring resources.

Improving the support to users with a new forum service

Michela Giusti, Kevin Marsh, Anabelle Menochet, Xiaobo Yang

Our user community has been growing thanks to the wide range of data and services made available by ECMWF. One of these services is the new forum, which builds on the success of the previous (Copernicus-focused) forum. The aim of the new forum is to reach a larger number of users, from ECMWF itself as well as from the Copernicus services run by ECMWF, and to help them build a user self-supporting community where they can share their knowledge and experience with others. A short project began in October 2023, and the new forum platform was released on 14 February 2024.

From Copernicus User Support to the new forum

Over the last few years, various support channels have been introduced to better support Copernicus users with the aim to encourage self-help. Until April 2019, the Copernicus User Support Ecosystem was only composed of the Jira-based service desk and the Copernicus Knowledge Base. However, these components were lacking a feature to enable users to help themselves. In order to connect with the growing Copernicus

community, the Copernicus User Support forum was launched in April 2019. It was primarily intended for users of the Copernicus Climate Change Service (C3S) Climate Data Store (CDS) and the Copernicus Atmosphere Monitoring Service (CAMS) Atmosphere Data Store (ADS), and it was built using the Atlassian Confluence platform.

After a review of the user journey carried out by Nomensa in the last quarter of 2021, the Data Support team introduced the ‘User Support Journey’ and restructured the forum platform following their recommendations. The ‘User Support Journey’ includes four channels: the Virtual Assistant, the Knowledge Base, the User Forum, and the Jira-based Support Portal. The aim is to direct users to find the information they are looking for as quickly and efficiently as possible.

While Confluence provided a good starting point, it became apparent over time that a dedicated forum platform with advanced capabilities would be needed. A new project started in October 2023 with Interloop, with the aim of moving the forum from Confluence to a new platform. In the first part of the project, a review of

the most popular forum platforms was carried out, and Discourse was selected as the best candidate. The main benefits of choosing Discourse are:

- It is a leading forum platform, supporting single sign-on (SSO), easy look and feel customisation and embedding, and quick content migration.
- It provides a much-improved user experience.
- It supports threading, categorisation and tagging of discussions.
- It allows us to better respond to the needs of a fast-growing user community, greatly boosting collaboration and self-help among users.

Having identified the platform, the Data Support team worked with the contractors and a number of ECMWF teams to fulfil the different requirements to make the platform work in the ECMWF infrastructure. In particular, the SSO was implemented with the ECMWF authentication and authorisation service, and the look and feel was



Forum screenshot. The new forum platform is now accessible at: <https://forum.ecmwf.int>.

customised to assure consistency with ECMWF and Copernicus stylesheets.

In order to widen the scope from the original target audience, the structure of the forum was completely reviewed to make it more general and accessible for all ECMWF users. The new structure includes categories for announcements, access and login, datasets and usage, software and tools, training, workshops and events, feedback, and general background information. There are also subcategories related to the Copernicus services and general services. All the content and related metadata of the previous forum platform were migrated and redirected/reorganised to comply with the new structure. A new ECMWF subdomain and related email address

were also implemented to improve consistency and allow notifications to be sent to users when there is activity in the forum.

The forum platform is now accessible at: <https://forum.ecmwf.int>. Here, users can engage in discussions and share information (when logged in using their ECMWF account), and access the latest news about Copernicus/ECMWF products.

What next?

The next steps to make the forum a key point of the user journey include:

- selecting and training moderators who can act as area specialists to

stimulate discussion (and will be identifiable as such from the badge icon next to their name)

- making the forum a useful and valuable resource for all users (e.g. training, open data, Integrated Forecasting System users) in the service of ‘one ECMWF’ – it is no longer a service only available to CAMS and C3S users
- developing a strategy to promote the forum widely. In this context, we are already working closely with collaborators to embed the content of the forum in the new CDS and ADS portal so that users can easily find it and better access the forum platform.

Improving the boundary forcing in reanalyses and seasonal forecasts

Magdalena A. Balmaseda, Tim Stockdale, Retish Senan, Souhail Bossetta, Angela Benedetti, Tanya Warnaars

Radiative forcing, land cover/use and vegetation are prescribed as boundary conditions in reanalyses and subseasonal–seasonal–decadal integrations. Their temporal variability impacts the solution. There are new high-accuracy observational datasets for the recent period (the last 20–30 years), which are not yet used in reanalyses and re-forecasts. To facilitate the seamless approach adopted by ECMWF, there is a need to create time-varying records which are consistent with those used in operational medium-range prediction. There is also a need to find the best and converging practices between the operational centres providing reanalyses and seasonal forecasts and the wider climate community. These have been the motivations for the activities in the EU Horizon 2020 project CONFESS (<https://confess-h2020.eu/>). CONFESS aimed to improve the representation of global trends and regional extremes in the next generation of Earth system reanalyses and seasonal forecasts made available by the EU-funded Copernicus Climate Change Service (C3S) implemented by ECMWF. It has done so by taking stock of observational datasets and model developments across different Copernicus Services on

vegetation, land cover, atmospheric composition and biomass burning.

CONFESS comes to an end in the spring of 2024, and it is time to review its legacy. The achievements of CONFESS have been possible thanks to the engagement and expertise of the project partners: Météo-France, ECMWF, Italy’s Institute of Atmospheric Sciences and Climate (CNR-ISAC), and the Barcelona Supercomputing Center (BSC).

Land cover/use and leaf area index

CONFESS has managed to identify and harmonise different datasets of land use/land cover (LULC) and leaf area index (LAI). The implementation of these datasets in different models has allowed coordinated experimentation in a multi-model framework, resulting in the quantitative assessment, for the first time, of the impact of time-varying land properties in multi-year land simulations and seasonal forecasts. The multi-model assessment has indicated that time-varying land properties, especially vegetation, have a significant impact on trends, skill and extremes. However, the impact is highly dependent on the model, and there is a further need to quantify and

constrain the land–atmospheric feedbacks. Therefore, the time-varying land properties will not be included in ECMWF’s next reanalysis, ERA6, or in its next seasonal forecasting system, SEAS6. They are, however, a key component of further developments for ERA6-land and ERA7. The most recent period of the CONFESS LULC and LAI datasets, treated as a fixed climatology, has proved beneficial for numerical weather prediction and will go in the next operational cycle of ECMWF’s Integrated Forecasting System (IFS), Cycle 49r1, and will therefore be inherited by ERA6 and SEAS6. The LAI has also been used in decadal integrations by partners in the Italian National Research Council (CNR), and to evaluate the quality of the prognostic vegetation models used by Météo-France.

Tropospheric aerosols

CONFESS has also enabled a major step in the treatment of time-varying tropospheric aerosols, by merging the capabilities for the EU-funded Copernicus Atmosphere Monitoring Service (CAMS) implemented by ECMWF, our current reanalysis (ERA5), and the latest repository of tropospheric aerosol emissions

LULC/LAI in IFS outdated and not time-varying

Recent Copernicus temporal records of LAI/LULC not used by models

CONFESS

- New time-varying LAI-LULC datasets implemented and evaluated in models
- Recommendations for future research directions
- New fixed LAI-LULC climatology in operations
- Impact on seasonal forecast skill, trends and extremes evaluated and quantified

Urgent need to update aerosols in ERA6-SEAS6

CAMS capabilities unexploited for climate
Divergent practices in climate and reanalyses community
Burning questions

CONFESS

- New tropospheric aerosols using CAMS/CMIP ready for ERA6-SEAS6
- New up-to-date volcanic model ready for ERA6-SEAS6
- Impact of large biomass burning event quantified in the IFS


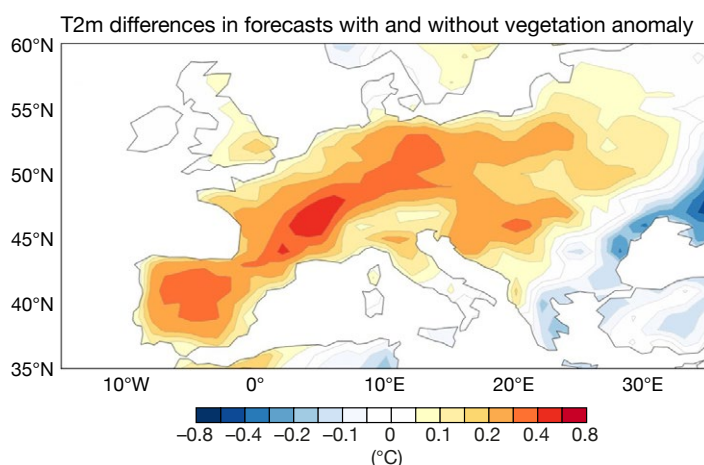

CONFESS in a nutshell. The situation regarding land and aerosols in the IFS before and after CONFESS.

(an update to CMIP6, likely to be used in CMIP7). The aerosol optical properties in the IFS have also been revised (see the article on a new time-varying tropospheric aerosol climatology in this Newsletter). This new dataset will be implemented in IFS Cycle 49r2, to be used in ERA6 and SEAS6. Having an up-to-date aerosol climatology that is consistent with the latest CAMS aerosols also helps us to explore the impact of interactive aerosols on numerical weather prediction.

Hazardous events

The activities in CONFESS have also made steps towards enabling prediction and monitoring systems to react to hazardous events, such as large volcanic eruptions and biomass burning. A new model for volcanic stratospheric aerosols (EVA_H) has been implemented in the IFS. This made it necessary to revise the aerosol optical properties. The EVA_H model has been thoroughly evaluated in seasonal and decadal forecasts, and it will be used in SEAS6 re-forecasts. It has

also been possible to quantify the impact of large biomass burning emissions in sub-seasonal and seasonal forecasts. The results, based on experimentation with prognostic aerosols in the IFS targeting three case studies, show that although there is a noticeable impact on predictions of local surface temperature, the amplitude is modest and it does not alter the large-scale atmospheric circulation. But this work was done before revising the IFS aerosol optical properties, and it should be repeated with more recent upgrade Cycles.



Impact of vegetation on a retrospective seasonal forecast of heat extremes. In the summer of 2003 (June–July–August), Europe experienced an intense and long-lasting heat wave, as indicated by the map of the ERA5 2-metre temperature (T2m) anomaly. The hot and dry conditions decreased the vegetation fraction (see the map showing the LAI anomaly in the same time frame), which in turn contributed to the intensity of the heat. This is shown in the map of 2-metre temperature differences in seasonal forecasts for June–July–August, initialised in May, with and without prescribed vegetation anomalies.

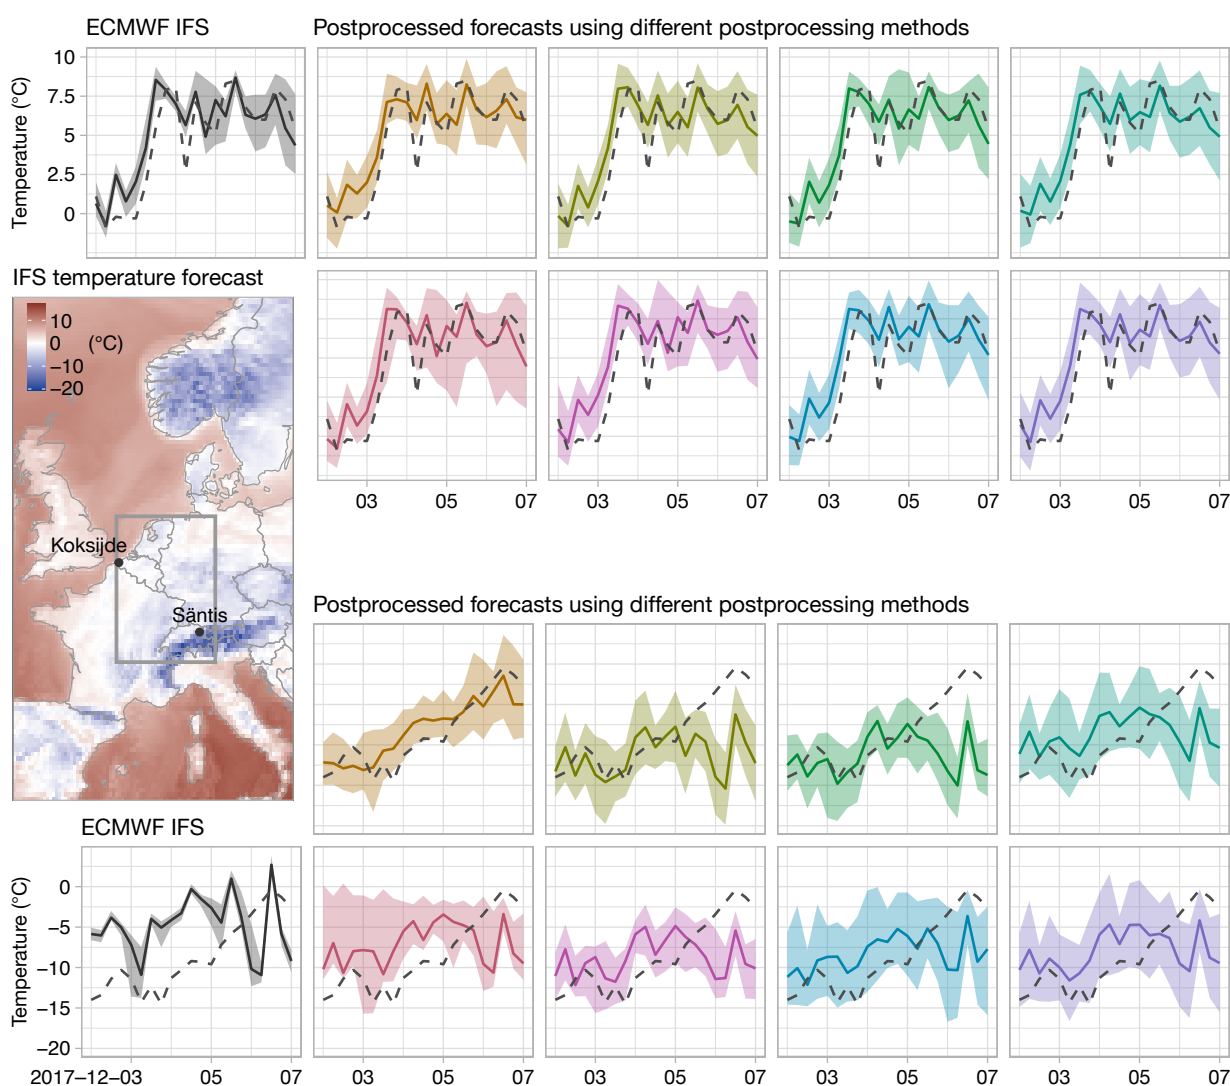
EUPPBench: A forecast dataset to benchmark statistical postprocessing methods

Cristina Primo (Deutscher Wetterdienst), Zied Ben Bouallègue (ECMWF), Jonas Bhend (MeteoSwiss), Sebastian Lerch (Karlsruhe Institute of Technology), Bert Van Schaeybroeck (Royal Meteorological Institute of Belgium), Jonathan Demaeyer (Royal Meteorological Institute of Belgium and EUMETNET)

During the last couple of years, several open-data and high-quality datasets have been published and made available online. An example is the WeatherBench initiative for data-driven medium-range weather forecasting. Such a benchmark dataset is useful to speed up progress in a field of research by enabling a fair comparison of different methods developed for the same goal. These datasets are

particularly convenient when large amounts of data are required to accomplish a task, and the endeavour of building a dataset can benefit a whole community. Statistical postprocessing – here the process of correcting systematic errors of medium-range weather forecasts – might also benefit from a benchmark dataset, since it is a data-intensive task by definition: based on historical data of forecasts and

observations, one aims to learn from past errors to improve future forecasts. Therefore, as part of the activities of the European National Meteorological Services Network (EUMETNET), the module on statistical postprocessing of weather forecasts, EUMETNET PP, has developed a benchmark dataset for statistical postprocessing named EUPPBench. The aim is to facilitate the comparison of rapidly evolving



First benchmark experiment. Example forecast evolution over seven days in December 2017, illustrating the extent of the EUPPBench dataset and of the first benchmark experiment (indicated by the box and the points in the left panel), along with example forecast time series (colours) and verifying observations (dashed lines) from the first benchmark experiment for Koksijde (top panels) and Sântis (bottom panels).

postprocessing techniques, especially new machine learning (ML) methods.

Contribution from ECMWF and NMHSs

EUPPBench is based on ECMWF Integrated Forecasting System (IFS) forecasts for the years 2017–2018 and corresponding reforecasts for 1997–2018, over Western Europe. Time-aligned observations from national meteorological and hydrological services (NMHSs) are available on a grid and at station points. The observations cover 22 years to match the reforecasts, which span 20 years in the past and 2 years for testing and validation. More than 40 meteorological variables are included in the benchmark dataset to allow experimenting with different targets and predictors as input for the postprocessing methods. Not only variables at the surface and at pressure levels are included in the dataset, but also more sophisticated ones, such as the Extreme Forecast Index (EFI). The data are stored in the Zarr format on ECMWF's and EUMETSAT's European Weather Cloud (EWC). A specifically designed *climelab* plugin was developed for easy access to the data.

Results

Can we improve 2 m temperature

ensemble forecasts at station locations with postprocessing? This simple challenge served as a basis to demonstrate how EUPPBench can be used for research. Different state-of-the-art methods for ensemble postprocessing were tested and their performance was compared using a set of verification metrics. This first comparative study was simple in the sense that the same variable was used both as sole predictor and as a target of the postprocessing methods. So, in our case, the goal was to predict 2 m temperature using only 2 m temperature forecasts (and metadata information) as predictors.

We learnt a number of lessons while generating and assessing this first benchmark challenge. Of course, we learnt about the methods themselves and their performance for postprocessing of ensemble forecasts at station locations. But we also gained insight into the whole process of building a benchmark dataset and how to approach this complicated exercise. It is worth mentioning that this benchmarking activity also provided a collaborative platform for discussions and exchange of ideas within the postprocessing community, spanning national meteorological services and research groups at universities.

The results of the inter-comparison exercise and the lessons learnt are published in Earth System Science Data (Demaeyer et al., 2023, <https://doi.org/10.5194/essd-15-2635-2023>). Not only the data (available on the EWC) but also the source code for generating postprocessed ensemble forecasts of 2 m temperature with a variety of methods are available through a GitHub repository, as indicated in the published article. Someone developing a new postprocessing approach can today compare their results with state-of-the-art methods using the EUPPBench framework.

Prospects

No doubt the lessons learnt during the first phase will prove valuable for the second phase of the EUMETNET PP module starting this year. The new ambition is to create a general benchmark for statistical postprocessing of precipitation ensemble forecasts, using high-resolution observations obtained from a collaboration with the European project RODEO. During this phase, we will thus seek to exploit all the information available in different predictors, and we will particularly encourage the testing of ML methods.

Copernicus Interactive Climate Atlas: a new tool to visualise climate variability and change

András Horányi

In February 2024, the Copernicus Climate Change Service (C3S), implemented by ECMWF on behalf of the European Commission, launched the Copernicus Interactive Climate Atlas (C3S Atlas in short). The C3S Atlas is a visualisation tool designed to provide complex climate information about the past (monitoring) and the future (projections) with only a few clicks. It is accessible at <https://atlas.climate.copernicus.eu>, and the first figure shows its landing page.

Background to the C3S Atlas


The most recent Assessment Report

(AR6) of the Intergovernmental Panel on Climate Change (IPCC) provided a state-of-the-art review of climate change that also included, for the first time, an Interactive Atlas (IPCC-IA, <https://interactive-atlas.ipcc.ch>). This IPCC-IA is, however, a static application, 'frozen' to reflect the content of the AR6 report, with no updates planned before the IPCC's seventh Assessment Report (AR7).

C3S embarked on a plan to operationalise the IPCC Atlas. This took into account C3S's vision and priorities, but it was also aligned with IPCC requirements and principles. Particular attention was paid to the

provision of regional climate information, the compatibility of data with the IPCC's quality standards, and the enhancement of the functionality and interactive features of the graphical displays.

Robert Vautard, the co-chair of IPCC Working Group I (<https://www.ipcc.ch/working-group/wg1/>) – the Working Group which hosted IPCC-IA in AR6, said: "I really welcome the initiative of C3S to extend the life of the IPCC-IA, because it will continue to provide this information for decision-makers, students and scientists, and also for IPCC authors of the 7th assessment cycle."



C3S Atlas landing page. The initial page of the Copernicus Interactive Climate Atlas, showing CMIP6 global climate projections for 2 m temperature for a Global Warming Level of 2°C. *Credit: C3S/ECMWF.*

What is new in the C3S Atlas compared to the IPCC-IA?

The C3S Atlas visualises datasets which are available in the C3S Climate Data Store (CDS). The first version deployed includes gridded observations (E-OBS); global reanalyses for the atmosphere, land, and ocean (ERA5, ERA5-Land and ORAS5, respectively); and climate projections (CMIP5, CMIP6 and CORDEX simulations). In particular, the CORDEX datasets include new information relative to the IPCC-IA: CORDEX-CORE, with a horizontal resolution of 25 km for the land areas of the globe, forms a global mosaic from individual simulations available from the various CORDEX domains; and CORDEX-EUR-11 includes an even higher resolution for Europe (12.5 km). The simulations for Europe were partly financed by C3S.

In terms of functionalities, the most exciting new feature of the C3S Atlas viewer is the customised selection of regional domains. That includes domains pre-defined in AR6, EUCRA (European Climate Risk Assessment) and European country regions, as well as any user-selected polygonal area of the globe.

What are the main visualisation features of the C3S Atlas?

The C3S Atlas displays global and

regional information for 30 different climate variables and indices; for various climate scenarios; for different time periods and Global Warming Levels (GWLs); and for absolute values (climatology) and change values for monthly, seasonal and annual timescales.

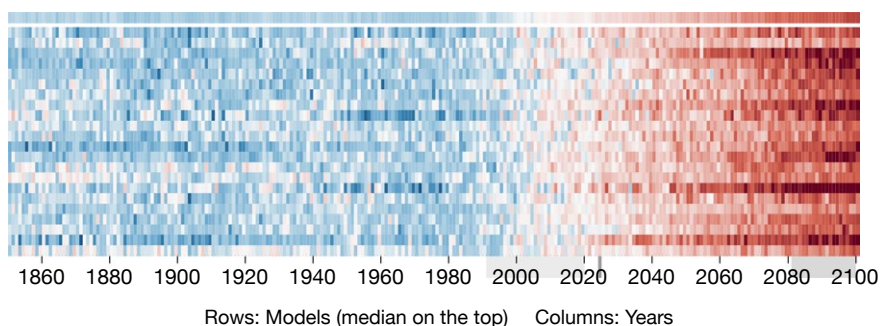
The main visualisation features include traditional global maps in different map projections; regional timeseries; various types of 'climate stripes' (including seasonal stripes); and annual cycle plots. A version of the climate stripes is shown in the second figure.

Most of the climate simulation plots display uncertainty information,

visualised by combining the relevant projections from all available models. In global maps, this uncertainty is shown in terms of robustness (agreement or disagreement between the models) and various visualisations of model spread in regional timeseries and climate stripe plots.

Where can I find more information about the C3S Atlas?

C3S published a web article for the launch of the C3S Atlas: <https://climate.copernicus.eu/copernicus-interactive-climate-atlas-game-changer-policy-makers>



Rows: Models (median on the top) Columns: Years

C3S Atlas climate stripes. An example of a version of the 'climate stripe' plots for annual mean 2 m temperature for Hungary (any European country can be selected). In the vertical direction, the median of the models is at the top, and below it includes all the available CMIP6 models (for the pessimistic SSP3-7.0 scenario). The horizontal axis shows the years from 1850 to 2100. The climate reference period is 1991–2020. The red colours indicate warmer and the blue colours colder climate with respect to the reference. *Credit: C3S/ECMWF.*

Ocean-wave-related changes in the next model upgrade

Jean-Raymond Bidlot, Peter Janssen

This article presents recent developments related to ocean wind waves within ECMWF's Earth system model. A revision of how waves are generated by wind is proposed. It improves the representation of air–sea momentum exchange under strong wind conditions, such as tropical cyclones and intense windstorms. Accounting for the role of the sea state (waves) in heat and moisture exchange has also been successfully introduced. These updates are informed by the strategic view that a better representation of coupled processes at the interfaces between components of the Earth system will lead to more realistic simulations of impactful weather conditions and, potentially, to a better use of observations sensitive to the ocean surface. These developments will be part of Cycle 49r1 of ECMWF's Integrated Forecasting System (IFS), which will replace IFS Cycle 48r1 later this year.

Revision of the wind input

ECMWF's Earth system model includes an ocean wind wave component (ecWAM) to represent how the generation of waves by wind modulates the exchange of momentum between the atmosphere and the ocean. In that framework, the characteristic length scale for momentum exchange is modelled by a sea-state-dependent Charnock relation, linking the surface aerodynamic roughness length to the surface stress. Because more momentum is always extracted from the wind when waves are growing, the drag coefficient (C_d) over the ocean tends to increase with wind speed, in line with the finding from observation campaigns (Edson et al., 2013).

However, recent experimental evidence suggested that C_d should generally reach maximum values for storm winds but should level or even decrease for particularly strong winds, such as in tropical cyclones or intense mid-latitude windstorms. The exact physical processes responsible are still a matter of active research. There is a decoupling between near-surface winds and the surface for strong wind situations. This could be due to flow separation, spray generation and wave dissipative effects, which are absent at lower wind conditions, and the impact of heavy rainfall, to name but a few factors, all of which reduce the transfer of momentum into the wave fields.

Based on this evidence, we aimed to adapt model parametrizations to reduce the drag coefficient for high winds. The latest revision was implemented in IFS Cycle 47r1, which became operational in 2020. The transfer of momentum from wind to waves was modified at the time by imposing an ad hoc sharp reduction of the parametrized contribution of the short gravity-capillary waves to the overall momentum flux. Note that those short gravity-capillary waves, with a wavelength of the order of centimetres, are not explicitly represented in the model, and their impact is therefore parametrized. Because the intention of that change was to tackle the intensity error in tropical cyclone conditions, it was only applied to 10 m winds above 33 m/s. The idea was that a strong decoupling between winds and short waves is happening in those conditions. The new model was found to yield a better maximum wind/minimum pressure relation, in particular at a high resolution (Bidlot et al., 2020; Majumdar et al., 2023).

Since the introduction of this simple reduction, we have worked on a more physically based model that would not require a fixed wind threshold. Based on the theoretical work of Janssen and Bidlot (2023), the current wind input parametrization can be extended to include a simplified model for the impact of short gravity-capillary waves on the surface stress, and for nonlinear feedback from the waves on further growth by wind. The new model tells us that, under strong wind conditions, the short gravity-capillary waves are less able to absorb momentum from the wind as hypothesised with the simple reduction scheme. But this is the case for storm conditions (wind > 20 m/s), well before the hardcoded threshold of 33 m/s of the simple reduction scheme. At the same time, as the sea state grows and larger and longer waves appear, a nonlinear feedback kicks in, reducing the growth rate of the waves. Essentially, further growth, and hence further transfer of momentum into the ocean, is stalled. For low wind conditions (< 10 m/s), the new parametrization behaves generally like the previous one. Only for a sea state approaching an equilibrium between wind input and dissipation (fully developed sea) as is quite common in the trade winds, the interaction with the wind becomes weaker, resulting in a slightly lower contribution from the local wind to the overall significant wave height. Note that, since waves propagate away from the generation areas into the rest of the ocean

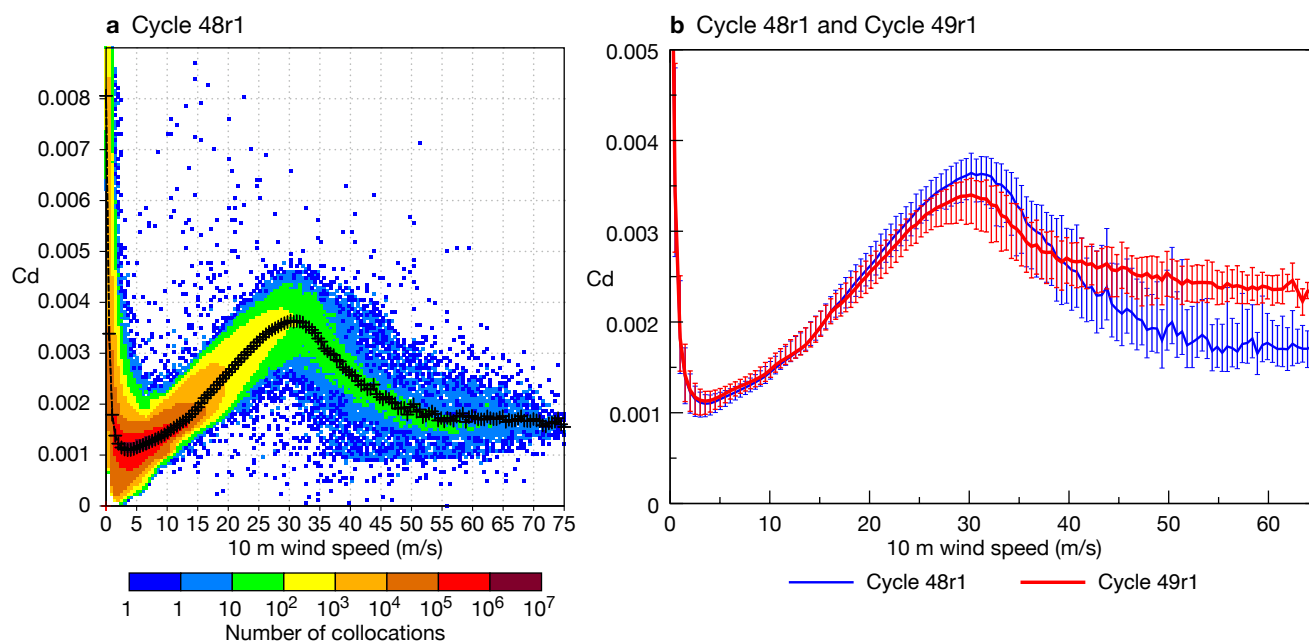


FIGURE 1 Drag coefficient (Cd) over the North Atlantic and corresponding 10 m wind speed during Hurricane Lee. Forecasts from the operational analysis of 8 September 2023, 00 UTC were performed at the experimental TCo2559 resolution (4.4 km). Results are shown aggregating all hourly forecast steps from 12 to 72 hours for (a) Cycle 48r1, with black crosses showing mean Cd values for given wind speeds, and (b) Cycle 48r1 (blue) and Cycle 49r1 (red) for the binned mean values and error bars for one standard deviation on either side of the mean.

basin as well, it was necessary to adjust slightly how swell is attenuated as it propagates over a large ocean expanse. This is to avoid flooding the tropics with excessive swell energy.

Figure 1a shows the distribution of Cd in IFS Cycle 48r1 plotted against 10 m wind speed for Category 5 Hurricane Lee in the North Atlantic in September 2023. This simulation was performed at the experimental TCo2559 resolution (4.4 km) because at that resolution the IFS is better able to intensify tropical cyclones and to yield very strong winds (Majumdar et al., 2023). A comparison with the new parametrization of IFS Cycle 49r1 is shown in Figure 1b. Here, for clarity the full distribution has been replaced by the binned mean values and error bars for one standard deviation on either side of the mean. The impact of the new model is visible for winds of about 20 m/s with a reduction of Cd with respect to the old system. Both versions produce the expected sharp decrease for wind above 30 m/s. The new parametrization, however, has a lesser decrease for winds above 40 m/s, in line with recent observations of Curic and Haus (2020). Note that tropical cyclones are controlled by momentum loss (drag) but also heat and moisture fluxes. The impact of waves on the latter is revisited in the next section.

Sea-state-dependent heat and moisture fluxes

Recent experimental evidence points to a sea state/wind dependency on the sensible and latent heat flux.

To capture this, the theory behind the wind input as described above can be extended to include thermal stratification (Janssen & Bidlot, 2018). It can then be shown that wave-induced motions in the air lead to an enhancement of the heat and moisture exchange between the ocean and the atmosphere. Otherwise, this exchange would be controlled only by molecular transfer at the sea surface and turbulence. The theory yields a sea-state-dependent parameterisation for the roughness length scales for heat and humidity. Note that the effects of spray as generated by whitecaps and breaking waves are still ignored.

Figure 2 shows a similar plot to Figure 1, but for the heat exchange coefficient (Ch) instead. Note that from scaling arguments, and based on the actual parametrization at ECMWF, Ch has always been proportional to the square root of Cd. This explains the similar mean dependency of both exchange coefficients on wind speed. However, in Cycle 49r1, with the added sea-state dependency, the heat exchange will tend to be increased with respect to previous cycles, for storm wind conditions. As mentioned above, the new model will have slightly higher drag for intense tropical cyclones, in line with recent observations, potentially slowing cyclone intensification. But at the same time, the heat and moisture exchanges will be enhanced, favouring storm intensification. It was found that the impact of combining both changes was fairly neutral on the scores for tropical cyclone track and intensity (Magnusson et al., 2021).

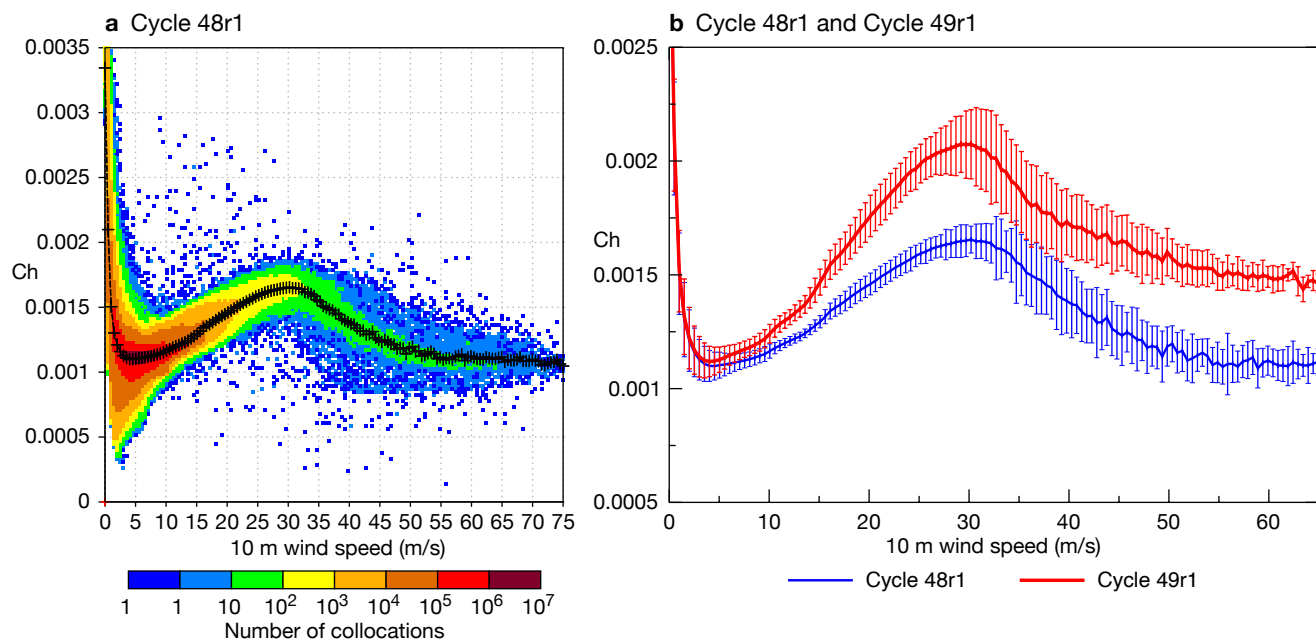


FIGURE 2 The same as Figure 1, but for the heat exchange coefficient Ch.

Other changes

In nature, waves and sea ice interact. These interactions are most prominent in the Marginal Ice Zone (MIZ). Waves can contribute to sea ice break-up and to mixing of warmer ocean water with the surface, leading to the continued existence of the MIZ. At the same time, sea ice attenuates and dissipates wave energy. This is an active current topic of research. However, these interactions are inexistent in the current operational system. Instead, sea ice is treated as if it were land. More precisely, when sea ice cover is above 30%, all wave spectral components are fully absorbed (i.e. set to 0), while for a lower sea ice threshold, sea ice has no impact on waves. For global wave forecasting, this was a fair assumption to make. However, the wave model is now part of ECMWF's Earth system model. It needs to supply relevant information to the other components (atmosphere and ocean). As a first step, over sea ice a new set of empirically derived default values for the relevant parameters has been introduced for Cycle 49r1 that were found to perform better than the previous set.

In addition, from Cycle 49r1 onwards, assimilation of altimeter wave height data will be done in hourly sequential windows rather than 6-hourly ones. This change has already been implemented in ERA5. It ensures a better use of the data as they are assimilated nearer to the time when the observations were made. It has a positive impact on short-range wave forecasts, but very little impact on other Earth system model components. Note that, since the wave data assimilation is still based on a simple sequential optimum interpolation scheme, the verifying analyses used in the forecast scores calculation will have seen less observations than before.

Resolution upgrade

The wave model is currently using a reduced latitude–longitude grid with a horizontal grid spacing about 1.5 coarser than that of the atmospheric model (0.125° or 14 km for waves and 9 km for the atmosphere). With Cycle 49r1, the wave model will be run on the same grid as the atmospheric model, simplifying the data exchange between the two systems. With higher resolution, the wave model is better able to capture the atmospheric forcing, and details of the coastal conditions are generally better represented. Note, however, that wave conditions in coastal areas will still be determined to a large extent by the quality of the forcing, the resolution of which has not changed. At 9 km, many small islands and bathymetric features are not resolved. It is still necessary to use the empirical parametrization that accounts for the impact of those small-scale features on the wave propagation. A small adjustment to the scheme was made for the new grid.

A few technical changes were necessary to reduce the cost of a higher horizontal resolution for the waves. The numerical scheme for the wave propagation has a strict constraint on the maximum time step that can be used. Finer resolution would imply a smaller time step. However, the scheme was modified to be split between long waves, which propagate the fastest and use the same time step as before, and shorter waves, which propagate more slowly and can use a larger time step. We use 36 frequency bins to discretise the wave spectrum, representing waves between 1,300 m and 1.5 m. The threshold between fast and slow waves is determined by the atmospheric model time step. For typical values for the atmospheric time step at

TCo1279 (ECMWF's triangular–cubic–octahedral grid with a grid spacing of 9 km), long waves are those with a wavelength longer than about 600 m, represented in the first five frequency bins. Only a small fraction of the wave spectrum then requires a smaller time step, the rest of the spectrum can be advected with a longer time, greatly reducing the overall run time.

In Cycle 48r1, the atmosphere and the waves exchange information every atmospheric time step (450 s at TCo1279). However, as the atmospheric resolution is increased and its respective time step is reduced, we found that it is no longer necessary to couple the two system so tightly. At TCo1279, coupling every two atmospheric time steps and integrating the wave physics over that longer interval was found to be sufficient, significantly reducing the running cost in Cycle 49r1.

At TCo1279, the wave spectrum is discretised with 36 frequencies and 36 directions. However, to reduce the data volume, please note that in Cycle 49r1 only the first 29 frequencies will be output spanning a range from ~0.035 Hz to ~0.5 Hz. Higher frequencies tend to be in balance with the wind forcing. Therefore, even when used to restart the model, the high-frequency part of the spectrum is re-established within one time step without the need to have that information loaded from the initial conditions. Moreover, most wave spectral observations have a cut-off frequency around 0.5 Hz or lower.

Impact

The performance of the new model was first tested in standalone mode, where ecWAM was run at operational resolution with prescribed 6-hourly atmospheric forcing from the operational analysis from August 2022 to August 2023. Altimeter wave heights from Jason-3, Cryosat-2, Saral, and Sentinels 3A, 3B and 6A were assimilated. Significant wave heights predicted by the model can be compared to altimeter data and independent in-situ observations not used by the wave data assimilation. The in-situ observations are those

used by the Lead Centre for Wave Forecast Validation (Haiden et al., 2019). Table 1 shows the performance gain in terms of scatter index with respect to Cycle 48r1, first using the same grid for waves as Cycle 48r1 (0.125° or 14 km, reduced latitude–longitude), and then with the new TCo1279 grid option. Note that these simulations only assess the impact of the change in the wind input and the change in resolution. Overall, both changes have a cumulative beneficial impact. The impact of the resolution change is most pronounced when a comparison is made to in-situ observations. These observation sites are mostly located in the northern hemisphere in near-coastal environments, where increased resolution can result in further improvement.

The changes have also been tested in coupled mode. A series of analysis experiments at TCo399 were performed covering one northern winter season (2021–2022) and one northern summer season (2022). Looking first at the broad atmospheric response, Figure 3a shows that the change to the wind input parametrization in ecWAM and the change in wave resolution alone had a neutral impact on atmospheric temperature scores. The same can be said about the impact on other variables. On the other hand, the sea state effect on heat and moisture exchanges is widely beneficial (Figure 3b,c). It results from a clear warming in the lower troposphere over the oceans in the tropics and the winter hemisphere (not shown). Near-surface temperatures are improved (Figure 4). Root-mean-square error scores for tropical temperature at around 850 hPa are slightly degraded as the model already exhibits a positive bias at those levels. This bias aspect will be partially addressed by other contributions to Cycle 49r1 (not shown) and will continue to be analysed in future model developments.

From the point of view of modelling the strong interactions between wind and waves, this set of changes will tend to reduce the dampening effect waves have on near-surface winds through weaker feedback on the momentum flux, and it will enhance the heat and

	SI (%) Cycle 48r1	Change in SI (%) compared to Cycle 49r1 at 0.125°	Change in SI (%) compared to Cycle 49r1 at TCo1279
In-situ	14.63	-0.93	-4.10
Altimeter NH	10.81	-1.39	-1.66
Altimeter Tropics	8.00	-0.38	-0.88
Altimeter SH	8.64	-2.78	-3.59

TABLE 1 Standard deviation of the difference between observations and the model normalised by the mean of the observations (also known as the scatter index – SI) for significant wave height for August 2022 to 2023. In-situ observations are compared to the analysis, and altimeter data are compared to the model's first guess. For altimeter data, separate findings are reported for the northern hemisphere (NH), the tropics (Tropics) and the southern hemisphere (SH). The two columns on the right show the relative change in SI in percentage terms with respect to Cycle 48r1. Negative values indicate an improvement.




FIGURE 3 Normalised change in temperature forecast root-mean-square (RMS) error, measured against own analysis, showing the impact of (a) the revised wind input, (b) sea-state-dependent heat and moisture fluxes, and (c) all wave-model-related changes, for combined winter and summer seasons. Cross-hatching indicates statistical significance at a confidence level of 95%. Blue areas indicate a reduction in RMS and hence a beneficial impact from the contributions.

moisture uptake over the oceans. As a result, more intense depressions with stronger winds during mid-latitude storms are produced. These stronger winds ultimately generate higher waves in the storm tracks. Figure 5 shows that, with the new formulation, significant wave heights tend to be higher in the storm track of the southern hemisphere for all seasons and in the northern hemisphere winter. Recall also that, for fully developed sea in trade wind conditions, the new model produces a lower contribution to the overall significant wave height. Hence, significant wave height in tropical areas is lower except if directly affected by strong swell emanating from the Southern Ocean, such as in the Indian Ocean during austral winter (Figure 5b).

Forecast scores for wave parameters against own analysis were found to be slightly degraded with the new wind input changes alone (Figure 6a) due to enhanced variability. However, against in-situ observations there is a small improvement like the one described in Table 1. Adding the sea state dependency on heat and moisture fluxes had a small beneficial impact on forecasts of waves via its positive effect on the atmospheric scores (not shown). The combination of all changes gives an apparent big gain in performance for short lead times when verified against own analysis (Figure 6b). This is partly an artefact of the verifying

analysis having seen less observations at the time it is compared to forecasts than in the reference experiment. Nevertheless, compared to observations, the impact is less but still positive.

So far, we have focussed on significant wave height as it is the only wave parameter that is reliably observed globally by altimeters which currently enter the wave data assimilation. Scores for mean wave periods are also indicative of the impact of the changes we make; however, the analysis is not constrained by any observations. Change in the mean state can potentially appear as having a large impact. The new system does indeed yield larger mean wave periods in the tropics linked to long swell contributing a bit more to the overall sea state than before. Consequently, scores appeared degraded. The new model is expected to produce more swell, and some retuning of the swell attenuation formulation was done. The present choice might still be sub-optimal. Future work should diagnose the model performance in tropical areas where in-situ observations are in short supply with wave spectra derived from Synthetic Aperture Radar (SAR) and recently available drifting buoys to steer new developments.

Finally, changing how wave parameters are estimated when sea ice cover is above 30% has a small positive




FIGURE 4 Normalised change in 2 m temperature forecast root-mean-square (RMS) error, measured against own analysis, showing the impact of sea-state-dependent heat and moisture fluxes for the combined winter and summer seasons. Blue areas indicate a reduction in RMS and hence a beneficial impact.

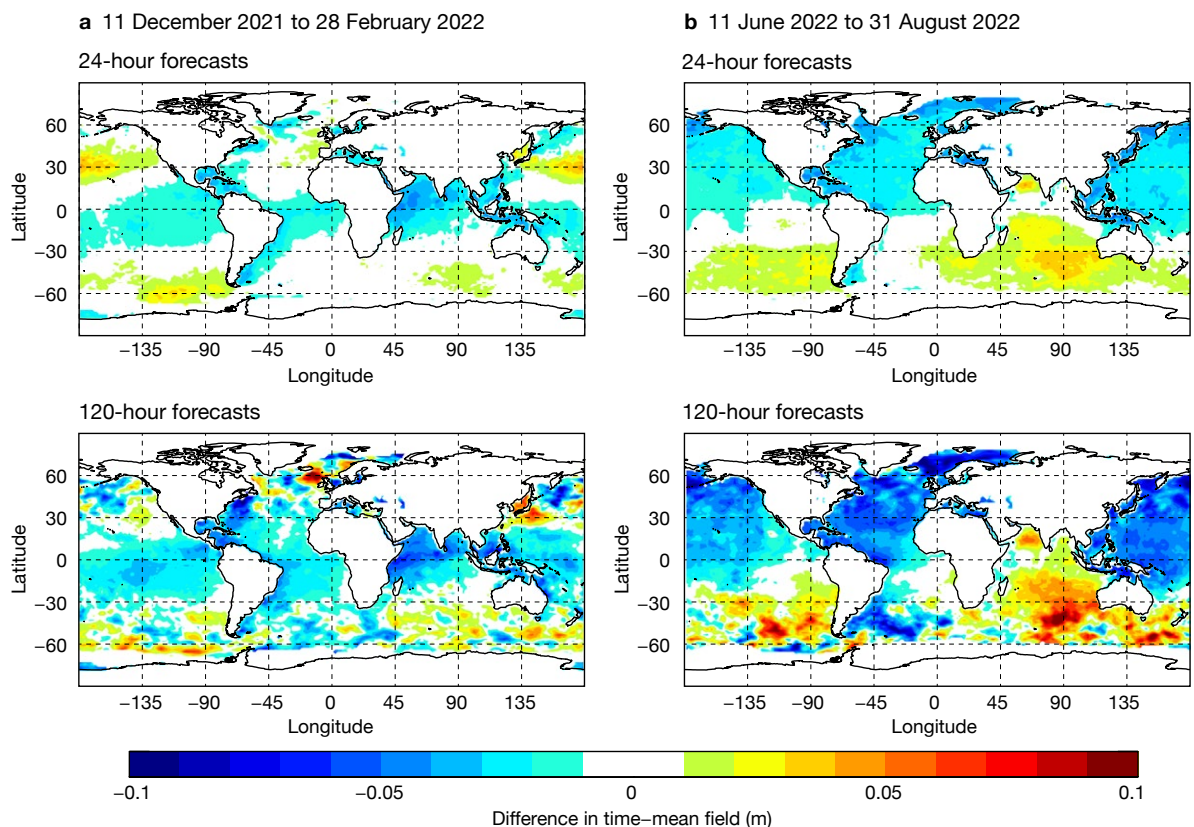


FIGURE 5 Mean difference in significant wave height for (a) 11 December 2021 to 28 February 2022 and (b) 11 June 2022 to 31 August 2022, showing the impact of all wave-model-related changes. Warm colours indicate higher mean significant wave height.

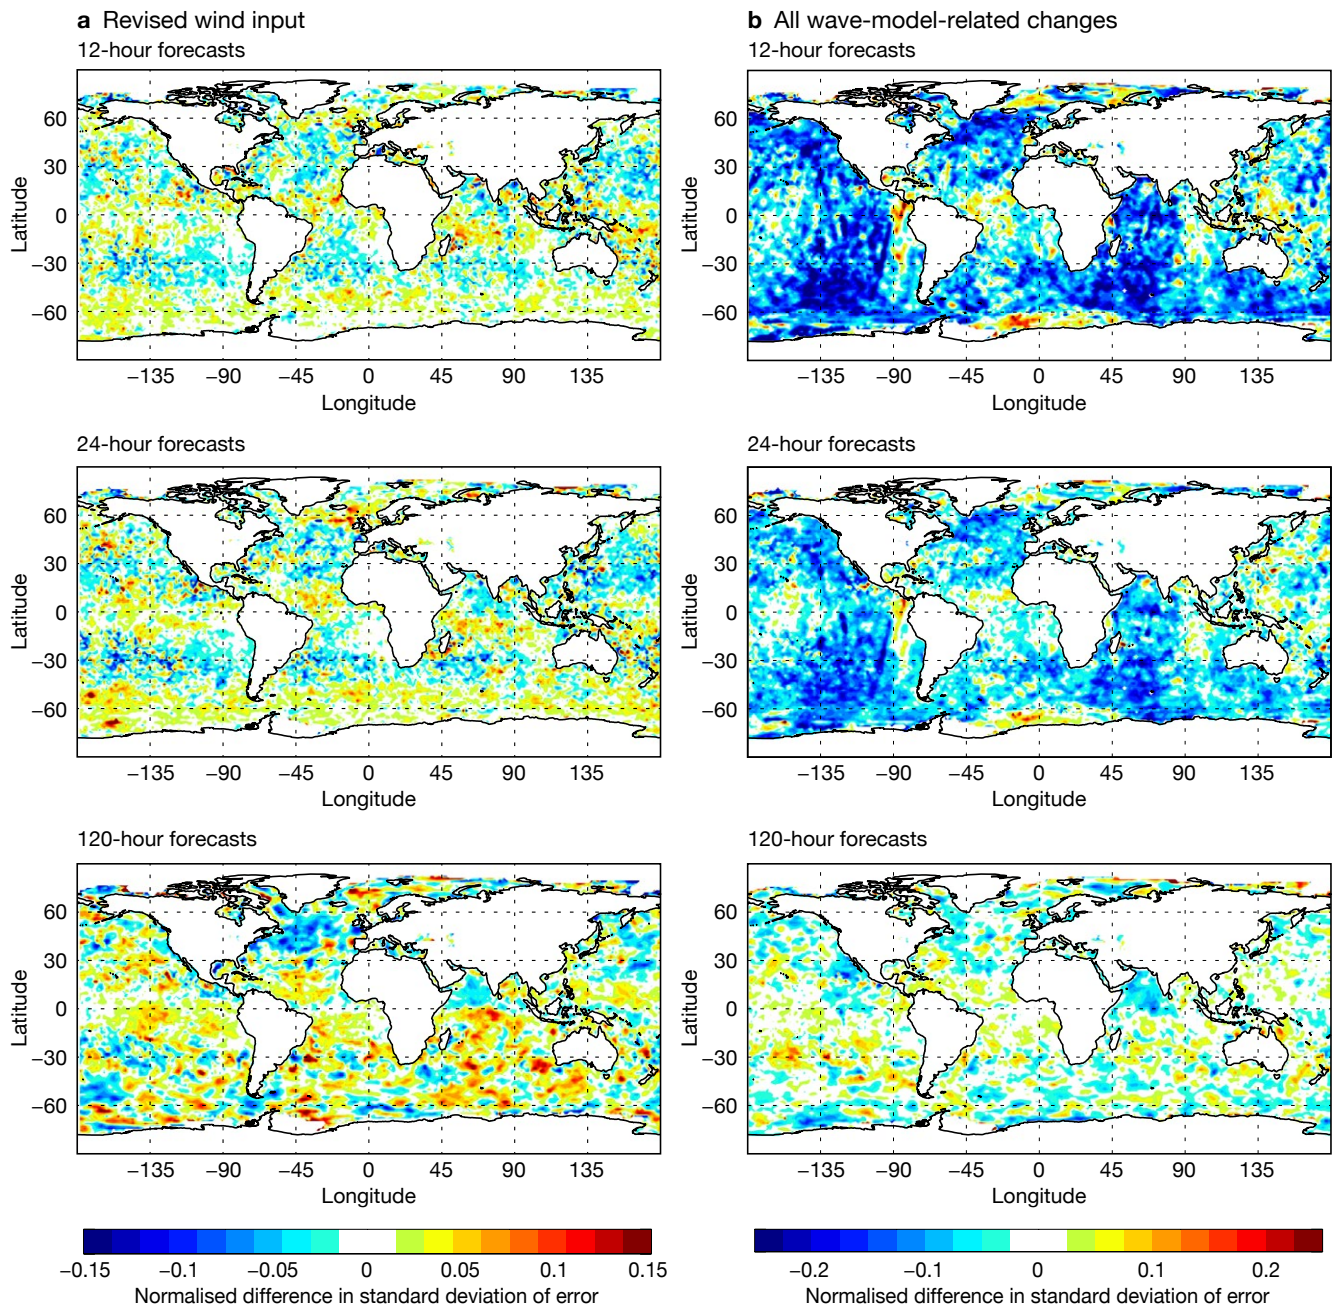


FIGURE 6 Normalised change in significant wave height forecast standard deviation of error (STDDEV), measured against own analysis, showing the impact of (a) the revised wind input and (b) all wave-model-related changes, for the combined winter and summer seasons. Blue areas indicate a reduction in STDDEV and hence a beneficial impact from the contributions.

impact on the sea ice cover forecasts (Figure 7). Ideally, the wave model should be run when sea ice cover is above this threshold. Recent progress has been made by the community on how to best represent wave–sea-ice interactions. Work is ongoing on implementing some of those findings in future model cycles.

Conclusion

The evolution of ocean wind waves determines ocean surface stress. A revision of how this momentum exchange is modelled in ECMWF’s Earth system model

is part of the upcoming Cycle 49r1. It is intended to be more consistent with recent estimates of drag coefficients under storm wind conditions (> 20 m/s). Moreover, the direct impact of waves on surface heat and moisture transfers was introduced in the atmospheric model. Cycle 49r1 will also see an increase in horizontal resolution for the wave products as the wave model will be running on the same grid as the atmospheric model component (i.e. TCo1279, which means 9 km for medium-range products, down from 14 km before). This will increase the consistency between waves and atmospheric parameters and provide added resolution in coastal environments.




FIGURE 7 Normalised difference in sea ice cover forecast root-mean-square (RMS) error, measured against own analysis, showing the impact of all wave-model-related changes for (a) 1 December 2021 to 28 February 2022 and (b) 1 June 2022 to 31 August 2022, for the southern hemisphere (SH) and the northern hemisphere (NH). Negative values indicate a reduction in RMS error and hence a beneficial impact.

Peak surface winds and significant wave heights during intense windstorms should generally be better represented with Cycle 49r1. Low-level temperatures over the ocean will also be slightly higher in the tropics and the winter hemisphere.

The current model does not explicitly account for the role of sea spray and active wave breaking in the calculation of ocean surface fluxes. These conditions tend to occur under extreme wind situations. Future

work should investigate whether introducing such processes can be beneficial, as ECMWF aims to run forecasts at kilometre scale, where the modelling of such extremes will be possible.

The small change in how wave parameters are passed to other components of the Earth system in sea-ice conditions hints at the possible importance of wave–sea-ice interactions. Work has started on adding these interactions, with potential for future progress.

Further reading

Bidlot, J.-R., F. Prates, R. Ribas, A. Mueller-Quintino, M. Crepulja & F. Vitart, 2020: Enhancing tropical cyclone wind forecasts, *ECMWF Newsletter* **No. 164**, 33–37.

Curcic, M. & B.K. Haus, 2020: Revised estimates of ocean surface drag in strong winds. *Geophys. Res. Lett.*, **47**. <https://doi.org/10.1029/2020GL087647>

Edson, J.B., V. Jampana, R.A. Weller, S.P. Bigorre, A.J. Pluedemann, C.W. Fairall et al., 2013. On the exchange of Momentum over the Open Ocean, *J. Phys. Oceanogr.*, **43**, 1378–1391.

Haiden, T., Z. Ben Bouallègue, R. Mládek & J.-R. Bidlot, 2019: WMO Lead Centre for Wave Forecast Verification established at ECMWF, *ECMWF Newsletter* **No. 161**, 13–14.

Janssen, P.A.E.M & J.-R. Bidlot, 2018: Progress in Operational Wave Forecasting, *Procedia IUTAM*, **26**, 14–29. IUTAM Symposium on Wind Waves, 4–8 September 2017, London, UK.

Janssen, P.A.E.M & J.-R. Bidlot, 2023. Wind–Wave Interaction for Strong Winds. *J. Phys. Oceanogr.*, **53**.

Magnusson, L., S. Majumdar, R. Emerton, D. Richardson, M. Alonso-Balmaseda, C. Baugh et al., 2021: Tropical cyclone activities at ECMWF. *ECMWF Technical Memorandum* **No. 888**.

Majumdar, S.J., L. Magnusson, P. Bechtold, J.-R. Bidlot & J.D. Doyle, 2023: Advanced Tropical Cyclone Prediction Using the Experimental Global ECMWF and Operational Regional COAMPS-TC Systems. *Mon. Wea. Rev.*, **151**, 2029–2048.

A new time-varying tropospheric aerosol climatology for the IFS

Tim Stockdale, Retish Senan, Robin Hogan, Zak Kipling, Johannes Flemming

Tropospheric aerosol has an important radiative impact on the atmosphere and wider Earth system. In configurations of ECMWF's Integrated Forecasting System (IFS) for medium-range and seasonal forecasting, it is represented as a fixed climatology, which was last updated some years ago. The present fixed climatology has two main weaknesses: it is unable to represent the large changes in anthropogenic aerosol that have occurred over recent decades, which have driven changes in the radiative balance relevant for both reanalysis and the calibration of seasonal forecasts; and it is incompatible with the representation of aerosol in the latest configurations of the EU-funded Copernicus Atmosphere Monitoring Service (CAMS), implemented by ECMWF, hampering work to assess the possibility of including interactive aerosols in our configurations for numerical weather prediction (NWP). We have therefore developed a new, decadal-varying climatology of tropospheric aerosol, derived from and compatible with the aerosol modelling used in CAMS, to support and improve our work in core NWP activities and the EU-funded Copernicus Climate Change Service (C3S) implemented by ECMWF. It is driven by the latest versions of emission datasets used in the World Climate Research Programme's CMIP6 project, and proposed for CMIP7. This new time-varying climatology is planned for IFS Cycle 49r2, for use in both the forthcoming ERA6 reanalysis and the SEAS6 seasonal forecast upgrade, and subsequently in other NWP configurations.

The importance of tropospheric aerosols

Aerosols in the atmosphere are important in many ways. They can absorb solar radiation, heating the air but reducing the sunlight at the surface; they can reflect sunlight directly to space, cooling the planet; they modify absorption and emission of longwave radiation; and they can modify the formation and properties of clouds, although cloud-aerosol interactions are not well represented by the IFS. There are many types and sizes of aerosol particles, whose optical properties can vary greatly with composition, size and the wavelength of light they are interacting with. Aerosol concentrations

vary over many orders of magnitude, and their radiative impact depends on both their intrinsic properties and their surrounding environment. The regional variation in importance of aerosols in the IFS can be demonstrated with an experiment in which all tropospheric aerosols are removed from the model: Figure 1 shows the impact of doing so on the net surface solar radiation, which increases by about 30% over parts of Africa, and on the 2-metre temperature, which increases by up to about 0.5 K in southern Europe.

A good model representation is hampered not only by the complexity of aerosols, but by the difficulty in accurately measuring their distribution and properties on a global scale. For many years, aerosols in the IFS were specified as a climatology based on external datasets, such as the Tegen et al. (1997) climatology used operationally until 2017. The most recent upgrade, used since IFS Cycle 43r3 and described by Bozzo et al. (2020), gave us for the first time a climatology based on aerosol modelling run at ECMWF. This climatology was created by first running the IFS and its CAMS aerosol model with realistic winds and specified emissions for the period 2003–2013 to create a model estimate of the climatology. The resulting 11 aerosol fields were then all scaled with a common factor at each horizontal grid point, such that the total aerosol optical depth summed over all aerosol species matched the values given by the CAMS interim reanalysis. This approach was taken because the reanalysis had enough satellite data to help constrain this vertical integral, but it was less successful in partitioning the optical depth between species, something which remains challenging today. Further adjustments, in particular an arbitrary reduction of aerosol over Africa in boreal summer, were found necessary to produce a climatology which could be used for NWP without damaging the scores.

Developing a new time-varying aerosol

Since this first CAMS climatology was produced, there have been major improvements in the chemistry and aerosol modelling used by CAMS. This motivates us to try to produce an improved version of the aerosol climatology. The approach we follow is similar, but with two key differences. Firstly, rather than use estimated aerosol values for a limited period of time to estimate a fixed climatology, we generate a dataset covering multiple decades, allowing us to calculate a decadal-




FIGURE 1 Climatological June–July–August (JJA) impact of removing all tropospheric aerosol from the IFS. We show (a) the percentage change in net surface solar radiation, and (b) the change in 2-metre temperature, calculated from 10 years of 10-member seasonal forecasts (2011 to 2020) starting on 1 May.

averaged climatology that is time-varying. Secondly, we rely entirely on the specification of emissions and the CAMS chemistry model to generate the aerosols, because satellite-constrained reanalyses of aerosol optical depth are not available over longer historical periods. The aerosol fields are constrained by meteorological and emissions data, but are not directly analysed, so we refer to the result as a pseudo-reanalysis. We had expected to need some final calibrations of total aerosol amount from a recent period for which CAMS analyses are available, but in fact the quality of the latest model means that this was not necessary.

We start by generating the multi-decadal pseudo-reanalysis. Since some aerosol species are strongly influenced by chemistry, we run the full Cycle 47r3 IFS-COMPO scheme with 123 chemical species and 16 aerosols. For affordability, we use a low-resolution TL255 system with a grid spacing of 80 km, which is appropriate as input for our final $3^\circ \times 3^\circ$ product, unchanged from the resolution of the previous climatology. Climatologies are inherently rather smooth: the radiative impacts of aerosol are most important on larger scales, and a higher model resolution would not reduce the many uncertainties in emissions and modelling.

Emissions of some aerosol types (sea-salt, mineral dust) are calculated by the IFS as it runs, but for most aerosols and chemical species, emissions must be specified from external data. Most of the emissions data comes from the latest version of the Community Emissions Data System (CEDS version v_2021_04_21), an extensive database developed to support climate change modelling such as CMIP6 and CMIP7. It contains very detailed information, specified as daily data on high-resolution grids, and with sector-specific information such as aviation, shipping, land transport, power plants, industry, agriculture etc. Natural

emissions sources, including biogenic sources and effusive (low-level) volcanic emissions, are also provided. Another important source of aerosol is fire emissions, either from human-directed burning or from forest fires. Here we use estimates from the historical global biomass burning emissions gridded dataset for CMIP6 (BB4CMIP6) for the period up to 2014 and from the Global Fire Assimilation System (GFAS) thereafter. A key chemical species for tropospheric chemistry is the hydroxyl radical (OH), and changes in this can drive changes in, for example, nitrate aerosols. Concentrations of OH in turn depend on methane (CH₄), which has increased substantially over the last 70 years and had to be properly accounted for.

The pseudo-reanalysis is created by initialising IFS-COMPO, running a 24-hour forecast, using the resulting chemistry and aerosol values to create initial conditions for the next 24-hour forecast, and repeating this continually. The meteorological initial conditions come from ERA5 and are reset every 24 hours, so the winds and aerosol transports are constrained to be accurate, but the chemistry and aerosols evolve freely according to the emissions and chemical and physical model processes. It would take too long to run this sequentially for the whole period, so the pseudo-reanalysis is run as a series of overlapping chunks, with a 6-month spin-up period discarded in each case. The pseudo-reanalysis runs from 1951 to 2019. For some aerosol types (sea-salt, mineral dust), we form a fixed climatology using the last 27 years of data, because we do not expect to resolve any significant long-term trends. Most aerosol types, however, have emissions dominated by human activity, and we calculate a time-varying climatology from a 9-year running mean. This is then sampled at 3-year intervals, giving an effective smoothing of the data closer to 10 years.

Figure 2 shows a few examples of the aerosol climatology for July in two different epochs, 1975

a

Aerosol species and global mean mass loadings in the new and old climatologies

Aerosol type	Old climatology	New climatology
Sea salt 1	1.1	4.4
Sea salt 2	63.0	47.0
Sea salt 3	76.4	15.4
Mineral dust 1	8.9	0.7
Mineral dust 2	27.7	6.5
Mineral dust 3	30.7	76.8
Organic matter (hydrophilic)	3.9	1.7
Organic matter (hydrophobic)	0.7	0.024
Black carbon (hydrophilic)	0.36	0.21
Black carbon (hydrophobic)	0.10	0.005
Sulphates	3.8	3.5
Nitrates (fine)		0.6
Nitrates (coarse)		3.0
Ammonium		1.1
Secondary organic (biogenic)		0.89
Secondary organic (anthropogenic)		0.61

This table compares the global-mean vertically integrated aerosol mass (in mg per m²) by species for the old and new climatologies for 2015. Salt and dust have three different size bins, and the last five species (formed from gases) are present only in the new climatology, although secondary organic aerosols were previously accounted for together with organic matter. The optical properties of the aerosols have also changed, so changes in radiative impact cannot be deduced solely from changes in the mass of aerosol.

and 2015. Industrial emissions of sulphate aerosol were high over Europe and the US in the 1970s, but have since reduced substantially, while emissions have increased in India, China and the Middle East over the same period. Some black carbon is from industrial emissions, whose growth in some regions can be seen, but it is dominated by biomass burning, some deliberate (as over Africa), but also increasingly by forest fires (Canada and Siberia). The increase in high latitude burning is related to a warming climate and is likely to continue to increase. Fine nitrate, an aerosol which originates from combustion processes and agriculture but further depends on chemical balances in the atmosphere, has increased in many areas. The figures here are only a snapshot: the full story of aerosol variation over time is complex.

The aerosol climatology represents aerosols which are formed or emitted in the troposphere, including those transported into the stratosphere. The stratospheric part of the solution requires longer spin-up times, achieved in a climatological sense with a multi-year 'pre-spin-up' run, although low frequency variation of tropospheric-origin aerosols in the stratosphere will not be properly represented. A bigger source of stratospheric aerosol is that formed after the injection of sulphur by large volcanic eruptions – this is represented separately in the IFS, either as a constant background (for NWP) or varying according to the history of volcanic eruptions (for reanalysis and seasonal forecasting).

Assessing the impact of the new aerosol climatology

The old aerosol climatology was adjusted to work well with the IFS, and subsequent IFS developments have been implemented so that they work well with these aerosols. The risk of degrading performance with a new aerosol climatology is thus real, even if it has a well-founded basis. Indeed, if we simply substitute the new climatology for the old with no further changes, the large changes in aerosol masses (see Box A) cause noticeable deterioration in both the model climate and NWP scores. However, alongside recent changes in aerosol modelling, CAMS has made improvements to the aerosol optical properties. When we use these new properties with the new aerosol climatology, the model behaviour becomes much closer to the model behaviour with the old climatology and with the old aerosol properties. NWP testing shows a slight overall cooling of the troposphere, with improved patterns of atmospheric heating in the tropical and sub-tropical lower troposphere, leading to a modest reduction in root-mean-square errors of medium-range temperature and wind. Longer runs show a variety of modest impacts on climate, either positive or negative depending on whether the change reduces or reinforces biases from other sources. One small but helpful improvement in the model climate is a reduction in the excessive easterly winds in the eastern equatorial Indian Ocean (see Figure 3). These erroneous winds drive

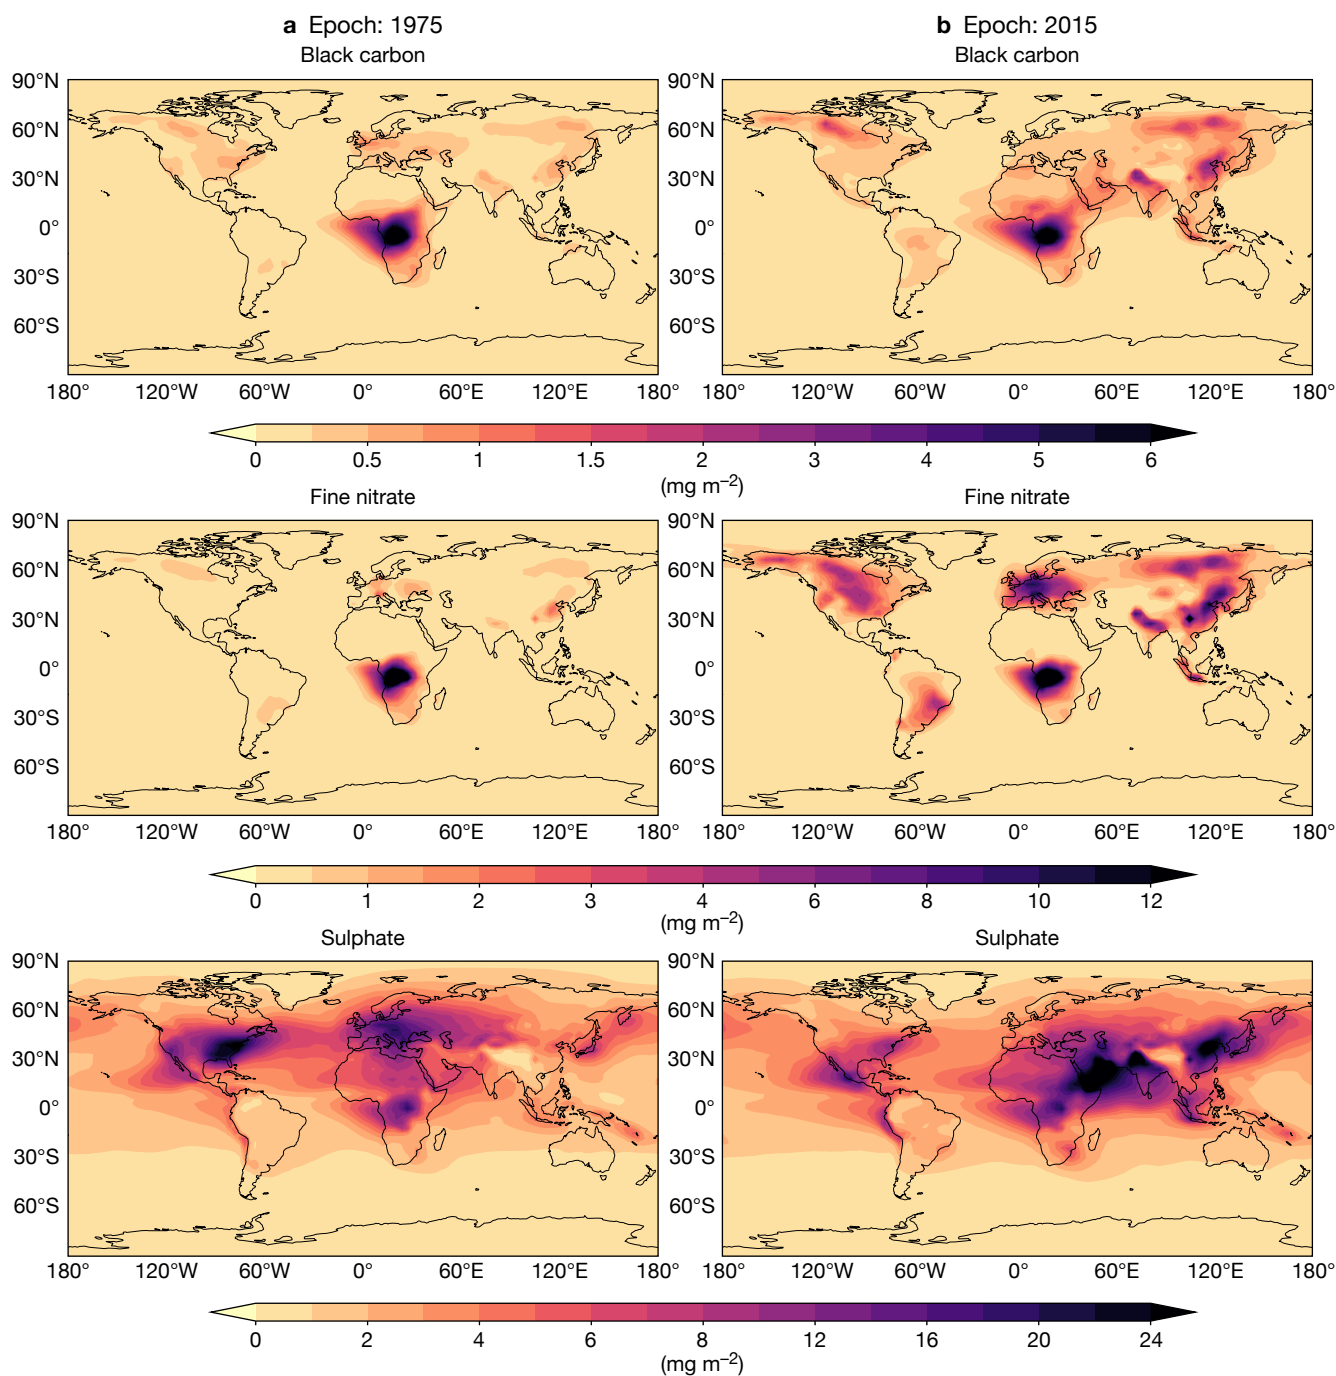


FIGURE 2 The time-varying vertically integrated July mass climatologies shown for (a) 1975 and (b) 2015 for three epoch-varying species.

excess upwelling, giving a cold sea-surface temperature (SST) bias. Since this is an error in a very non-linear part of the tropics, it has been a high priority to try to reduce it. The improvement from the new aerosols is only moderate, but still welcome.

Trends in aerosols and interaction with clouds

One of the aims of including time-varying aerosols in the IFS is to improve the representation of temperature trends in seasonal forecasts. Operational seasonal forecasts of temperature, presented as the expected

temperature for a coming season relative to past decades, have become largely dominated by the effects of global warming, and our model must reproduce these well if the real-time forecasts are to be accurate.

Aerosols interact with radiation in many and complex ways, but on seasonal timescales changes in surface temperature are typically dominated by changes in top-of-the-atmosphere (TOA) net solar radiation. We illustrate the impacts of time-varying aerosols by looking at the difference between the recent 20-year period 2001–2020 and the preceding 20 years, 1981–2000, which we will refer to for brevity as ‘the




FIGURE 3 The bias in 850 hPa zonal (east–west) winds in September–October–November, in seasonal forecasts from 1 May, with (a) the old and (b) the new aerosol climatology, in the period 2001–2020.

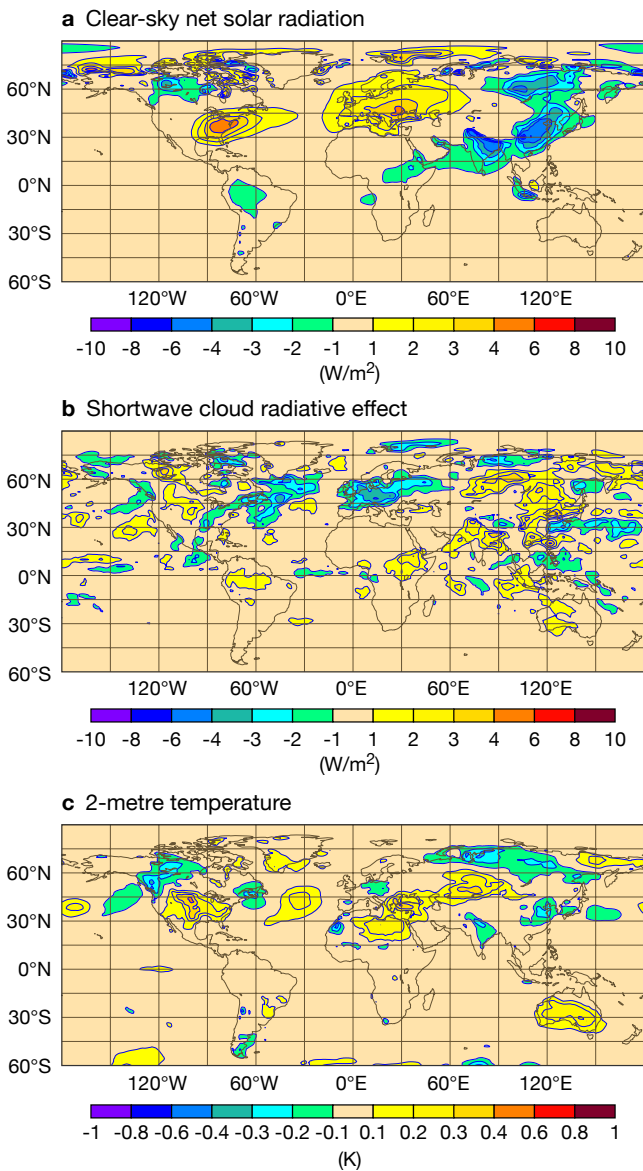


FIGURE 4 The difference in June–July–August (JJA) ensemble-mean 20-year trends between IFS experiments with the new time-varying and the old fixed aerosol climatology, for (a) top-of-the-atmosphere clear-sky net solar radiation, (b) the shortwave cloud radiative effect, and (c) 2-metre temperature.

trend’. From the earlier to the later period, there are many changes in aerosols, including a substantial reduction in reflective sulphate aerosol over Europe and eastern North America, together with an increase over South and East Asia.

Figure 4a shows the impact of the aerosol change on the trend of northern summer clear-sky TOA (top-of-the-atmosphere) net solar radiation (that is, the net solar radiation calculated at each grid point assuming no clouds are present). Changes of up to 5 W/m² can be seen in the regions of largest sulphate aerosol change. However, we also need to account for how aerosol changes affect the clouds. The shortwave cloud radiative effect is the difference between the net TOA solar radiation with the actual cloud cover, and that calculated assuming no clouds. It is thus the effect of the clouds on the net top solar radiation. The difference in the trend of the cloud radiative effect, Figure 4b, shows that in the IFS the reduction in aerosol over Europe has led to an increase in the reflection of sunlight to space by clouds; this surprising result is explained below. In much of northern Europe, this cloud change entirely offsets the reduction in reflection by aerosol itself, and in places it results in the aerosol change causing a slight cooling in surface temperature (Figure 4c). In the Eastern Mediterranean, the absence of clouds means that aerosol-induced temperature changes are mostly driven by direct radiative aerosol forcing changes, which drive additional warming of several tenths of a degree. The result is that here the overall temperature trends due to all causes, including greenhouse gases and aerosols (Figure 5a), are closer to those from ERA5 (Figure 5b). Note that the impact of aerosols on IFS cloud trends over Europe is on top of a reduction in cloud over time due to other causes, most likely a positive feedback from global warming.

Comparison of the trends in 2-metre temperature between seasonal forecasts and reanalysis in Figure 5 suggests we have a problem over Europe, though.




FIGURE 5 20-year trends in June–July–August (JJA) 2-metre temperature from (a) the 1 May ensemble mean seasonal forecasts with time-varying greenhouse gases and aerosol and (b) the ERA5 reanalysis.

The forecast trend is based on averaging 10-member ensembles, each of which has a slightly different trend pattern. The ERA5 values, on the other hand, are based on the single observed outcome of what happened, so we expect ERA5 to be noisier, and not to exactly match the forecast ensemble mean. Nonetheless, it is striking that the real warming trends in ERA5 have been very strong across the whole of Europe in this period, in which aerosol pollution has decreased. There are well understood mechanisms whereby a reduction in sulphate aerosol should give a reduction in cloud brightness and hence stronger solar input, which is consistent with the strong warming seen in ERA5. So why does the IFS cloud response go the wrong way? The problem is that the IFS cloud microphysics takes no account of the aerosol climatology we specify for the radiation – we are missing the positive feedback between sulphate aerosol and cloud brightness that is known to exist. On the other hand, the reduction in sulphate aerosol and other absorbing aerosols means that more solar radiation hits the surface, and less is absorbed in the troposphere. This means the troposphere is slightly cooler, and more conducive to creating clouds and extending their lifetime. This is the semi-direct radiative effect of aerosols, which the IFS faithfully represents. Since the IFS has this negative impact of aerosol on clouds but is missing the larger positive one, the net impact of aerosol-induced changes over the cloudier parts of Europe has the wrong sign. When added to the clear-sky effect, the impact of reducing aerosol on TOA solar radiation is very close to neutral, instead of being a large positive effect.

This is disappointing, but the IFS cloud physics are being further developed to include aerosol indirect effects on clouds, and we expect better results in the future. In cloud-free areas such as the Eastern Mediterranean, we are seeing clear improvements in forecast temperature trends. NWP experiments suggest that the semi-direct effect is most relevant on longer forecast timescales, where tropospheric temperatures

have time to respond to weak radiative forcing, and the semi-direct effect should thus be less relevant for reanalyses.

Conclusion and outlook

A new time-varying climatology of tropospheric aerosols has been created for the IFS. It is clearly derived from the CAMS chemistry and aerosol model and can be used in the IFS with the same aerosol optical properties as are used by CAMS. This is a big step forward in the coherence of our aerosol representation across different IFS configurations and is a step towards our longer-term strategy of representing highly variable species of aerosol, such as dust or fire-related emissions, interactively in NWP configurations.

The new aerosol climatology is an improved representation of aerosol, but there is still scope for development. The Cycle 47r3 aerosol and chemistry scheme was assessed by CAMS using in-situ data from the AERONET validation network. The assessment shows that the aerosol in a free-running version of the model is biased low in terms of total aerosol optical depth. This is most noticeably the case in areas with relatively low aerosol loading. This suggests that more widely dispersed and longer-lived aerosol is underestimated. Coincidentally, the IFS has long had an additional spatially constant tropospheric background aerosol term, on top of the explicit representation of various types of aerosol. There had been some hope that the new aerosol climatology might allow this term to be removed or reduced, but it turns out that we need to retain it to preserve both NWP scores and the fidelity of the model climate. We will need further improvements in IFS aerosol modelling before we can eventually remove this additional term.

The second aim of the new aerosol climatology was to represent anthropogenically driven time-variation, and hence improve our ability to model temperature changes over time. We have succeeded in creating a plausible

time-evolving aerosol climatology, and the direct and semi-direct radiative impacts of changing aerosol are likely to be reasonably represented. However, the lack of indirect cloud aerosol feedbacks within the IFS means that the resulting changes in net radiative forcing are deficient in some regions, and thus the impact on surface temperature trends is incomplete. Work to allow cloud properties to explicitly depend on aerosol has started. This will both be helped by the improved climatology and, in time, contribute to a more realistic time evolution of clouds in response to changing aerosols. Correctly representing the impact of time-varying aerosols on clouds (and hence temperature) is a key challenge for the wider climate community.

The new aerosol climatology is useful in its own terms and will contribute to both ERA6 and SEAS6 as well as future NWP cycles. It also sets a foundation for further improvements in our treatment of aerosol and cloud-aerosol interactions.

Acknowledgement

The development of the time-varying aerosol climatology was partially supported by the CONFESS project, which received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 101004156.

Further reading

Bozzo, A., A. Benedetti, J. Flemming, Z. Kipling & S. Rémy, 2020: An aerosol climatology for global models based on the tropospheric aerosol scheme in the Integrated Forecasting System of ECMWF. *Geosci. Model Dev.*, **13**, 1007–1034.

Tegen, I., P. Hoorig, M. Chin, I. Fung, D. Jacob & J. Penner, 1997: Contribution of different aerosol species to the global aerosol extinction optical thickness: Estimates from model results, *J. Geophys. Res.*, **102**, 23895–23915.

Medicane Daniel: an extraordinary cyclone with devastating impacts

Tim Hewson, Abdelwanees Ashoor (Omar Al-Mukhtar University, Al Bayda, Libya), Souhail Boussetta, Kerry Emanuel (Massachusetts Institute of Technology), Kostas Lagouvardos (National Observatory of Athens/Institute for Environmental Research), David Lavers, Linus Magnusson, Fatima Pilloso, Ervin Zsoter

In early September 2023, a set of extreme rainstorms led to devastating flooding in parts of Greece, Bulgaria and Türkiye. These events related to development of a surface cyclone nearby on the night of 4 September, assigned the name ‘Daniel’ as part of a EUMETNET cyclone naming initiative (Cusack et al, 2017). In subsequent days, Daniel meandered slowly across the Mediterranean before adopting an east-south-eastward trajectory near northern Libya late on 8 September, whereupon it became a medicane. Landfall was near Benghazi around 23 UTC on 9 September. Remarkably, the medicane deepened further over land. The resulting intense rainfall over the Akhdar (Green) mountains of northern Libya on the night of 10–11 September drained into the small Wadi Derna catchment and went on to cause catastrophic flooding in the city of Derna. Two dams burst, and there were 5,000–15,000 fatalities as buildings were swept away. This was likely the deadliest rainfall-related flooding disaster since ECMWF started producing operational forecasts in the late 1970s and the second most deadly dam-related disaster of all time. This article examines some of the meteorology and hydrology at play, whilst providing insights into predictability aspects, forecast quality and event rarity. In so doing we synthesise many aspects of ongoing work at ECMWF, and beyond.

Supplementary forecast and observation charts for this episode can be found in two case links in ECMWF’s online severe event catalogue: <https://confluence.ecmwf.int/display/FCST/Severe+Event+Catalogue>.

Defining Daniel’s full history

Figures 1a,b depict the full eight-day lifecycle of Daniel – its track and central pressure. This is based on UK Met Office surface analyses, augmented, from 12 UTC on 9 September onwards, by reassessment of 15-minute interval satellite imagery sequences (from the EUMETView tool: <https://view.eumetsat.int/productviewer>) and other observations, alongside ECMWF high-resolution (HRES) analysis fields.

Reassessment led to a deeper, repositioned centre over northern Libya on 10 and 11 September, and track extension across Egypt (beyond the Met Office chart domain) on 11 and 12 September.

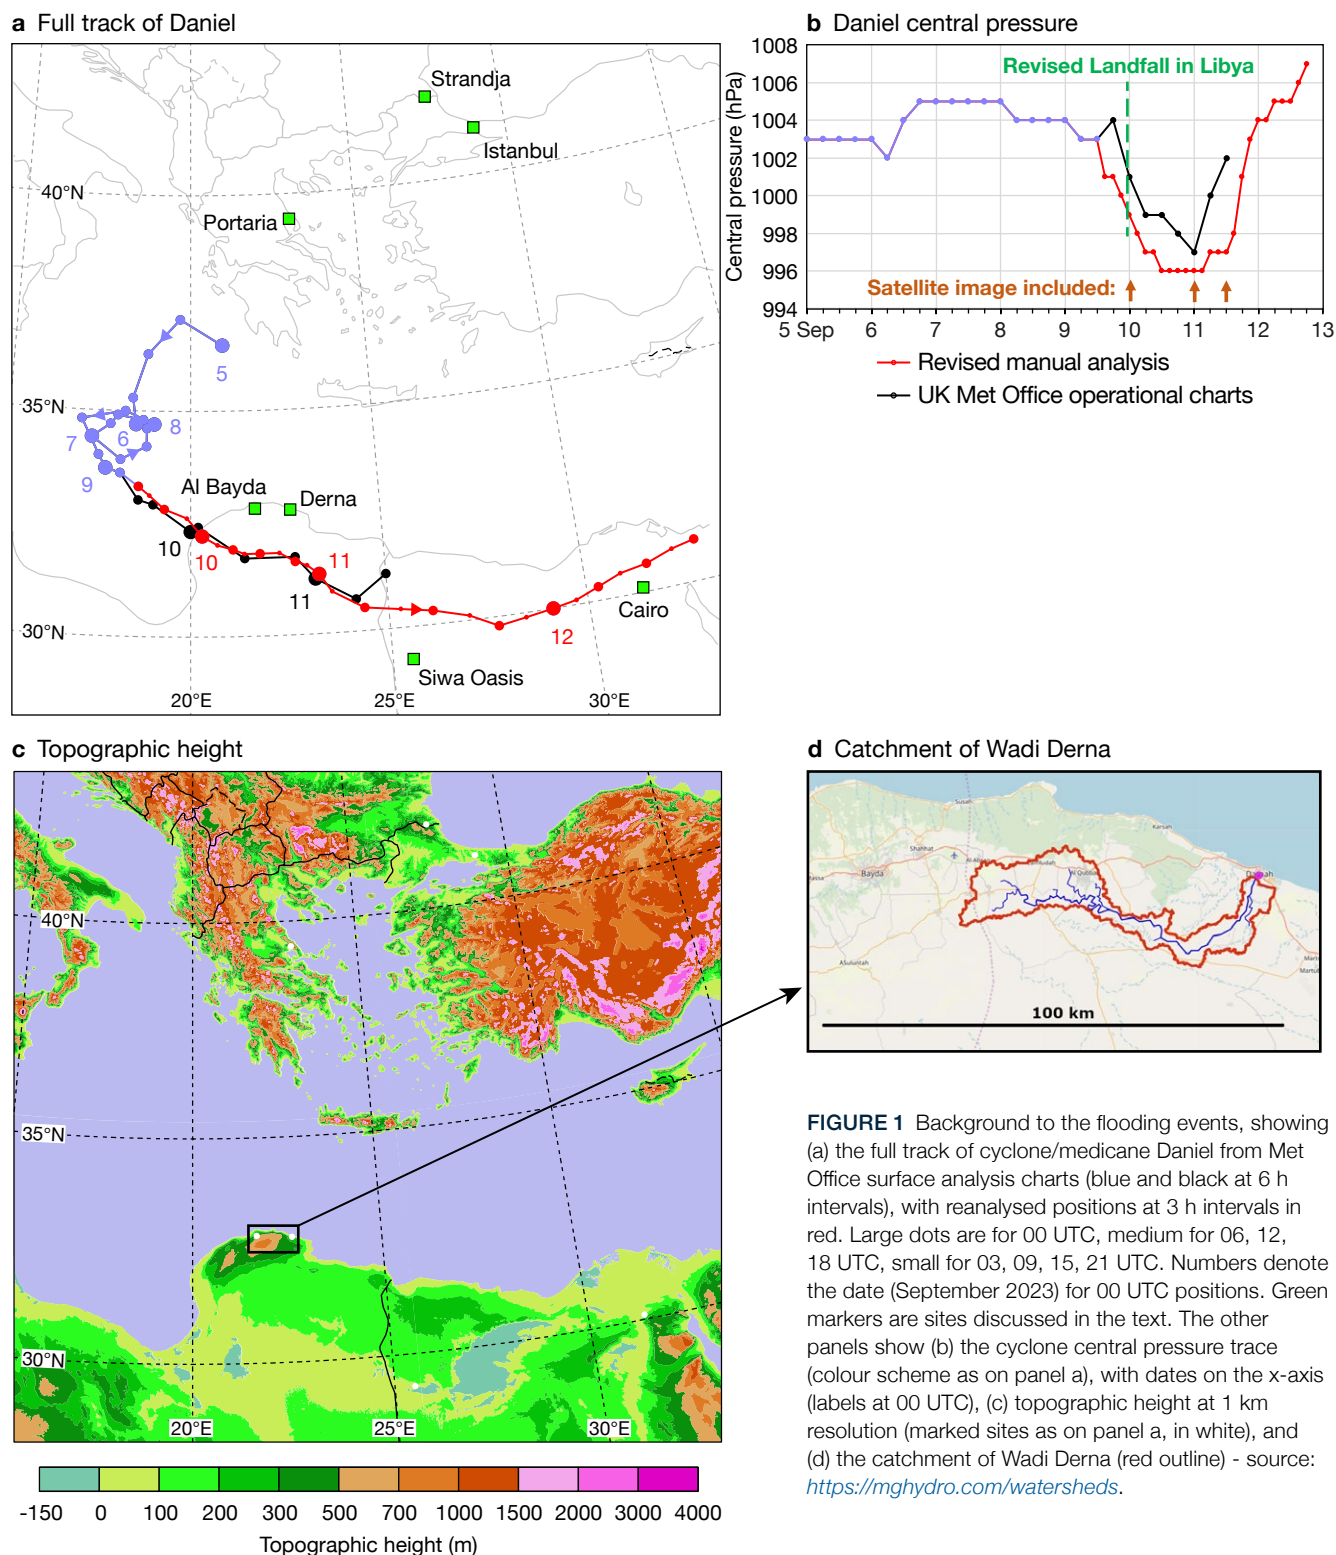
Evidence for the adjustments included hourly METAR observations of mean sea-level pressure and 10 m mean wind, from Siwa Oasis on 11 September and Cairo on 12 September (Figure 1a), that conflicted with model data. Siwa Oasis winds reached 34 kts at 11 and 12 UTC on 11 September as the low passed by. Meanwhile, visible imagery early on 11 September showed a substantial tower of elevated dust/sand in the cyclone centre (Figure 2c), which would have needed near-centre 10 m wind speed in excess of 20–25 kts to accrue, which even 9 km resolution model solutions did not show. Likewise, aerosol optical depths from a Copernicus Atmosphere Monitoring Service (CAMS) numerical model run (40 km resolution), used to pinpoint dust and sand, showed no sign of near-centre concentration (compare Figures 2c,d). This all pointed to a deeper low with a tighter core, which had probably existed for some time.

Further investigations showed Siwa Oasis METAR data were not available for assimilation at ECMWF; pressure would have been used otherwise. Over land, 10 m winds are not assimilated, although for flat terrain a case could potentially be made for their use, especially for sites in data-sparse areas such as this one.

ECMWF runs an ensemble of 50 data assimilations (EDA members), to try to represent analysis uncertainties. These can provide options for reassessment when the main analysis is in doubt. There were typically 5 hPa ranges in EDA central pressures for Daniel from 10 September 00 UTC onwards; our revisions were mostly within these ranges. Only at 12 UTC on 11 September did our revised pressure step outside: 997 hPa versus 998 hPa in the deepest EDA member, justified by the aforementioned METAR non-availability.

Why did Daniel develop?


We now examine the broadscale synoptic drivers for Daniel. Conceptually, the synoptic-scale environment conducive to development and persistence of such a




feature is a quasi-stationary upper-level trough or cut-off low. This was present over many days in early September 2023 (blue and green contours on Figure 3). This is unusually far south for September, and indeed requires, in turn, a quasi-stationary omega block or strong upstream ridge over central Europe for 'maintenance' (a later snapshot of this can be seen in Figure 10b – green contours). For rainfall associated with Storm Daniel over Greece and the Mediterranean

on 4–8 September, the signal for an extreme event appeared in the ECMWF ensemble (ENS) between 30 August and 1 September. A dramatic change can be seen in the EFI (Extreme Forecast Index) and SOT (Shift Of Tails) for total precipitation comparing forecasts from the two dates, separated by 36 hours (Figure 4). The 30 August errors were on a large scale, and related to prediction of the upstream blocking over Europe. Tracking this error back in time, we found a connection


a Infrared image at 00 UTC on 10 September 2023



b Infrared image at 00 UTC on 11 September 2023



c Visible image at 11:45 UTC on 11 September 2023



d 12 h forecast for 12 UTC on 11 September 2023

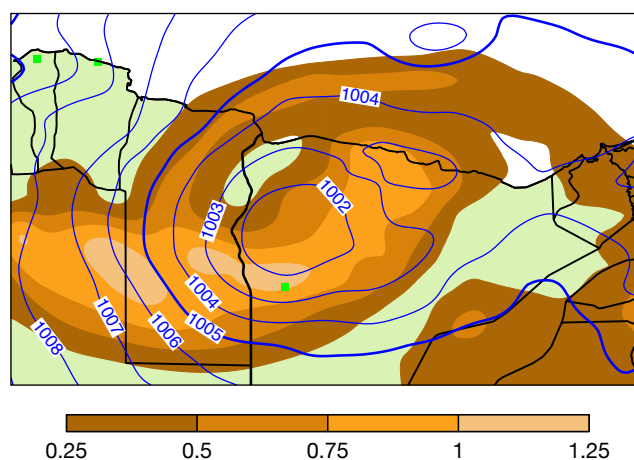


FIGURE 2 Imagery-related products at selected times during Daniel's life cycle (marked also on Figure 1b). We show infrared $10.6 \mu\text{m}$ channel images at (a) 00 UTC on 10 September and (b) 00 UTC on 11 September. The other panels show (c) a visible image from the US National Oceanic and Atmospheric Administration (NOAA) at 11:45 UTC on 11 September and (d) the 12 h forecast of mean sea-level pressure (blue, 1 hPa interval) and aerosol optical depth (shading) for 12 UTC on 11 September, from a Copernicus Atmosphere Monitoring Service (CAMS) operational run, with marked sites as in Figure 1.

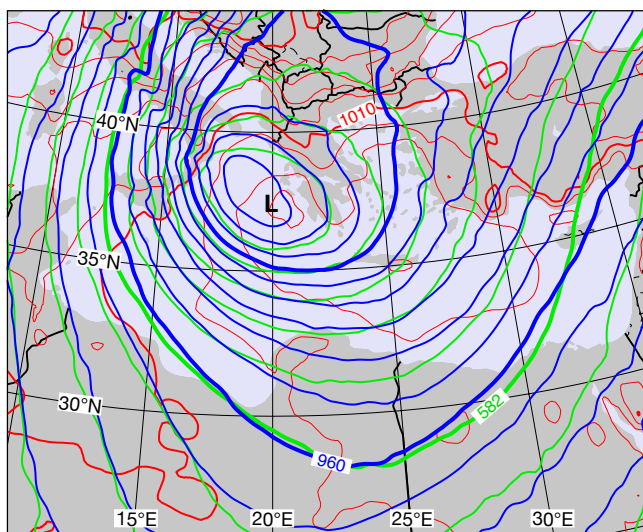
to the interaction between hurricane Franklin and the mid-latitude wave guide, which took place around 1 September over the central Atlantic. Small phasing errors between Franklin and an upper trough quickly amplified around that time and led to significant errors downstream in the block and the cut-off low. Indeed, this extreme sensitivity was manifested in an extremely large difference emerging between the HRES forecast and the ENS control forecast, which currently have minute configuration differences (e.g. cm in topographic heights) for accidental technical reasons, but are otherwise identical. Once large uncertainties in the extratropical transition had gone, the large-scale predictions of the environment for Daniel improved and the signal for extreme rainfall increased.

What drove the local behaviour of Daniel?

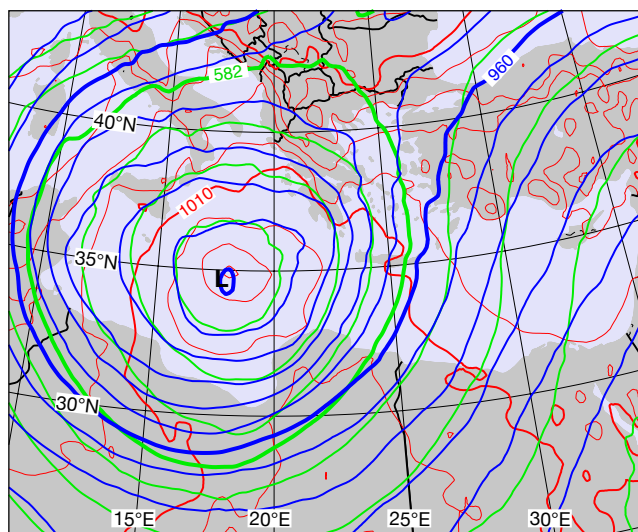
Here we discuss Figure 3, which shows mean sea-level pressure, and geopotential height in the mid- to upper troposphere, at four representative moments in Daniel's lifecycle.

- On the night of 4–5 September, genesis occurred, apparently in response to upper-level left-exit forcing, ahead of jet cores (or stronger geostrophic gradients) which are themselves co-located at 500 and 300 hPa west of Daniel (Figure 3a). This is a typical juxtaposition for extratropical cyclone development.
- Several days followed with little variation in minimum pressure (Figure 1b). One can infer minor pulses of upper forcing, but mostly the atmosphere was in a barotropic or equivalent barotropic state, as demonstrated by concentric contours at all three levels in Figure 3b. In this scenario, the latent heat release from active bursts of convection around the low (seen on imagery), if symptomatic of strong enthalpy fluxes over the high autumn sea-surface temperatures (SSTs), will balance out the tendency for frictionally induced decay, and keep the low going. In other medicane cases, when organised convection encircles the low, one tends to see deepening as the low feeds off strong fluxes.

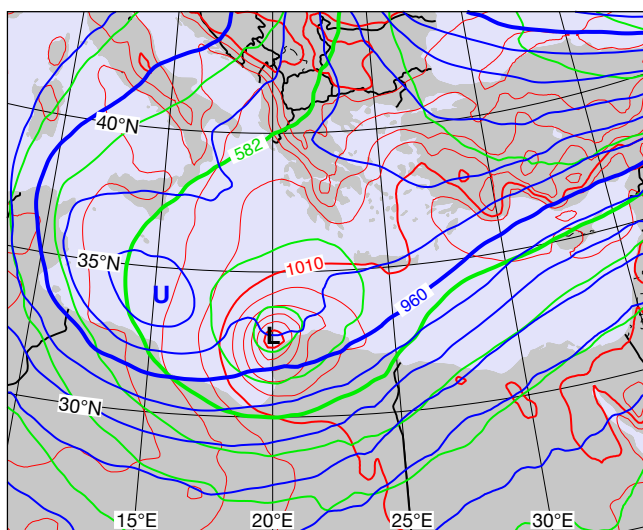
a 00 UTC on 5 September



b 12 UTC on 6 September



c 00 UTC on 10 September



d 00 UTC on 11 September




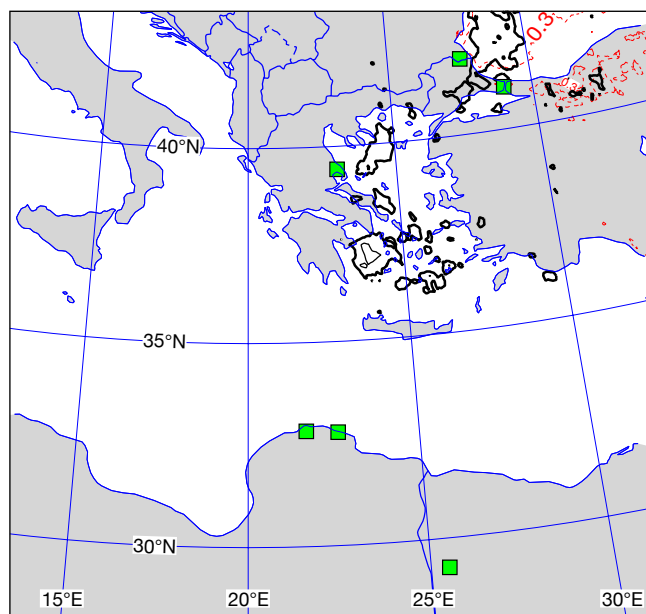
FIGURE 3 Mean sea-level pressure (red, 2 hPa interval) and geopotential height (green for 500 hPa, blue for 300 hPa, both 3 dm interval) from HRES analyses at selected times during Daniel's lifecycle. L marks the surface low centre, U an important upper-level (300 hPa) feature (see text).

However, for Daniel this organisation was lacking until 8 September. This was in spite of SSTs being $\sim 1^\circ\text{C}$ above normal, and in spite of 'potential intensity' here being high (computed to be 50–60 m/s following Bister & Emanuel, 2002; see Figure 5b for a related construction for a later time). The night of 8–9 September then brought some convective organisation and the onset of modest deepening (Figure 1b).

- On the night of 9–10 September, Daniel made landfall. Whilst the imagery signal had again become more muted (Figure 2a), and whilst the low and middle troposphere remained in a relatively barotropic state (Figure 3c), at 300 hPa a marked upper-level low was approaching from the west (marked 'U'). In response to the associated

vorticity/potential vorticity advection, and a saturated column beneath which promotes a large Rossby penetration depth, the surface low deepened further over land. This might seem somewhat strange when medicanes commonly decay over land, but in truth Daniel was a hybrid system at this stage, positioned somewhere within a broad 'grey zone' that spans surface-flux-driven development to baroclinically-driven development. Likewise, forcing from just very high levels is also rather atypical, and different to the situation in Figure 3a. During 10 September, cold-topped cloud became much more widespread (compare Figure 2b with 2a), commensurate with further deepening, and the intensity of the flooding rains probably peaked.

a Forecast from 12 UTC on 30 August



b Forecast from 00 UTC on 1 September

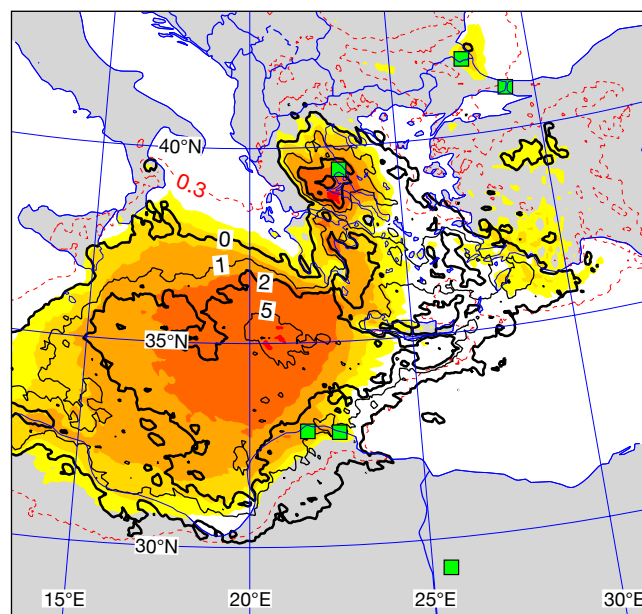


FIGURE 4 Extreme Forecast Index (EFI, shading and red contours) and Shift Of Tails (SOT, black contours; 0, 1, 2, 5, thicker for 0 and 2) for 5-day accumulated rainfall on 4–8 September 2023, from operational forecasts with start times of (a) 12 UTC on 30 August 2023 and (b) 00 UTC on 1 September 2023. Marked sites (green squares) are as in Figure 1a.

- Some time after 00 UTC on 11 September, when cyclonic features were almost vertically stacked again throughout the troposphere (Figure 3d), the surface low began to fill, as frictional decay took over in the absence of contrary dynamical and thermodynamical drivers.

The ‘unexpected’ deepening of tropical cyclones over land, which occasionally takes place over Australia (where such cyclones have been called Agukabams) and exceptionally over the US (section I in Emanuel, 2018), has been hypothesised to be due to uptake of heat from hot soil/sand, promoted by precursor rainfall wetting the sand and accordingly increasing its thermal conductivity. Such situations are finely balanced, with the right amount of rain needed ahead of the cyclone, on the right timescales. There was a distinct possibility that such a mechanism was active in the current case, alongside upper-level forcing. Within its soil component, ECMWF’s Integrated Forecasting System (IFS) includes a representation of thermal conductivity dependence on moisture content. We ran ENS experiments with this turned off to ascertain whether the Agukabam mechanism could have been at play. Although the soil seemed to be sufficiently hot (of order 30°C at the outset), there were no systematic differences between operations and the dry soil experiment in terms of either Daniel’s depth or the associated enthalpy fluxes. This means that either the mechanism was unimportant for this case, or that imperfections in the IFS soil physics,

soil classifications or other factors meant that we could not elucidate an impact. The fact that Daniel appears to have been at the deeper end of the EDA range hints that an active process may have been under-represented or unrepresented, but we do not know what.

Rainfall and forecasts for southeast Europe

Whilst the extreme rainfall in Libya in September unequivocally relates to the passage of Daniel, flooding rains further north owed their existence more to uplift of an influx of intrinsically warm, moist air from marine sources to the northeast: the Black Sea in the case of Bulgaria and Türkiye, and the Aegean in the case of Greece (note the airflow implied by the red contours on Figures 3a,b).

Uplift can have perhaps three drivers: dynamical forcing (related in this case to advection of vorticity filaments around the nearby upper lows); surface triggering of unstable air (provided in this case by high sea-surface temperatures – a common autumn rainstorm-driver in the Mediterranean); and the projection of steady airflow onto topographic slopes. In the current case all appear to have been active, to varying degrees, and interactive.

Three sites where particularly high intensities/quantities of rainfall occurred, and where severe flooding impacts were reported in the media, are marked in Figures 1a,c and shown in Figure 6b (labels B, C, A): Strandja in Bulgaria,

330 mm on (just) 5 September; Istanbul in Türkiye, 125 mm on 5 September; Portaria in Greece, 762 mm on 5 September and 1,096 mm 5–8 September. Flooding over Greece was very widespread, due to rain over many days, and focussed on the Thessaly region in central Greece in particular. Here large lakes were created or re-created (Figure 6b: land had been reclaimed from Lake Karla in 1962 for cultivation purposes).

The HRES model sounding for Portaria in Figure 5a is believed to be indicative for the three locations (Strandja, Istanbul and Portaria). One can identify high moisture content (12 g/kg at low levels), a propensity for SST (26°C) triggering of ‘skinny’ Convective Available Potential Energy (CAPE) up to the tropopause, and strong onshore flow (~35 kts) that projects onto a mountain chain east of Portaria (Figure 1c, ridge height reaches 1,600 m). The topographic height/forcing is less at Strandja and Istanbul, but likely still relevant. Moreover, the persistence of the flow, allied to slow-moving Daniel (Figure 1a), was a key factor for Greece.


On the broadscale, we see from Figure 4 that there is a signal regarding extreme rainfall from 1 September onwards, which is 3–4 days in advance. The EFI/SOT signal strengthened at shorter leads (not shown here).

Within the EU-funded Destination Earth (DestinE) initiative of the European Commission, ECMWF has been running in experimental mode higher-resolution (4.4 km) global simulations (Vannière et al., 2024). We expect these to deliver more accurate rainfall forecasts for sites due to better representation of

convection (albeit still parametrized to some extent) and topography, and reduced sub-grid variability, all because of the smaller grid-box size. Here we examine and illustrate that capability, by comparing concatenated short lead-time rainfall forecasts for 4–7 September (Figure 6) between these DestinE runs and the unperturbed operational HRES (9 km resolution). We use short leads to try to illustrate potential predictability, by disassociating from broader-scale errors that can contaminate the picture at longer leads (e.g. as demonstrated with Figure 4).


Figure 6 shows that the overall pattern, with the most extreme rainfall in central Greece (b), is well captured in HRES (a) and even better in DestinE (c). Indeed, DestinE does better near Portaria (see label A on Figure 6b) with higher peaks, and a tendency to shift the maximum inland, away from coastal waters. Note the strip of cyan over the sea in HRES – this is symptomatic of a known model deficiency, linked to convective parametrization and an inability to move convective cells. Near Strandja (B), there are hints of similar advantages in DestinE, but both realisations fall well short of the observation (80–100 mm versus 330 mm). Neither simulation provides any indication of extreme rainfall in Istanbul (C), which is disappointing. The operational ENS was also not very helpful here (Figure 4b) and this signal did not improve at shorter lead times. The cluster of high totals near Athens (D: 50–150 mm) is also not well represented (~20–40 mm forecast). Finally, note that there are hints of the aforementioned coastal waters issue in the DestinE runs near the three ‘fingers’ of the Chalkidiki peninsula (E). In summary, DestinE performed a bit

a Portaria, 12 UTC on 5 September



Location: Lat: 39.4, Lon: 23; Model orography: 562 m

b Al Bayda, 00 UTC on 11 September



Location: Lat: 32.8, Lon: 21.7; Model orography: 397 m

FIGURE 5 Model soundings (red) from HRES analyses for (a) Portaria, 12 UTC on 5 September 2023, and (b) Al Bayda, 00 UTC on 11 September 2023, during their respective extreme rainfall events. Both sites had seas upwind, nearby, at these times: blue dots show sea-surface temperature (SST); blue lines are related parcel constructions up to lifting condensation level, based on sounding and SST data. The sites are marked in Figure 1. The purple curve on (b) is a parcel construction, from saturation at SST, and cyclone central pressure: the resulting CAPE (hatched) is proportional to medicane ‘potential intensity’: here ≈ 60 m/s (= maximum possible sustained 10 m wind).

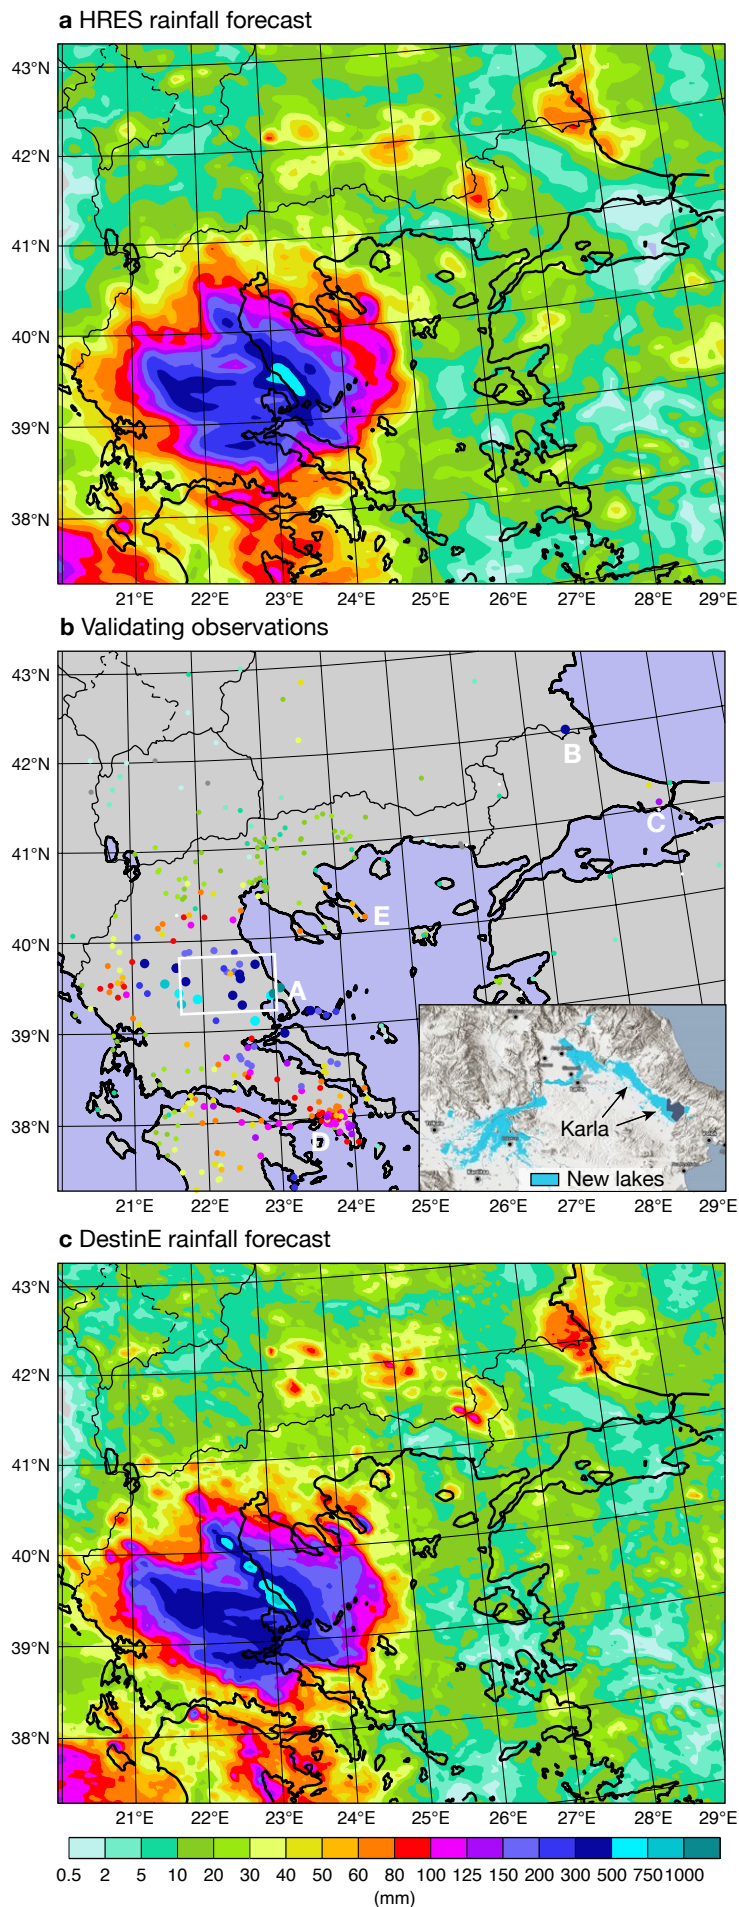
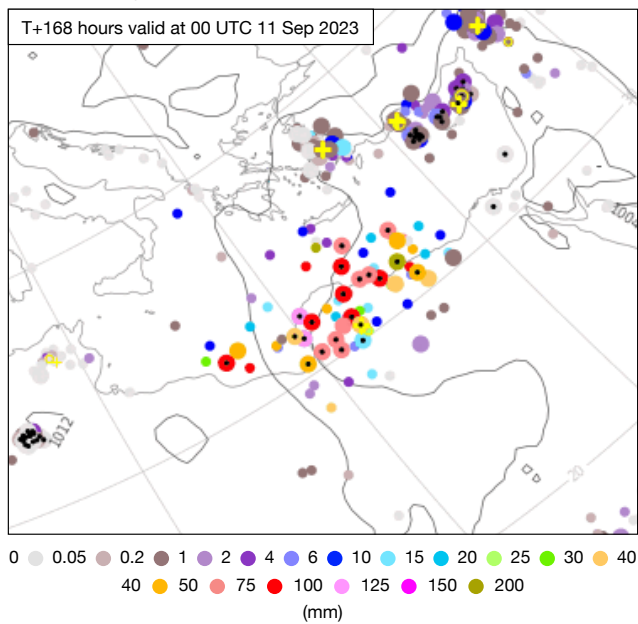


FIGURE 6 The charts denote 4-day short-range rainfall forecast verification, to investigate resolution-dependent model capability. They show 4-day rainfall from (a) HRES (9 km resolution) and (c) ECMWF Destination Earth (DestinE) runs (4.4 km resolution), summing together in each case four 0–24 h forecast values from 00 UTC runs on 4, 5, 6 and 7 September 2023. They also show (b) validating observations, with quality control and manual gap filling applied for gauges in the Thessaly region in central Greece, where a few data outages occurred. Labelled regions are referenced in the text. The inset box in (b) shows where new lakes formed due to the extreme rainfall.

a Cyclonic centres from HRES and ENS members, coloured using a rainfall metric



b 24-hour rainfall according to HRES and five ENS members

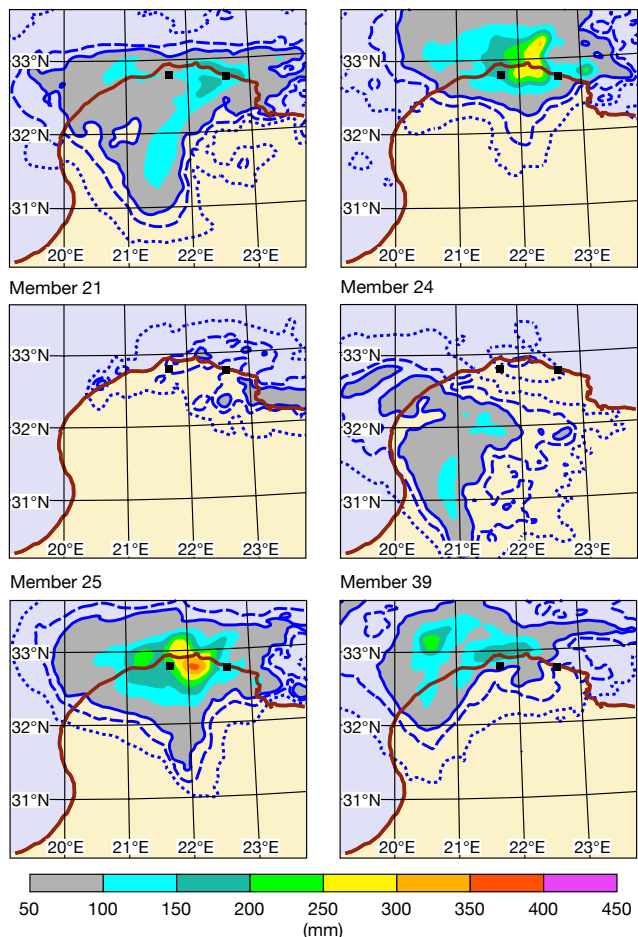


FIGURE 7 IFS forecasts for 10–11 September 2023, (a) from 00 UTC on 4 September and (b) from 00 UTC on 8 September. Panel (a) shows a standard ‘cyclone database’ forecast product, in which a spot denotes a cyclonic centre in HRES (yellow cross) or in an ENS member, and the colour indicates, for that forecast realisation, the attribute of ‘maximum 12 h rainfall’ between 18 UTC on 10 September and 06 UTC on 11 September within a 300 km radius of the centre. Black dots signify ‘barotropic’ low centres. In (b), six selected forecasts (HRES and numbered ENS members) of 24 h rainfall from 06 UTC on 10 September to 06 UTC on 11 September in northern Libya are shown; dotted, dashed and solid blue contours denote 10, 30, 50 mm respectively, with shading as on the legend, and with marked sites (black squares) as on Figure 1a.

better than HRES, but the need to still parametrize convection to a certain extent still creates issues.

Rainfall and forecasts for Libya

In contrast to the absence of any early signal of extreme rainfall in southeast Europe at ~6 day leads (Figure 4a), for events over Libya the potential for extreme medicane-related rainfall was already clearcut in ENS output seven days before. This is apparent in Figure 7a, which shows an IFS cyclonic feature forecast for T+168 h, from the operational cyclone database suite, in which the ‘maximum 12 h rainfall’ cyclone attribute is denoted by spot colour (see legend and caption). We can note many features over or near northern Libya; 21 of these (40% of members) have black spots at the centre, denoting a low centre and no clearcut fronts (the ‘barotropic low’ class). For these, the maximum 12 h rainfall lies between 50 and 200 mm. The higher predictability, relative to the southeast Europe events, is likely due to the predictability barrier related to Hurricane Franklin having already been crossed, and the necessary broadscale pattern over Europe (omega block + cut-off low) having been established.

Figure 7b then shows examples of spatial 24 h rainfall forecasts from four days later, for a 24 h period, in HRES and various ensemble members. Verifying rainfall totals (the only ones officially available) are shown in Figure 10d. In these forecasts, whilst a cyclone track over Libya or very close by has become almost certain, rainfall amounts, and regional distribution, remain highly uncertain. Relative to the Wadi Derna catchment (Figure 1d), two of these members (8 and 25) have very extreme rainfall (say > 200 mm/24h); HRES is not as extreme, member 39 keeps rainfall mainly over the ocean, whilst members 21 and 24 do not show enough rainfall to even cause any streamflow in the ordinarily dry river valley. The local limit before runoff starts has been calculated to be about 30 mm when soil is initially dry. From a warning perspective, this situation would be challenging to manage – we have a huge range of possible outcomes, despite high confidence in the broadscale pattern.

As regards *actual* rainfall generation mechanisms, the representative sounding in Figure 5b exhibits similar characteristics to the Greek sounding in Figure 5a. There is deep instability to nearby SST triggering, again with skinny CAPE and parcels potentially reaching the tropopause. Moisture content low down is higher, 15 g/kg




FIGURE 8 3D-map of part of northern Libya, looking approximately southwestwards. Derna flooded due to extreme rainfall and the collapse of two dams (labelled). Wadi Derna is visible as a steep-side valley extending south-southwest from Derna, which rises as it curves away to the west (see arrows); Figure 1d shows the full extent of its catchment, which has a top height of 765 m above sea level (asl). Valley base is 188 m asl at Bu Mansur dam (dam height was 75 m), and 24 m asl at El Bilad dam (height was 40 m). Map © Google (13/09/2023): includes data from Airbus, Landsat / Copernicus, Data SIO, NOAA, U.S. Navy, NGA, GEBCO.

versus 12 g/kg in Figure 5a, whilst lower tropospheric winds are a bit stronger. These impinge on the coast from the north, and indeed onto north-facing escarpments (one is visible a little way inland in Figure 8). Thus winds would bring convection inland, convection which can also be augmented by ascent over the slopes. So whilst topography is not as high as it is near to Portaria, other ingredients favoured more rain. In addition, we have strong forcing related to the medicane deepening, and to the aforementioned upper-level structures. On the other hand, the cyclone is now mobile, while in Greece conditions were relatively static, prolonging the rainfall there.

Strong forcing would tend to mitigate against the erroneous focusing of coastal water convection that we saw in Figure 6; indeed the seemingly comparable member 39 pattern on Figure 7b was atypical.

Flood forecasting and dam collapse

To set the scene for hydrological considerations, Figure 8 shows a 3D rendering of the topography of the Wadi Derna valley and catchment, and surrounding areas. Whilst Wadi Derna is ordinarily dry, one can nonetheless infer that the landscape has, over geological timescales, been shaped by major flooding events, and associated landslides, with many steep-sided valleys etched into the elevated terrain.

Prior to Daniel there were two dams, marked in Figure 8. The Bu Mansur dam had a capacity of 22.5 million m³, whilst El Bilad, downstream, had a capacity of 1.5 million m³. Clearly, then, any failure of the upper dam would overwhelm the lower dam. Note also that the full catchment is relatively small, spanning just 575 km².

Under the auspices of the European Commission's Copernicus Emergency Management Service (CEMS), managed by the Joint Research Centre (JRC), ECMWF runs the CEMS hydrological forecasting centre. As part of the CEMS remit, the Global Flood Awareness System (GloFAS) aims to alert users to the possibility of flooding anywhere in the world, based on hydrological modelling at ~5 km. Many layers are publicly available to display potential flood situations interactively with a web tool, with the predictions relying on ensemble rainfall forecasts from ECMWF (9 km resolution) as their primary input. Meanwhile, weather analyses from ERA5 (31 km resolution) are used in a retrospective mode as forcing for the GloFAS hydrological model, to create climatological reference points for river discharge, expressed as return periods.

We should add some important caveats regarding GloFAS usage for this case. Firstly, although GloFAS simulations include the impact of about 600 of the world's largest reservoirs (and their dams), there is

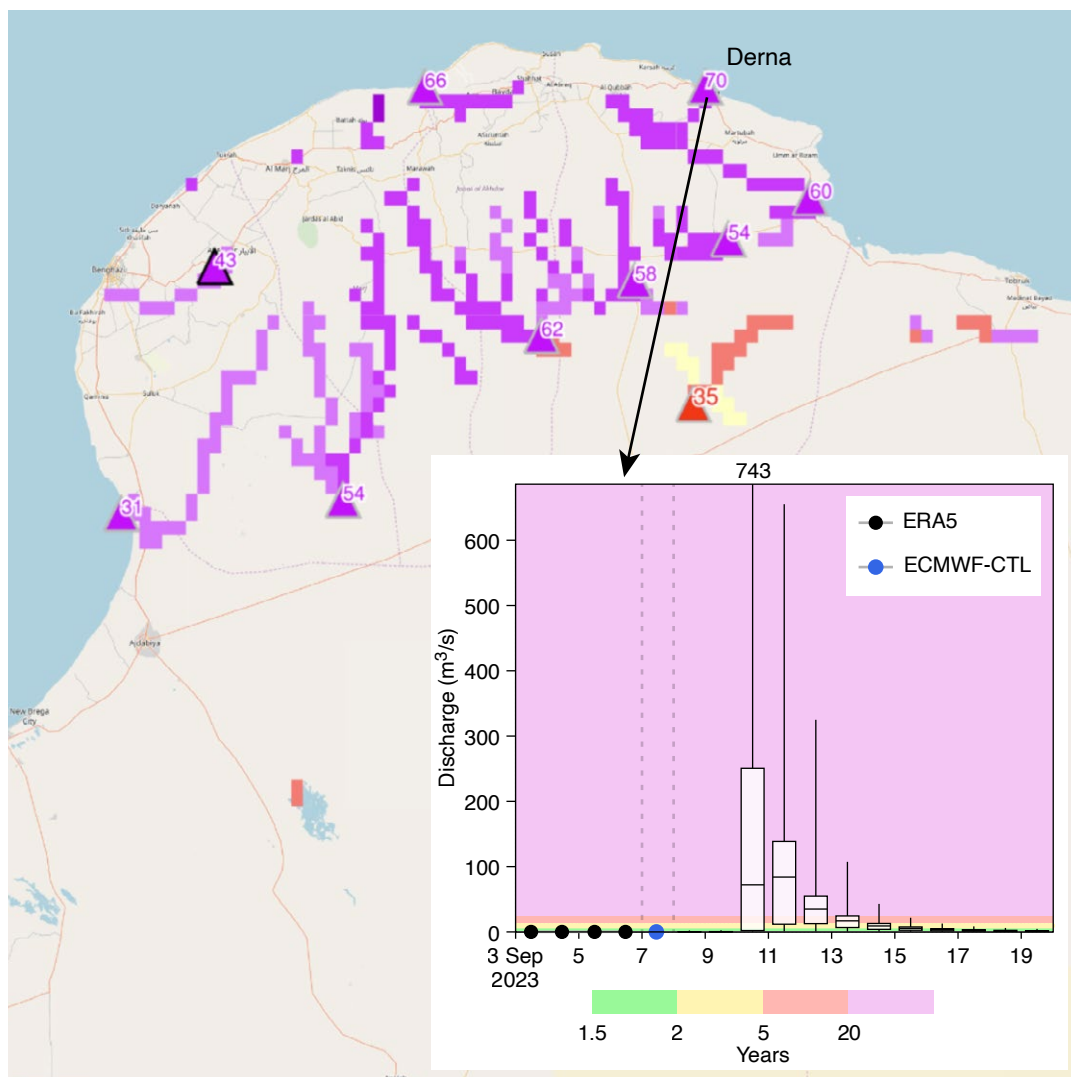


FIGURE 9 Global Flood Awareness System (GloFAS) forecast output snapshot from 00 UTC on 8 September 2023. Red and purple block shading denotes rivers/wadis where flooding with return periods respectively > 5 and > 20 years are foreseen in the next 30 days, whilst triangles with numbers denote the corresponding exceedance probabilities at selected points (see text). The inset shows an ENS discharge hydrograph of 24 h average values forecast for Derna town; box-and-whisker style denotes ENS-based ranges for percentiles 0, 25, 50, 75, 100, with shading signifying return periods in years according to an ERA5-based (discharge) climatology. Fourteen days have been cropped from the right of the hydrograph product – all have zero discharge in all members. (Source: Copernicus Emergency Management Service)

nothing of this type for the small catchment of Wadi Derna. Moreover, GloFAS certainly does not allow for dam collapse. In reality, the Wadi Derna valley was dry at the start of the event (i.e. no reservoir water), whilst spillways for each dam were kept fully open throughout. Secondly, whilst the scale mismatch between ERA5 (31 km) and ENS (9 km) resolutions is not that critical for large catchments, it can become problematic for small catchments, due to rainfall sub-grid variability and catchment rainfall representativeness issues. This is very relevant here because, in areal extent, Wadi Derna spans about 7 ENS grid boxes, but only 0.6 ERA5 grid boxes. In practice, this means that whilst the ENS-based forecast discharge values can be reasonable given the right rainfall, the ERA5-based return period referencing is sub-optimal. The third and final point is that ERA5 quality over time, over Libya rather than for example central

Europe, will also have been compromised by a dearth of conventional observations for assimilation.

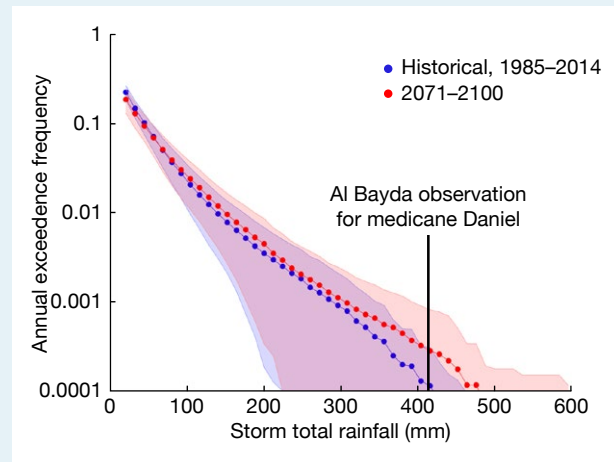
Figure 9 shows a set of GloFAS predictions for northern Libya, along with a Derna-based hydrograph, all from the same data time of 00 UTC on 8 September (as for Figure 7b). Normally the valley is dry, denoted by zero discharge from 3–9 September. However, many ensemble members quickly ramp up the discharge well into the (nominal) > 20 year return period class (purple background) on 10 and 11 September. On the map, purple shading shows modelled river pixels where the probability of exceedance of a 20-year return period discharge, during a 30-day forecast window, is $> 30\%$ for at least one day (numbers give maximum probability at selected points). So, from this product alone (albeit with the small catchment caveats) we could infer a high

a Long-term medicane risk in Libya

To assess the rarity of Storm Daniel in Libya, one can also use climate model simulations coupled with precipitation downscaling. We downscaled 20 CMIP5-generation simulations to produce 3,045 medicanes for each model, for both the historical period 1986–2005 and the period 2081–2100 under representative concentration pathway (RCP) 8.5. A total of 121,800 synthetic medicanes resulted. For more information on the downscaling, see Romero & Emanuel (2017) and Feldmann et al. (2019); here we give a brief overview.

Beginning with daily climate model output, we apply principal component analysis (PCA) for each month of each year during the two periods. The PCs are calculated using the spatial-field correlation matrix, calculated from the daily gridded fields over the Mediterranean region, of potential intensity (a function of SST and vertical profiles of temperature and humidity: see example parcel curve in purple on Figure 5b); temperature and relative humidity at 600 hPa; and geopotential heights at 850 and 250 hPa (as wind surrogates via the geostrophic relation). We then perform a second PCA to account for the temporal structure of the data. For these temporal correlations, we consider all the moving 10-day periods within the month. Random draws plus slight perturbations of this final set of extracted PCs, once converted back into physical space, allow one to generate 10-day sequences of spatiotemporally coherent fields compatible with the reanalysis or global climate model climate. These sequences behave as genuine analogues of truly simulated synoptic evolutions.

Finally, we apply a genesis potential index to the synthetic fields, finding local maxima. If the maximum values exceed a pre-defined threshold, we construct a candidate medicane track beginning at that maximum and using a beta-and-advection model based on the synoptic-scale winds. Finally, we run a deterministic intensity model along the tracks (CHIPS



Annual exceedance frequency of medicane storm total rainfall at Al Bayda, Libya, based on 20 models. The blue and red curves show multi-model exceedance frequencies for recent (1986–2005) and future (2081–2100) periods, and the shading shows one standard deviation from the mean.

= Coupled Hurricane Intensity Prediction System), and if maximum wind speeds exceed 34 kts, we include it in the final set of synthetic medicanes.

To estimate flooding potential in the region surrounding Derna, we used the synthetic medicanes to estimate long-term medicane-related rainfall exceedance frequencies for Al Bayda (where 414 mm was recorded). These are displayed above. The blue and red curves show the multi-model mean exceedance frequencies for the historical and future periods, respectively, and the shading represents one standard deviation among the 20 models. We see little if any significant difference between the historical and future periods, but large inter-model spread. Observing 400 mm from a pure medicane would be, by this calculation, an exceedingly rare event, with the most extreme model in the future climate having a return period over 500 years. Thus Daniel was either a freak event or affected by physics not included in our standard model of medicanes.

risk of flooding, with about a 70% chance for a severe situation to develop, but still a 25–30% chance, say, of something much less severe. Nonetheless, the range of possible discharge levels (on the graph) is still immense.

To try to assess what the true total discharge was for this event (Libya has no gauged rivers), an offline analysis was undertaken, using the customary Soil Conservation Service approach (SCS method: NRCS, 2008). This gave a 31 h average discharge rate of 890 m³/s, based on an

estimated catchment rainfall value of 200 mm. This discharge can be compared with GloFAS values in the Figure 9 graph (admittedly for 24 h periods). Whilst the SCS value is very high, it is still almost within the GloFAS forecast range (max = 743 m³/s). This comparison thus lends credence to the GloFAS discharge forecasts.

Total storm discharge was computed, via the SCS method, to be 69 million m³, which is more than three times the upper dam capacity. Thus, even allowing for spillway

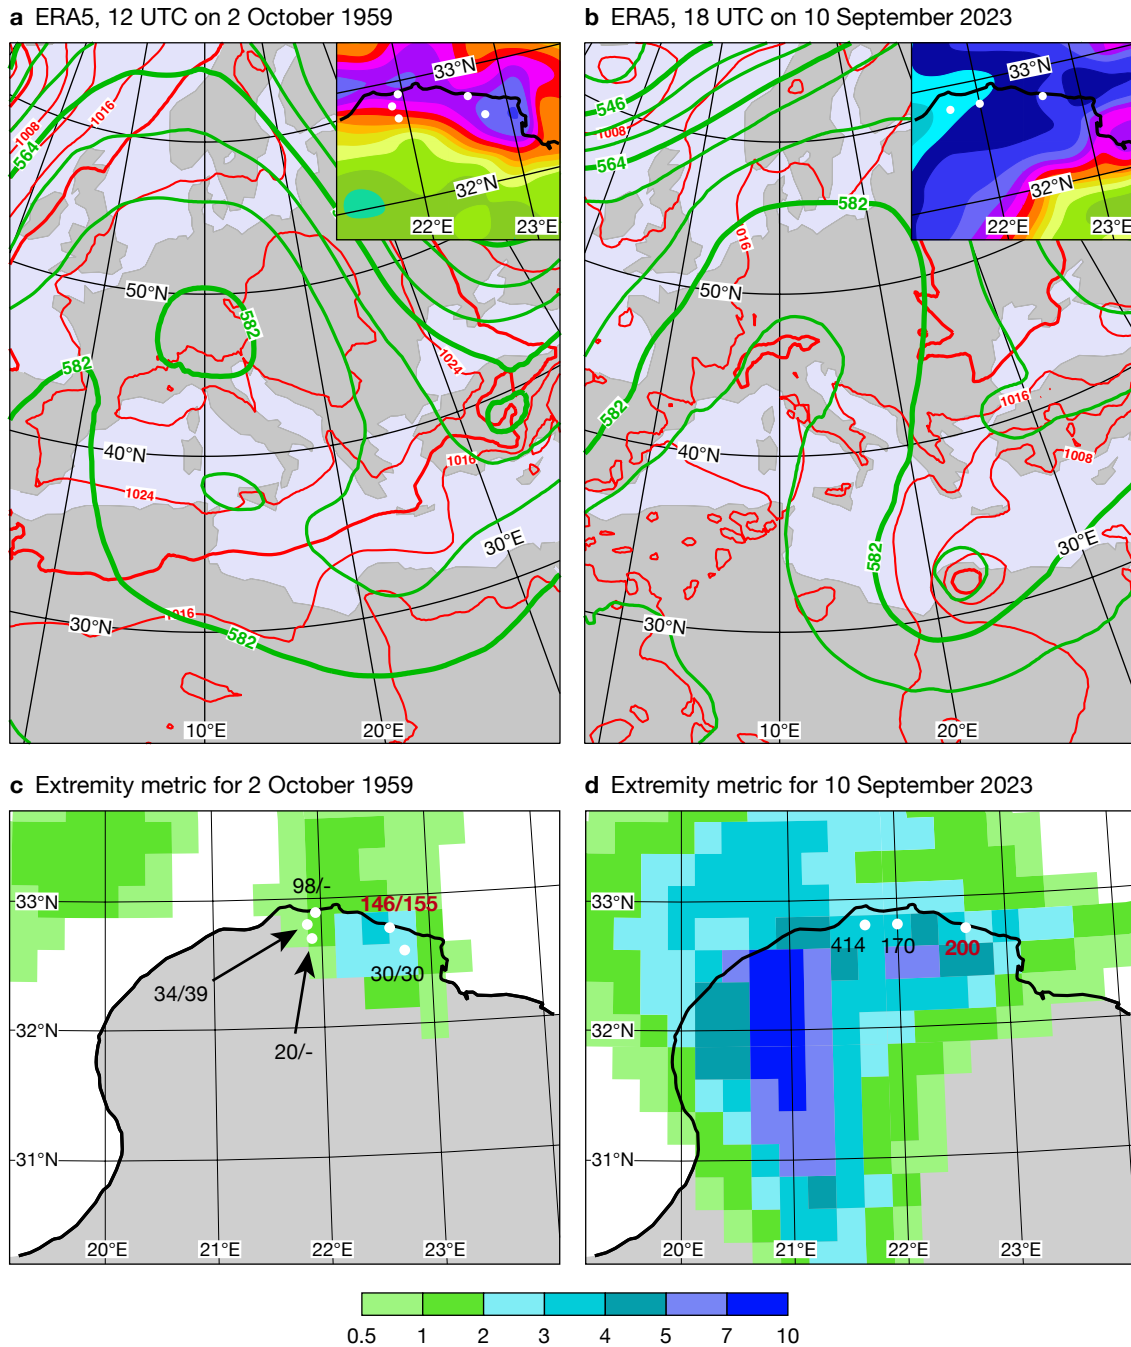


FIGURE 10 Reference points for two extreme rainfall events in ECMWF's ERA5 reanalysis. The top two figures show 500 hPa height (green, 6 dm interval) and mean sea-level pressure (red, 4 hPa interval) at (a) 12 UTC on 2 October 1959, and (b) 18 UTC on 10 September 2023. Insets show the 99th percentile of 24 h rainfall from ERA5-ecPoint for those days over northern Libya (legend as in Figure 11). The bottom two figures show an extremity metric (shaded) for 24 h grid box rainfall on those two dates, computed, for each grid box separately, as raw ERA5 rainfall for that day divided by the median of the 30 annual maximum 24 h rainfall values (1991–2020, 00–24 UTC). Dots/numbers in (c) and (d) show gauge rainfall (mm) in 24 h ending at 06 UTC on (c) 2/3 October 1959, (d) 11 September 2023; red numbers are for Derna. The small Wadi Derna catchment intersects the three ERA5 grid boxes due west of the Derna observation point (reference Figure 1d). ERA5-ecPoint data was created within the HIGHLANDER project, funded by the Connecting Europe Facility (CEF) Telecommunications sector under agreement number INEA/CEF/ICT/A2018/1815462.

release, overtopping became very likely. Indeed, on-the-ground data has indicated that substantial overtopping of the Bu Mansur dam did occur, causing disintegration and collapse, with the resulting wave of water then overwhelming the much lower capacity El Bilad dam downstream, around 01 UTC on 11 September, and going

on to cause catastrophic impacts in Derna itself. Another important factor may have been the short duration; rain in Derna commenced at about 17 UTC, implying that 200 mm accumulated there in under 13 h.

Separate calculations (Ashoor & Eladawy, 2024) suggest that the lower El Bilad dam *could have* coped with

discharge from its own much smaller catchment, via spillway release, had the upper dam not collapsed. This reiterates the fact that dam collapse acts to dramatically and dangerously amplify the nominal peak discharge well beyond what would have been observed had no dams been built.

Just as important as creating predictions is the communication of warnings to those that would manage any mitigating action (Golding, 2022). Very sadly, the needed actions, such as evacuation, did not take place. This was for multiple reasons. Golding refers to communication chain inadequacies as the “valleys of death”, a term which has a particular poignancy here.

Evaluating event rarity

We now outline some ways to contextualise actual (extreme) events. One can examine:

1. Rainfall observation records
2. Synthesised rainfall (based on free-running model simulations)
3. Past extreme rainfall events
4. Gridscale rainfall in reanalyses
5. Point-scale rainfall inferred from reanalyses
6. Simulations of the future.

Option 1 was not used here; in any case discontinuities in observation records make this challenging. Options 2 and 6 are discussed together in Box A, focussing on synthesised medicane-related rainfall, based on climate model runs. Whilst skilful simulations of the future become vital if risk is changing sufficiently fast to diminish the relevance of real-event analysis, Box A shows that multi-model uncertainty is still much greater

100-year return period 24 h rainfall from ERA5-ecPoint 2000–2020

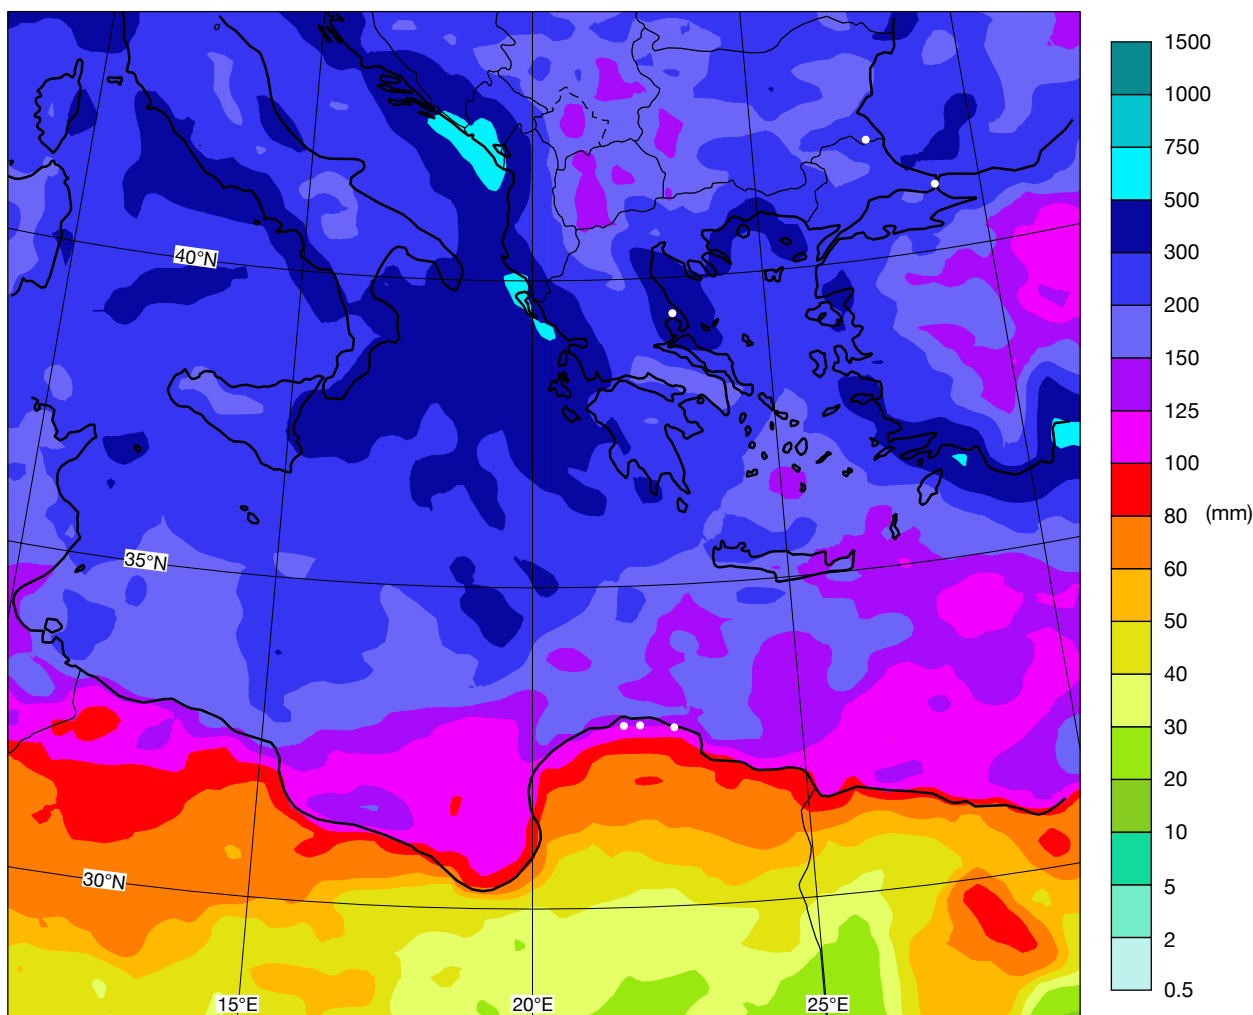


FIGURE 11 An ERA5-ecPoint-based reference metric for extreme 24 h rainfall (mm). This is the 99.998th percentile, which equates to approximately a 100-year return period. It is derived from the 99 percentiles created each day for point rainfall realisations within each ERA5 grid box, over the 21 years 2000–2020: for each grid box, it roughly denotes the realisation ranked 15th out of 760,000. The white dots indicate the places given in Figure 1a plus the additional Libya measuring station shown in Figure 10d. ERA5-ecPoint data was created within the HIGHLANDER project (as on Figure 10).

than the change in the multi-model mean, and that, even within this framework, Daniel was very much an outlier. All this then suggests a need for more work, because in the end climate change could potentially produce large positive or negative changes in rainfall. Further work should also encompass medicanes that are driven, in part, by baroclinic forcing (like Daniel – see the upper-level feature U in Figure 3c).

For option 3, we cross-reference an extreme flooding event in the Derna area in October 1959. Event rainfall values in mm are shown in Figure 10c (1/2 October), with Derna in red. In Derna itself 301 mm fell in 1959, versus 200 mm in 2023, although there was seemingly more spatial variability in 1959. Derna flooding in 1959 from the wadi peaked on the night of 2 September, and whilst there were considerable human and material losses, these were at a much lower level than in 2023. This was at least partly because there were then no dams, and (self-evidently) no dams collapsed. Figures 10a,b compare the synoptic situations for the 1959 and 2023 events using ERA5 data. The upper-level patterns (green) were strikingly similar, with a central Europe ridge/block, and an extended trough downstream reaching Libya. The surface pattern (red) is, however, very different, with no medicane in 1959. Whilst this could have been an early era reanalysis weakness, two other reanalyses do support the ERA5 picture, so maybe it is not. The absence of a medicane but presence of an upper trough, plus the time of year, suggest that deep, moist convective cells were responsible in 1959. At the same time, contour patterns imply high vertical (geostrophic) wind shear, which would have been conducive to convective organisation and longevity, whilst slack flow in the trough axis could also promote slow cell movement, and accordingly higher localised extremes. So, whilst the broadscale upper-level driver had a similar look in 1959, the surface weather picture was likely very different to 2023, with convective storms delivering extreme rainfall locally. For small catchments, this scenario can be more dangerous than widespread rainfall.

For option 4, we compare the grid-scale rainfall in ERA5 with a longer-term ERA5-based reference point, namely the median annual maximum over a (recent) 30-year period. The ratio of the two is shaded in Figures 10c,d. In comparing in this way, without showing actual values, we tried to cancel out model-resolution-related issues (which are discussed, with examples, in Hewson, 2024). The truly widespread nature of the 2023 event is clear, but in the immediate vicinity of Derna (the two adjoining grid boxes, say) this ratio was much the same in 1959. Whilst such local details are not ordinarily that reliable, there is nonetheless a reasonable correlation between the ratios and observations we do have for the two cases. This assessment supports conclusions from option 3: much more localised rainfall in 1959, but possibly as extreme as in 2023 on the Wadi-Derna catchment scales.

For option 5, we use ERA5-ecPoint output, a downscaled version of ERA5 that aims to deliver a probabilistic representation of rainfall at point scale (akin to gauge measurements) within each ERA5 grid box (see Botazzi et al., 2024). Insets in Figure 10a,b show 99th percentiles for point rainfall so derived, for 2 October 1959 and 10 September 2023 respectively. Whilst values in the latter case are clearly higher, in both instances observed values seem to be in range. Indeed, in the 1959 case, whilst 155 mm in Derna would rank as very unlikely, it was still rated a plausible outcome from the synoptic pattern parameters encapsulated in the ecPoint post-processing (~1% probability or slightly less). Finally, Figure 11 provides, from 21 recent years of post-processed ERA5-ecPoint data, a 100-year return period (RP) estimate for 24 h gauge-based rainfall. Whilst other datasets could admittedly give different RP estimates, an advantage of ecPoint is that it allows n-year RPs to be estimated from m years of data, where $n \gg m$, without using extreme value theory/tail parametrization. At the three marked sites in Libya, which provided data for the Daniel case, the value is about 125 mm, suggesting that observed rainfall for Daniel, and the Derna value for the 1959 case (Figures 10c,d), were indeed very extreme (RP > 100 years), particularly at Al Bayda in 2023 (RP \gg 100 years). Interestingly, the 100-year return period value for Al Bayda is almost identical to the values given by climate model downscaling in the Figure in Box A. However, the presence of strong gradients across northern Libya in Figure 11 also suggests that intrinsic uncertainty in these values is locally elevated. On top of this, we know that topographically-enhanced convection can be associated with systematic model rainfall biases (Hewson & Pilloso, 2021). Such factors should be taken into account should any such data ever be used for design specification.

Finally, referring back to the European rainfall extremes discussed earlier (and reiterating that Figure 11 is not medicane-specific), we can also infer that, with this dataset, the referenced 24 h rainfall values recorded at Portaria, Strandja and Istanbul had RPs, respectively, of \gg 100 years, ~100 years and \ll 100 years.

Summary

Key findings of this study are as follows:

- The Daniel cyclone/medicane lasted for eight days, meandering over the Mediterranean for six days before crossing northern Libya and then decaying over Egypt.
- The later track behaviour of Daniel was reassessed using multiple sources, including METAR observations and dust in imagery; the lowest central pressures (996 hPa) occurred over Libya.

- The strange intensification over land was attributed to forcing from anomalies at very high tropospheric levels. Earlier behaviour followed more conventional pathways.
- Model reruns suggest that release of heat from sand/soil over Libya did not aid deepening.
- Extreme rainfall/flooding over Bulgaria, Türkiye and Greece occurred whilst Daniel was developing as a discrete cyclone; flood devastation over Libya was much more directly related to Daniel's passage.
- High moisture content, instability, high SSTs, topographic uplift and flow persistence all contributed to the rainfall/flooding events in Europe; in Libya, mechanisms were similar but persistence was less and dynamically forced uplift stronger.
- Medium-range forecasts for Europe were compromised by a predictability barrier related to upstream interactions between Hurricane Franklin and the mid-latitude waveguide. Later ensemble forecasts for Libya were able to raise alarms seven days in advance.
- ECMWF and DestinE runs at 9 km and 4.4 km resolution captured the European rainfall well at short leads; although both exhibited systematic and random errors, the DestinE 4.4 km runs were slightly better.
- Two-day lead-time discharge predictions for Wadi Derna from GloFAS exhibited extreme uncertainty, from a dry valley to extraordinarily high levels, reflecting different spatial rainfall patterns and the small catchment size.
- The actual (inferred) discharge was comparable with the wettest two-day-lead GloFAS outcome.
- As runoff was about three times greater than dam capacity for Wadi Derna, dam collapse was almost inevitable; the resulting wave of water dramatically exacerbated impacts.
- Another devastating flooding event in October 1959 was also examined via reanalysis data – whilst the upper-level pattern and rainfall near Derna were similar to the 2023 case, rainfall was much more localised.
- Different ways of assessing event rarity for 24 h Libya rainfall gave relatively consistent answers; nominal event return periods were > 100 years in 1959 and >> 100 years in 2023.

Further reading

Ashoor, A. & A. Eladawy, 2024: Watch and Upgrade or Deconstruct and Relocate: Derna Catastrophe Lessons Amid the Climate-change Era of Unpredictable Flash Floods, PREPRINT (Version 1) available at *Research Square*. <https://doi.org/10.21203/rs.3.rs-3858769/v1>

Bister, M. & K.A. Emanuel, 2002: Low frequency variability of tropical cyclone potential intensity 1. Interannual to interdecadal variability. *Journal of Geophysical Research: Atmospheres*, **107**. <https://doi.org/10.1029/2001JD000776>

Bottazzi, M., L. Rodríguez-Muñoz, B. Chiavarini, C. Caroli, G. Trotta, C. Dellacasa et al., 2024: High performance computing to support land, climate, and user-oriented services: The HIGHLANDER Data Portal. *Meteorological Applications*, **31**(2), e2166. <https://doi.org/10.1002/met.2166>

Cusack, E., L. Paterson, W.J. Lang & C. Csekits, 2017: WGCEF Task Team on Storm Naming in Europe. *The European Forecaster*. Vol. **22**. P. 48-50. (euroforecaster.org/newsletter22/full.pdf)

Emanuel, K., 2018: 100 Years of Progress in Tropical Cyclone Research. *Meteorological Monographs*, **59**. <https://doi.org/10.1175/amsmonographs-d-18-0016.1>

Feldmann, M., K. Emanuel, L. Zhu & U. Lohmann, 2019: Estimation of Atlantic tropical cyclone rainfall frequency in the United States. *Journal of Applied Meteorology and*

Climatology, **58**(8), 1853–1866. <https://doi.org/10.1175/JAMC-D-19-0011.1>

Golding, B. (ed.), 2022: Towards the “Perfect” weather warning: bridging disciplinary gaps through partnership and communication. *Springer Cham*. <https://doi.org/10.1007/978-3-030-98989-7>

Hewson, T.D., 2024: Capturing Extreme Rainfall events. *ECMWF Newsletter* **178**, 2-3.

Hewson, T.D., & F.M. Pilloso, 2021: A low-cost post-processing technique improves weather forecasts around the world. *Communications Earth & Environment*, **2**, 132. <https://doi.org/10.1038/s43247-021-00185-9>

NRCS, 2008: *National engineering handbook: Part 630—hydrology*. USDA Soil Conservation Service: Washington, DC, USA.

Romero, R. & K. Emanuel, 2017: Climate Change and Hurricane-Like Extratropical Cyclones: Projections for North Atlantic Polar Lows and Medicanes Based on CMIP5 Models. *Journal of Climate*, **30**(1), 279–299. <https://doi.org/10.1175/JCLI-D-16-0255.1>

Vannièrè, B., I. Sandu, P. Düben, M. Maier-Gerber, J. Schrötle & J. Denissen, 2024: A daily forecast with the prototype global Extremes Digital Twin of Destination Earth. *ECMWF Newsletter No. 178*, 10–11. <https://doi.org/10.21957/1a8466ec2f>

Introducing earthkit

Iain Russell, Tiago Quintino, Baudouin Raoult, Sándor Kertész, Pedro Maciel, James Varndell, Corentin Carton de Wiart, Edward Comyn-Platt, Olivier Iffrig, James Hawkes, Simon Smart

Earthkit is an exciting new open-source Python development led by ECMWF, providing powerful tools for speeding up weather and climate science workflows by simplifying data access, processing, analysis, visualisation, and much more. Earthkit is at the start of a multi-year development, but many aspects are already open for testing and feedback.

Earthkit in context

In 2023, ECMWF published its Software Strategy and Roadmap for 2023–2027 (Quintino et al., 2023), outlining the plan for software development over the next five years. Some of the core principles which underpin the software strategy are:

- improving reusability and componentisation of software
- further use of external software
- adoption of open development
- ensuring data scalability to new higher resolutions
- modernisation to new standards for higher interoperability.

Earthkit helps us realise this vision by offering multiple interoperable software components built on top of well-established open-source Python libraries, such as NumPy, pandas and Matplotlib. It also integrates and leverages the robust and operations-ready software stack that is familiar to ECMWF production systems (e.g. ecCodes, the Fields DataBase (FDB), etc.).

ECMWF has decades of experience in developing software that handles high data volumes with timeliness as an imperative. With this in mind, we have designed earthkit with certain principles, such as the ability to perform an entire workflow without reading any data from or writing any data to storage. We are also responding to the trend towards GPU architectures in high-performance computing facilities (HPCF), by ensuring that processing pipelines are GPU-ready.

The breadth of scope of earthkit is such that it is intended to eventually replace Metview, ECMWF's meteorological workstation software, as ECMWF's all-in-one environment for retrieving, inspecting,

processing and visualising meteorological data. This is no small task, and we acknowledge that it will take several years, considering that Metview has been developed continually over the past three decades (for a look back at the beginning, see: Daabeck et al., 1995). Another major software package that will eventually be replaced by earthkit will be Magics, ECMWF's plotting package with a 35-year history of development (Lamy-Thépaut et al., 2023). Its functionality will be covered by the components earthkit-maps and earthkit-plots. Developments of Metview and Magics have already slowed in order to put more focus on earthkit, and user migration paths will be planned as more earthkit functionality becomes available.

Earthkit architecture

Earthkit consists of several Python components, each specialising in a certain area, such as data access, regridding or plotting. Each of these may depend on several Python and binary libraries developed at ECMWF, which, in turn, may depend on established third-party Python and binary libraries. One of the driving forces behind the architecture is reusability and componentisation of code – each package should do one job and do it well.

In addition to end users directly working with earthkit, ECMWF has many applications and frameworks that will be migrated to use earthkit over time. These include ECMWF's post-processing framework PProc, its verification software Quaver, its graphical forecast web application ecCharts, its hydrological post-processing framework Danu, and the Common Data Store (CDS) Engine. This is all summarised in Figure 1.

The following sections will present some detail about selected earthkit components.

Earthkit-data

The purpose of earthkit-data is twofold. Firstly, it provides an interface for accessing data, which includes loading files from disk, from URLs, directly from ECMWF's FDB, and from remote services such as ECMWF's Meteorological Archival and Retrieval System (MARS). Secondly, it represents data in a format-agnostic way and offers data extraction, inspection, slicing, filtering and conversion. By doing so, it provides a data layer making all the earthkit components interoperable. All is provided via a high-level interface,




FIGURE 1 High-level overview of the software components that make up earthkit. The dependency chain includes many well-established third-party libraries covering, for example, data representation and processing (NumPy, pandas and Xarray), plotting (Matplotlib) and compression (libaec).

where details of data transfer, format and storage are abstracted and transparent to the user.

To simplify data access, users can invoke it via a single method `from_source()`. There are various built-in sources available, for example 'file', 'url', 'fdb' or 'mars', supporting a wide range of data formats including GRIB, NetCDF, BUFR, ODB and many more. Some sources, such as URLs and FDB, can be read as a stream straight into memory, allowing for highly effective workflows not involving any storage I/O. Additionally, users can define their own sources as plugins, which can then be published and shared with others without the need to update and release earthkit-data itself.

Earthkit-data also works comfortably with tar.gz files, returning a combination of their unpackaged contents.

The code in Figure 2 shows how to load such data and display the contents of the packaged GRIB files, and the result is also shown in Figure 2.

An additional benefit for users is that data retrieved from remote sources can be cached to disk, providing immediate access to it when the same retrieval is run again. The cache is highly configurable and can be switched off as per the typical requirements for operational environments, where disk access should be minimised and explicit.

All data objects returned by `from_source()` allow a few common methods, including the conversion of the underlying data into familiar scientific Python formats via `to_numpy()`, `to_pandas()` and `to_xarray()`. GRIB data and suitable NetCDF files are represented as

```
import earthkit.data

path = "https://get.ecmwf.int/repository/test-data/earthkit-
data/examples/test_gribs.tar.gz"

ds = earthkit.data.from_source("url", path)

ds.ls()
```

	centre	shortName	typeOfLevel	level	dataDate	dataTime	stepRange	dataType	number	gridType
0	ecmf	2t	surface	0	20200513	1200	0	an	0	regular_ll
1	ecmf	msl	surface	0	20200513	1200	0	an	0	regular_ll
2	ecmf	t	isobaricInhPa	500	20070101	1200	0	an	0	regular_ll
3	ecmf	z	isobaricInhPa	500	20070101	1200	0	an	0	regular_ll
4	ecmf	t	isobaricInhPa	850	20070101	1200	0	an	0	regular_ll
5	ecmf	z	isobaricInhPa	850	20070101	1200	0	an	0	regular_ll

FIGURE 2 The `ls()` function displays columns of per-field information; extra keys can be specified by the user.

a FieldList (i.e. a list of fields: horizontal slices of the atmosphere at a given time or period), offering a more specialised and richer functionality.

By design, memory usage in earthkit-data is kept at a minimum by implementing lazy loading whenever possible. When iterating over a FieldList created from GRIB, only one field is kept in memory at a time by default. This technique is vital when larger-than-memory datasets must be processed, and it can deal with an arbitrary number of GRIB fields. We can also represent a field as a combination of a metadata object and a NumPy array storing the values. With this we can build whole FieldLists decoupled from the original data format. This representation is particularly useful for computations, and support for other array backends, such as the GPU-ready PyTorch, is planned for the near future.

For examples, see the earthkit-data documentation, including its user guide and API reference at <https://earthkit-data.readthedocs.io/en/latest/>.

Earthkit-maps

Earthkit-maps is the geospatial visualisation component of earthkit, empowering users to produce publication-quality weather and climate maps in just a few lines of code. It is a convenient tool for taking a quick look at any geospatial data from a wide range of formats, whilst also offering the depth to fully customise every aspect of a chart for a publication or a presentation.

Earthkit-maps is built on the shoulders of giants, taking some of the features and philosophies of ECMWF's Magics visualisation software and applying them to a backend built on top of the feature-rich, open-source visualisation libraries Matplotlib and cartopy. It builds upon this strong foundation with domain-specific knowledge to offer a super simple API, which removes much of the code one would normally have to write when using either Magics or Matplotlib.

Earthkit-maps also leverages the power of earthkit-data and earthkit-regrid to greatly simplify the process of taking data from its source through to a visualisation. Users can easily write input-agnostic code that will work just the same with GRIB data as it does for NetCDF. They can also visualise data on an octahedral reduced Gaussian grid in exactly the same way as for a regular latitude-longitude grid.

Finally, earthkit-maps provides some extra convenience features, enabling:

- cropping a domain to fit a named country, region or continent, optimising (optionally) the map projection for the given region of the world
- automatic use of a suitable style based on input metadata

- creation and use of style libraries based on an organisation, project, or personal choice
- formatting of titles and labels using templates which read directly from the data's metadata
- built-in support for a wide range of ancillary layers from the Natural Earth cartography library, such as cities, urban areas, rivers, lakes and more.

The quickest way to get started using earthkit-maps is with `earthkit.maps.quickplot()`, which will attempt to visualise data the best it can based on its metadata. Figure 3 shows an example plot of 2-metre temperature and mean sea-level pressure.

We have a global field here, but what if we want to zoom in on Europe? All we need to do is include a `domain="Europe"` argument to obtain the plot shown in Figure 4.

For many more examples, see the earthkit-maps documentation, including its user guide and gallery at <https://earthkit-maps.readthedocs.io/en/latest/>.

Earthkit-regrid

Earthkit-regrid is a component that specialises in mapping one arrangement of geographical points onto another, for example transforming an octahedral reduced Gaussian grid onto a regular latitude-longitude grid at a specific spatial resolution.

Earthkit-regrid leverages the principle that interpolation is fundamentally a geometrical relationship between input and output grid resolutions and formats. This relationship is modelled as a linear operation $\mathbf{y} = \mathbf{Ax}$, where \mathbf{x} and \mathbf{y} represent the input and output fields respectively, and \mathbf{A} denotes the interpolation operator as a matrix that maps the input to the output. This matrix is generally sparse (lightweight) and multiplication can be efficiently implemented by a variety of optimized backends, including both CPU- and GPU-based architectures. The advantage of separating the matrix computation from the interpolation process is that it allows for a simpler interpolation package that can be deployed anywhere with minimal Python-only libraries.

The concept of calculating interpolation weights, caching, and performing multiplication was pioneered with the Meteorological Interpolation and Regridding (MIR) software package (Maciel et al., 2017). This is the post-processing engine of ECMWF's archive and product generation services (MARS/pgen), and it is key to the services' scalability and efficiency. Earthkit-regrid utilises pre-calculated MIR interpolation weights for a large selection of input/output grid combinations and a limited number of interpolation methods, which are fetched from a designated online source and cached upon the first request. MIR supports many advanced

```
import earthkit.data

import earthkit.maps
source_filename = "era5-2t-mslp-christmas-2010.grib"
earthkit.data.download_example_file(source_filename)
data = earthkit.data.from_source("file", source_filename)
earthkit.maps.quickplot(data)
```

2 metre temperature and Mean sea level pressure at 12:00 on 2010-12-25

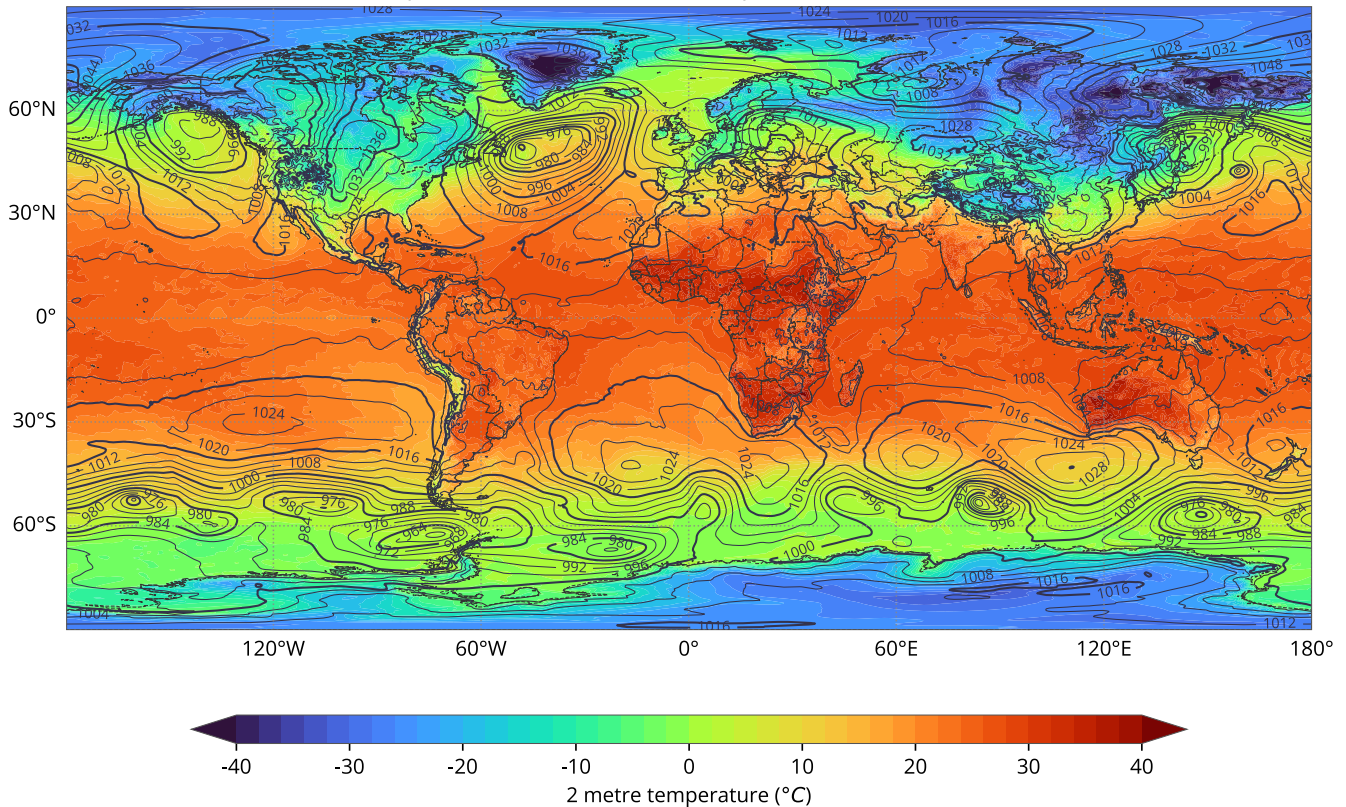


FIGURE 3 The quickplot() function of earthkit-maps is able to automatically assign styles and combine layers.

interpolation methods (e.g. conservative interpolation or methods that support climate files of ECMWF’s Integrated Forecasting System), is extensively configurable, and has followed the ever-increasing requirements of high-resolution model runs. By sharing interpolation weights, MIR’s functionality can be made increasingly accessible to earthkit-regrid over time.

The strategy of separating the generation of interpolation weights from performing interpolations through a high-level Python package is highly advantageous for both users and ECMWF as a data provider. Future iterations of earthkit-regrid will be able to directly call the MIR engine in order to facilitate regriddings that have not been precomputed. There will thus be facilitated support for new input/output grid combinations and more sophisticated interpolation methods, without necessitating changes in the user environment. The synergy between earthkit and MIR is a testament to the transformative potential of collaborative software development.

For more information and examples, see the earthkit-regrid documentation at <https://earthkit-regrid.readthedocs.io/en/latest/>.

Earthkit-meteo

The earthkit-meteo component is in its design phase and will contain a series of meteorological and mathematical functions that are widely used in our field. It aims to offer clear APIs coupled with extensive testing. This is to ensure their safe application across various contexts, such as operational forecast post-processing, integration into user applications, and usage within Jupyter notebooks, among others.

The earthkit-meteo component will be organised into several sub-modules, each dedicated to a specific post-processing category:

- **thermo**: contains thermodynamic functions, including computation of humidity, temperature conversions, etc.

```
earthkit.maps.quickplot(data, domain="Europe")
```

2 metre temperature and Mean sea level pressure at 12:00 on 2010-12-25

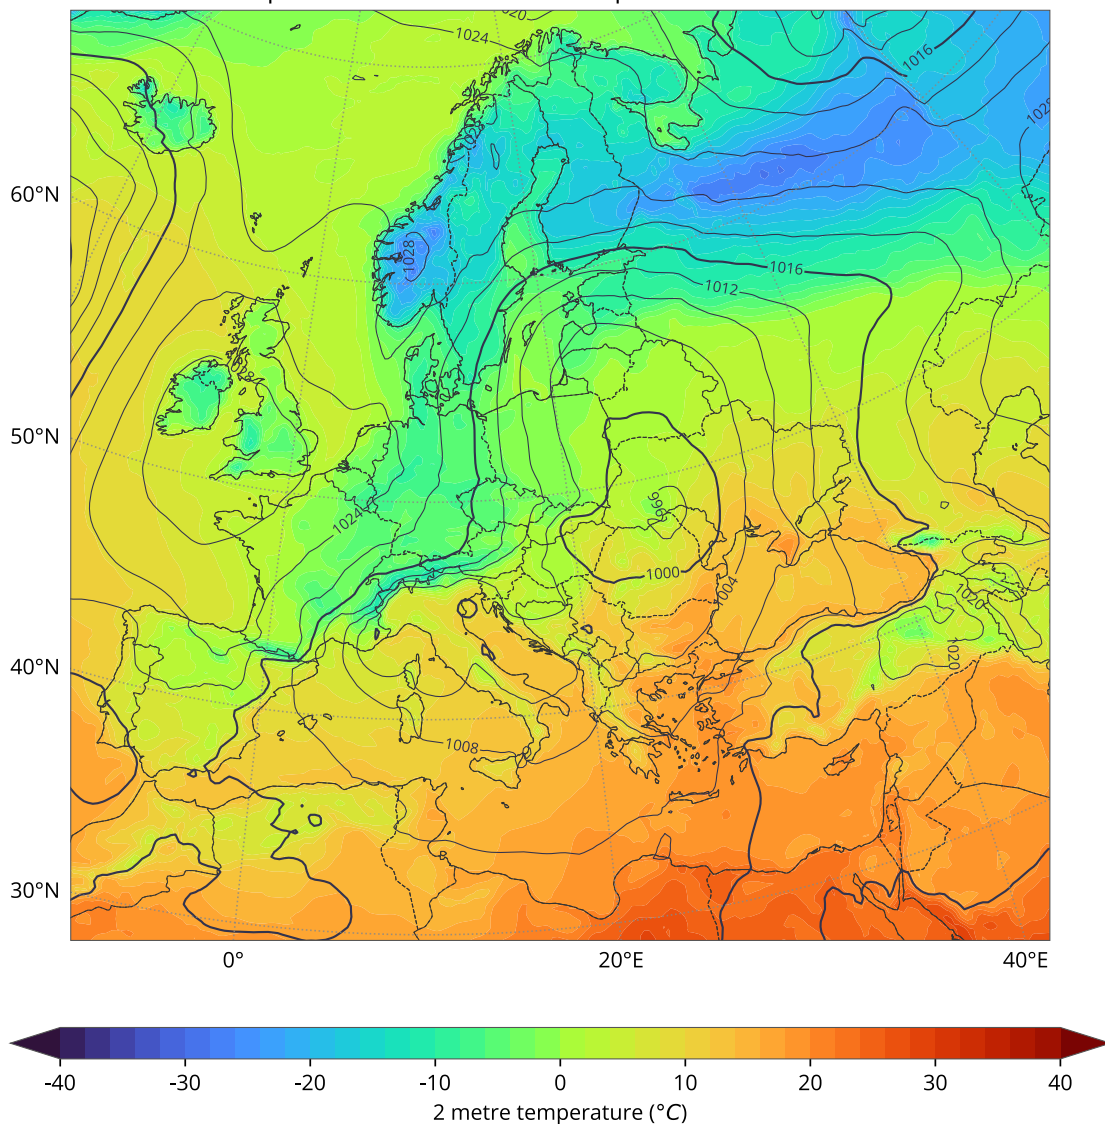


FIGURE 4 The result of specifying one of the pre-defined domains as an argument to the quickplot() function.

- **wind:** calculates wind-related quantities, such as wind directions, and conversions between polar coordinates and Cartesian coordinates
- **stats:** offers essential mathematical and statistical functions, including quantiles and extensions of NumPy functions
- **extreme:** includes Extreme Forecast Index (EFI) and Shift of Tails (SOT)
- **score:** provides metrics such as the Continuous Ranked Probability Score (CRPS)
- ... and many more coming!
- **A NumPy low-level Python API.** This interface directly processes NumPy arrays, facilitating efficient post-processing without the complexities of data management.
- **An earthkit high-level Python API.** This interface maintains metadata throughout the processing, ensuring outputs are valid earthkit-data objects with appropriate metadata, constructed from the inputs.

Earthkit-meteo will undergo some consolidation before its initial release, which will initially introduce the NumPy-level API.

Collaboration

All source code for earthkit is openly available on GitHub, and we have avidly embraced open

The earthkit-meteo modules will be available through two main interfaces, each designed for different use cases:

development. With this approach, we mean to develop the codebase publicly. This means making the process transparent and, from the outset, available for interaction with the community, and in particular with ECMWF Member and Co-operating States. We seek early feedback on the design and features of the earthkit component system.

Early collaborations have already been set in motion. For example: (1) MeteoSwiss is seeking to adopt the earthkit components as a basis for its new generation post-processing system; (2) the German Meteorological Service (DWD) has provided a design and requirements for earthkit-data to support accessing data from the Amazon Simple Storage Service (S3); and (3) the Australian Bureau of Meteorology (BoM) is contributing new plot styles to earthkit-maps.

Through these and future collaborations, we believe that earthkit will strengthen our operational workflows as it becomes a building block for applications such as PProc (for post-processing), ecCharts, or new machine learning frameworks.

Conclusion and future

Earthkit is a large part of ECMWF's software strategy published in 2023, and it embodies all its principles. As more services and applications migrate to use earthkit, we will create an ever-more efficient and consistent software stack. The immediate priorities will be to allow basic and scientific computations via earthkit-meteo and the planned earthkit-hydro, and to bring earthkit-regrid to a level that more fully leverages the power of the MIR regridding engine. Developments have a robust future with further earthkit components planned or in early development stages. Earthkit-plots will complement earthkit-maps by providing ways to generate non-map plots, such as time series, meteograms and cross sections. Earthkit-climate will provide methods for calculating climate metrics and risk factors. It is expected to be highly community-driven. This is also because contractors of the EU's Copernicus

Climate Change Service (C3S), which is implemented by ECMWF, will contribute their expertise whilst producing fully traceable project outputs, e.g. derived datasets and web applications.

We envisage a future where earthkit is widely adopted and contributed to, from both inside and outside ECMWF, and where it serves as a platform for stronger interaction with our Member States. Earthkit is still under development, but beta versions of earthkit-data, earthkit-maps, earthkit-plots and earthkit-regrid are available through PyPI. We welcome feedback on these packages through issues raised on GitHub.

A general introduction to earthkit is available here: <https://earthkit.readthedocs.io/en/latest/>, with installation instructions here: <https://earthkit.readthedocs.io/en/latest/install.html>. The GitHub code repositories for earthkit-data, earthkit-maps and earthkit-regrid are available here: <https://github.com/ecmwf/earthkit-data>; <https://github.com/ecmwf/earthkit-maps>; <https://github.com/ecmwf/earthkit-regrid>.

Further reading

Quintino, T., U. Modigliani, F. Pappenberger, S. Lamy-Thépaut, S. Smart, J. Hawkes, et al., 2023: Software Strategy and Roadmap 2023–2027. *ECMWF Technical Memorandum No. 904*.

Daabeck, J., B. Norris & B. Raoult, 1995: Metview – Interactive Access, Manipulation and Visualisation of Meteorological Data on UNIX Workstations. *ECMWF Newsletter No. 68*, 9–28.

Lamy-Thépaut, S., S. Siemen & J. Varndell, 2023: What next for Magics visualisation? *ECMWF Newsletter No. 177*, 22–26.

Maciel P., T. Quintino, U. Modigliani, P. Dando, B. Raoult, W. Deconinck, et al., 2017: The new ECMWF interpolation package MIR. *ECMWF Newsletter No. 152*, 36–39.

ECMWF publications

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

Technical Memoranda

915 **Geer, A., C. Lupu, D. Duncan, N. Bormann & S. English:** SURFEM-ocean microwave surface emissivity evaluated. *February 2024*

ESA or EUMETSAT Contract Reports

Rennie, M. & L. Isaksen: The NWP impact of Aeolus Level-2B winds at ECMWF. *January 2024*

Kirsti, S., P. Weston & P. de Rosnay: Quarter 4 2023: Operations Service Report. *January 2024*

Kirsti, S., P. Weston & P. de Rosnay: Annual SMOS brightness temperature monitoring report 2022/23. *November 2023*

ECMWF Calendar 2024

Apr 23	Webinar: Using Jupyter on the EWC to interact with ECMWF datasets	Jun 17–20	7th C3S General Assembly
Apr 23	Policy Advisory Committee (virtual)	Jun 19–20	Council
Apr 24	Finance Committee (virtual)	Sep 9–12	Workshop on Diagnostics for Global Weather Prediction
May 7–10	ECMWF–ESA Machine Learning Workshop (Italy)	Sep 30–Oct 3	Training course: Use and interpretation of ECMWF products
May 13–16	OpenIFS User Meeting	Oct 7–9	Scientific Advisory Committee
May 13–16	Online training course: Atos HPC training	Oct 10–11	Technical Advisory Committee
Jun 4	Online workshop: Using ECMWF Web Map Services (WMS)	Oct 21–22	Finance Committee
Jun 5–6	Using ECMWF's Forecasts (UEF2024)	Oct 22	Policy Advisory Committee
Jun 10–13	8th CAMS General Assembly	Nov 4–8	Training course: Predictability and ensemble forecast systems
Jun 11–13	8th EUMETSAT ROM SAF workshop on GNSS radio occultation measurements	Nov 11–15	Training course: Numerical methods for weather prediction
		Dec 10–11	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 firstname.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. 179 | Spring 2024

European Centre for Medium-Range Weather Forecasts

www.ecmwf.int