

Newsletter

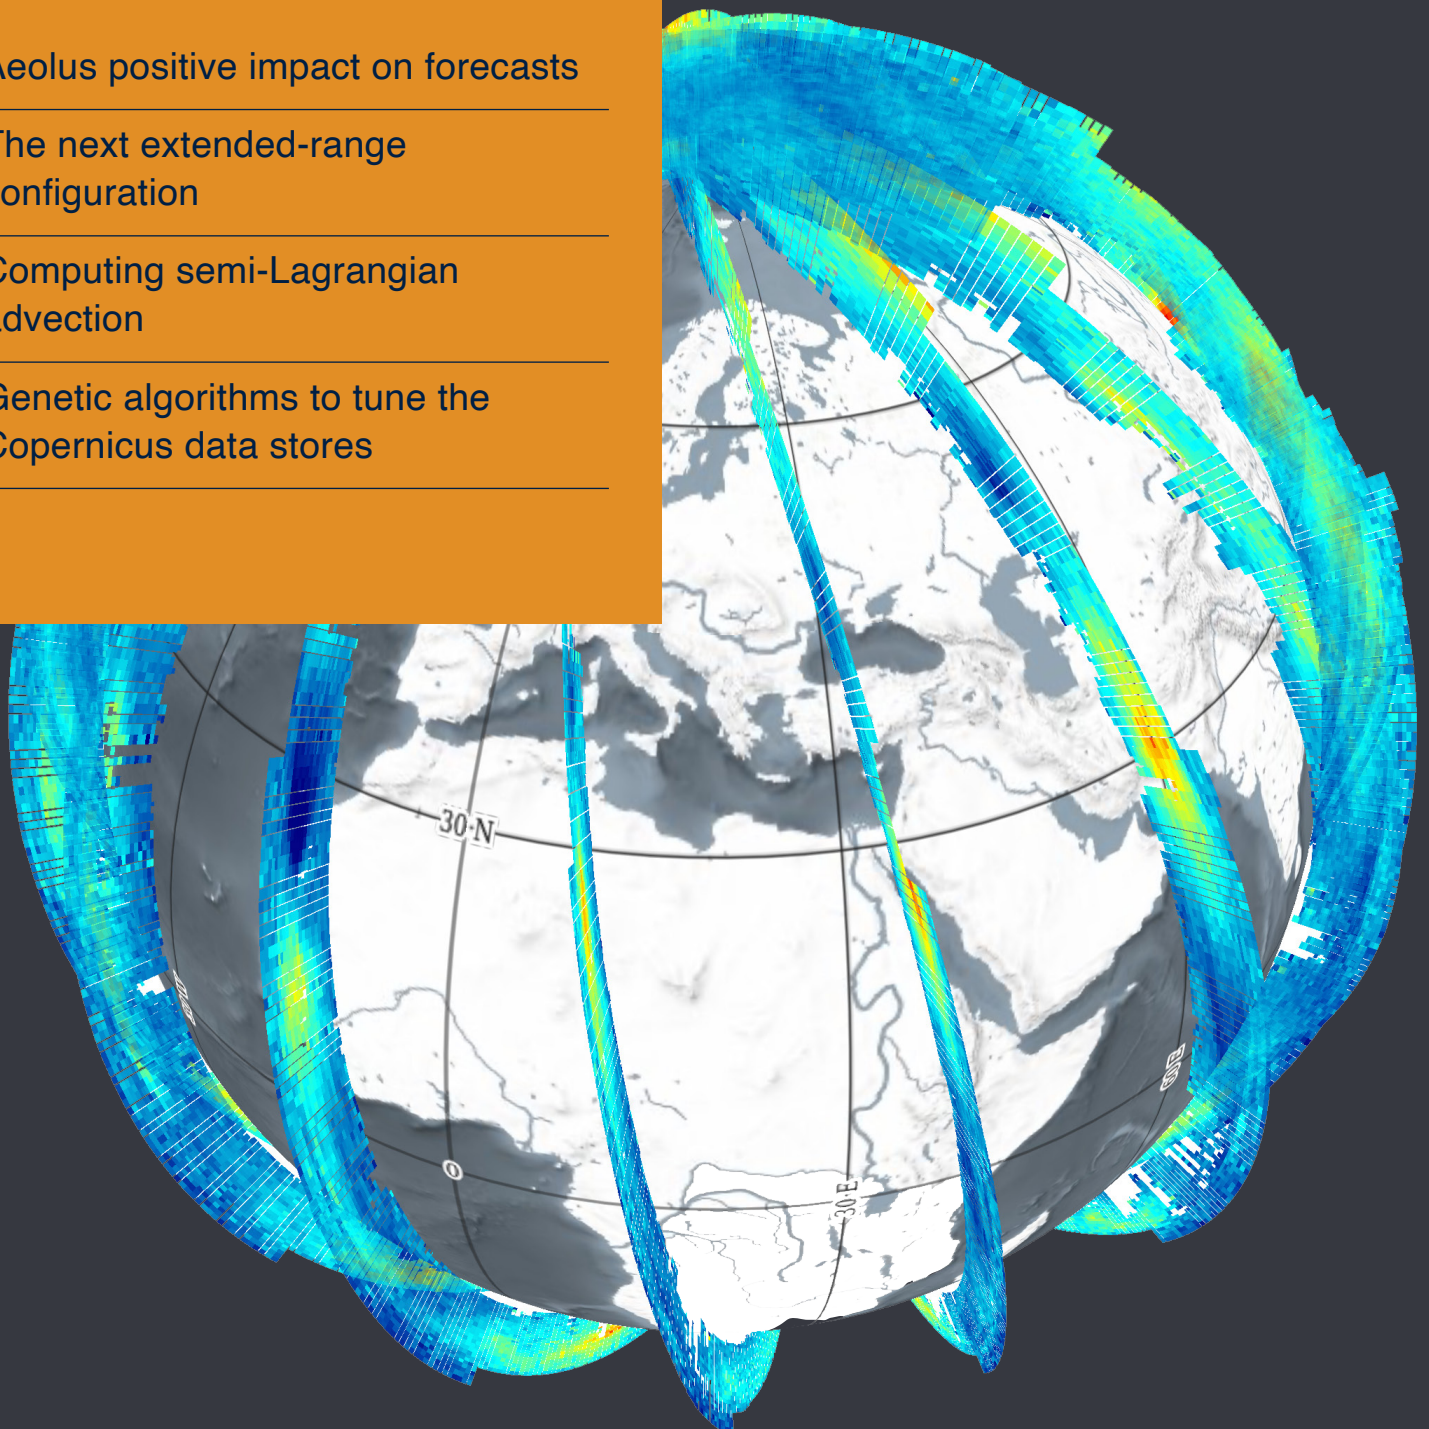
No. 173 | Autumn 2022

Aeolus positive impact on forecasts

The next extended-range
configuration

Computing semi-Lagrangian
advection

Genetic algorithms to tune the
Copernicus data stores



© Copyright 2022

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The ECMWF Newsletter is not peer-reviewed.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

New satellites

Numerical weather prediction crucially depends on the Global Observing System, and in particular on satellites providing timely weather-related information across the globe. That is why ECMWF works closely with the European Organization for the Exploitation of Meteorological Satellites (EUMETSAT) and the European Space Agency (ESA) to ensure that satellites will provide highly relevant information for weather prediction. In December, EUMETSAT will launch the first new geostationary satellite from the Meteosat Third Generation (MTG) series. This imager satellite, MTG-I1, will go into orbit at 36,000 km. It will provide more details about current weather conditions than Meteosat Second Generation (MSG) satellites, which will gradually be replaced by MTG. The MTG constellation comprises two imaging satellites and one sounding satellite. They complement a series of EUMETSAT polar-orbiting satellites, which circle the globe at an altitude of 817 km. A new series of such satellites, called EPS Second Generation, will be launched from the mid-2020s onwards with new and enhanced instruments.

While these satellites will provide much coverage, ECMWF uses data from a host of other satellites. A particularly interesting one was launched by ESA in 2018. The ground-breaking Aeolus satellite provides Doppler wind lidar observations. ECMWF started to assimilate observations from this pioneering ESA Earth Explorer Mission in January 2020. The aim of the mission was to demonstrate new spaced-based Doppler wind lidar capabilities. As the Aeolus article in this Newsletter shows, this aim has been fulfilled: statistically significant and good-magnitude positive

impact on a number of forecasting parameters has been demonstrated. An operational follow-on mission is now being considered by ESA and EUMETSAT.



These developments come at a time when ECMWF switches its data centre operations from Reading (UK) to a new data centre in Bologna (Italy), which houses a new high-performance computing facility (HPCF). The ECMWF numerical weather prediction suite went into operations on the Atos HPCF on 18 October 2022. It will enable us to halve the grid spacing of our ensemble forecasts from 18 km to 9 km in an operational forecast upgrade planned for next year.

Other developments to be introduced in next year's upgrade of the Integrated Forecasting System (IFS) are detailed in this Newsletter. They include a new way of computing semi-Lagrangian advection, and a completely new configuration for extended-range forecasts. The latter change also benefits from the new HPCF. It makes the grid spacing across the full time range of extended-range forecasts the same, it increases their frequency, and it includes twice as many ensemble members as before. The new configuration is supported by users, and that must be the ultimate test for any changes in suites and products we introduce at ECMWF.

Florence Rabier
Director-General

Contents

Editorial

New satellites 1

News

Predicting the heat over western Europe in the summer of 2022 2

ECMWF launches Massive Open Online Course on Machine Learning in Weather and Climate 4

ECMWF contributes to APEC Climate Center's BSISO prediction system 5

Developing an agriculture application for Africa 6

Fire can now be found on MARS 8

Progress on emulating the radiation scheme via machine learning 9

ECMWF real-time data benefit the wider research community 10

Seasonal forecasts of extreme rain in Pakistan 12

New observations since mid-June 2022 13

Earth system science

Aeolus positive impact on forecasts with the second reprocessed dataset 14

The next extended-range configuration for IFS Cycle 48r1 21

A new way of computing semi-Lagrangian advection in the IFS 27

Computing

Using genetic algorithms to tune the Copernicus Climate and Atmosphere Data Stores 32

General

ECMWF publications 35

ECMWF Calendar 2022/23 35

Contact information 35

Predicting the heat over western Europe in the summer of 2022

Linus Magnusson, Ivan Tsonevsky, Estibaliz Gascon, Rebecca Emerton, Emmanuel Rouges

During the summer of 2022, Europe saw several heatwaves and much drier conditions than normal in western Europe. For example, on 19 July the UK for the first time reached temperatures above 40°C. Other parts of the northern hemisphere also saw extreme weather this summer, with severe flooding in Pakistan (see the article by Antje Weisheimer et al. in this Newsletter) and severe droughts in parts of China.

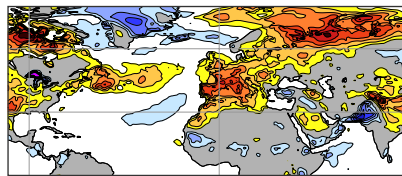
Extended-range forecasts

How well did ECMWF’s extended-range forecasts predict the anomalous weather this summer? To answer this question, the first figure shows composites of weekly temperature and precipitation anomalies for the period 6 June to 28 August from the ERA5 reanalysis, and of all ensemble mean forecasts that verified during that period, with different lead times. For 2-metre temperature, the strong anomaly over western Europe is evident in ERA5 as well as in the composite of week-2 forecasts. In the week-6 composite, the warm anomaly over western Europe is still present, but it is weaker.

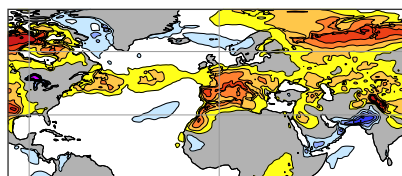
Regarding precipitation (right-hand side of the first figure), the dry anomaly over western Europe was well predicted in week-2 forecasts and seems to be connected to a drier-than-normal band from Florida up to western Europe. This pattern, alongside the anomalously active Indian monsoon and consequent flooding in Pakistan, was captured in the week-6 forecasts.

Looking at a time series of the weekly average temperature over western Europe from ERA5 (second figure, black), three heatwaves are apparent. All lead times shown predicted a warmer-than-normal summer, but it was only the week-1 and week-2 forecasts that captured the intra-seasonal variability in terms of ensemble mean. Week-3 forecasts showed some intra-seasonal

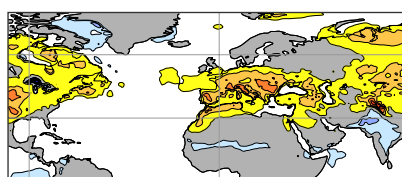
ERA5 temperature anomaly



Ensemble mean, week-2 forecasts

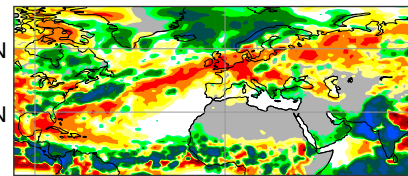


Ensemble mean, week-6 forecasts

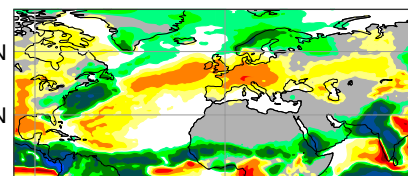


Two-metre temperature anomaly (°C)

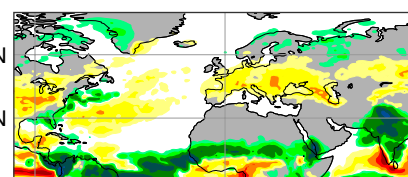
ERA5 precipitation anomaly



Ensemble mean, week-2 forecasts



Ensemble mean, week-6 forecasts



Precipitation anomaly (mm/day)

Two-metre temperature and precipitation anomalies. Composites of weekly 2-metre temperature anomalies (left) and precipitation anomalies (right) for the weeks commencing 6 June to 28 August from the ERA5 reanalysis (top), the ensemble mean from week-2 forecasts (middle), and the ensemble mean from week-6 forecasts (bottom).

variability but with a delayed signal.

Medium-range forecasts for the July heatwave

Focusing on the extreme period when the UK record was set, the heatwave started from the Iberian Peninsula on 17 July and moved north-east to reach Germany and southern Scandinavia on 20 July. The third figure (left) shows the SYNOP stations where the maximum temperature on 19 July exceeded the 99th percentile of the climatology, when the heatwave covered the United Kingdom and large parts of north-western Europe. The spatial extent of the heatwave was well predicted in the forecast from 13 July, as is visualised by showing the probability of exceeding the 99th percentile of the model

climatology in the right-hand panel.

For the prediction of the heatwave, we here focus on maximum temperatures in London on 19 July, when London-Heathrow and St James’s Park in central London reached 40.2°C and London-Northolt measured 40.0°C. The ensemble median in both medium-range and short-range forecasts was well above the model climatology, with around 35°C (and similar values for the high-resolution forecast). However, this was far below the observed values. To further evaluate a specific forecast, we show the 3-day forecast from 17 July based on hourly values for these locations. Note that the maximum reported station value is based on continuous observations while the values here are the hourly SYNOP. For this 3-day period, the

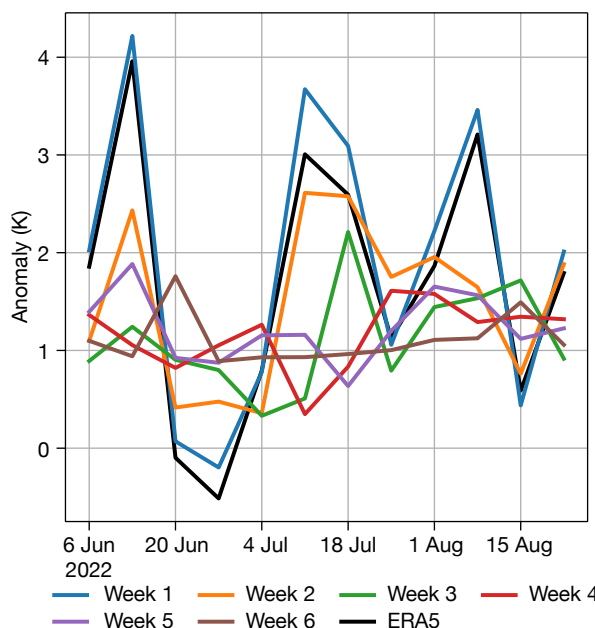
temperature was rising. This was captured by the forecasts. For the daily minimum, the forecasts captured the extremes well for Heathrow and St James's Park and overestimated the temperature at Northolt, which was colder than the other two stations. For the daily maximum temperature, we see that the forecast underestimated the peak for all three stations and all three days. For the hottest day, none of the ensemble members were close to reaching the maximum temperature at any of the three observation stations. This too low amplitude of the diurnal cycle is a recurring feature during heatwaves.

Summary

To summarise the skill of forecasts for European heatwaves in the summer of 2022 on different timescales, a signal for a warmer-than-normal summer was consistently present in 6-week forecasts from the extended-range system. On a 2-week timescale, the forecast started to distinguish between the extreme weeks and the more normal weeks. However, regarding the example of the record-breaking day over the UK, the forecasts failed to capture the magnitude of the extreme. These conclusions are similar to heatwaves in 2015, 2018 and 2019, also reported in ECMWF Newsletters.

From the predictability perspective, an area for further research is the relationship between European heatwaves and other extremes over the northern hemisphere, such as the precipitation anomaly over Pakistan/northwest India. In 2010, Pakistan also experienced flooding while western Russia had a severe heatwave. The relationship between these two regions has been investigated in the EU's Horizon 2020 CAFE project, where a connection between the BSISO index (as a proxy for the monsoon) and Russian heatwaves was established, while the connection to western Europe is less clear.

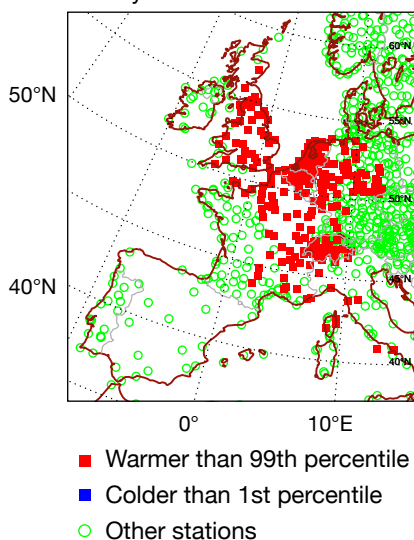
More details on temperatures in Europe in the summer of 2022 can be found on the website of the EU's Copernicus Climate Change Service (C3S) run by ECMWF: <https://climate.copernicus.eu/surface-air-temperature-maps>.



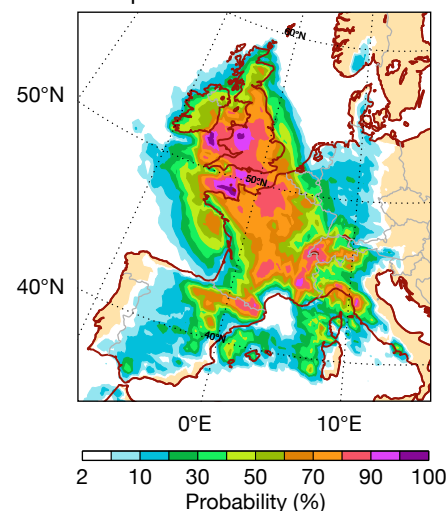
Weekly anomalies of 2-metre temperature over western Europe.

The chart shows weekly anomalies in the ERA5 reanalysis and according to ECMWF extended-range forecasts one to six weeks ahead over land points inside 35°N–55°N and 15°W–10°E.

Observed extremes on 19 July 2022

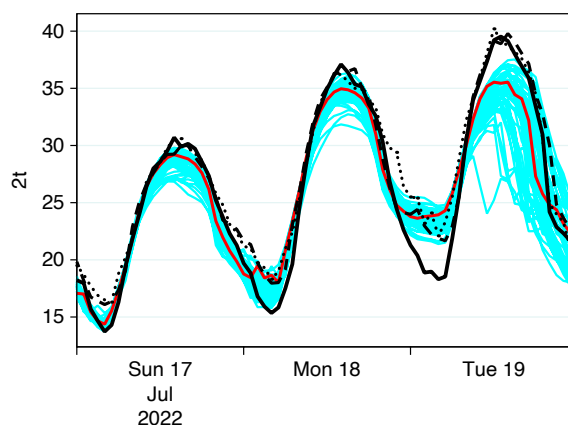


Probability of exceeding 99th percentile of model climate



Maximum 2-metre temperature on 19 July in the 99th percentile of climatology.

The left-hand chart shows observations warmer than the 99th percentile (red), colder than the 1st percentile (blue) and other stations (green), while the right-hand chart shows the probability of exceeding the 99th percentile of the model climate on 19 July in a forecast from 13 July 00 UTC.



Forecast for 2-metre temperature for London and corresponding observations.

The chart shows the high-resolution forecast (red), ensemble members (cyan), and three observation stations (Northolt – black solid, St James's Park – black dashed, Heathrow – black dotted).

ECMWF launches Massive Open Online Course on Machine Learning in Weather and Climate

Chris Stewart, Matthew Chantry (both ECMWF), Christian Weibel (University of Luxembourg Competence Centre)

Registrations are now open for a new Massive Open Online Course (MOOC) on Machine Learning in Weather and Climate! The MOOC will be launched in January 2023 by ECMWF in partnership with the International Foundation on Big Data and Artificial Intelligence for Human Development (IFAB).

The objective is to train a wider community on the impact and use of machine learning in numerical weather and climate predictions. It will bring together experts throughout Member and Co-operating States and beyond to provide a shared vision across the communities of Earth system sciences, high-performance computing and machine learning.

What is the MOOC and who is it for?

It is a fully online free training course, accessible to a global audience. In order to maximise inclusivity, the MOOC is distributed in time with a three-tiered approach. Learners may choose to follow one or more tiers. Within each tier, the programme is modular. Participants may select one or more modules within a particular topic subset, according to their interests.

The first tier is an introduction to machine learning in weather and climate. It is aimed at anyone interested in the topic, be they academics, from operational services, policymakers or the general public. Only a basic knowledge of weather and climate science, statistics and computing is assumed of learners.

Tier two takes a deeper look at the concepts of machine learning, and tier three demonstrates practical machine learning applications in weather and climate. These two tiers are more hands-on and include coding assignments, in which learners have the opportunity to apply machine learning to real-world problems. Tiers two and three are therefore more suited to technical data users from Member

The banner features the ECMWF logo on the left, followed by the text 'in partnership with IFAB' and 'INTERNATIONAL FOUNDATION ON BIG DATA AND ARTIFICIAL INTELLIGENCE FOR HUMAN DEVELOPMENT'. The main title 'MOOC Machine Learning in Weather & Climate' is prominently displayed in the center. Below it, a call to action reads 'Register now for our FREE online course!'. On the right side, there is a stylized globe with a network of white lines and dots representing data connections. The background is a vibrant blue with abstract white and yellow geometric patterns.

States, academia or industry across different sectors.

To allow for self-paced learning, the MOOC runs over ten weeks, with around three to four hours of study per week, for a total of 36 hours of training. The end of the MOOC coincides with the application phase of ECMWF's Summer of Weather Code (ESoWC), in which learners will be encouraged to submit proposals for coding projects with ECMWF mentoring and the chance to win a cash stipend.

About the training

The MOOC mixes interactive e-learning with webinars. Learning will be consolidated through a wide variety of activities, from videos and commuter-friendly podcasts to discussion forums and e-learning modules. Guided by a host of domain specialists from around the world, learners will be able to explore the subject from different angles. The webinars, for instance, are perfect occasions to interact with

leading experts from the field.

The acquired knowledge will then be honed in practical exercises. Using interactive Jupyter notebooks, learners will have the possibility of gaining hands-on practical experience with machine learning algorithms.

Throughout their whole learning journey, participants will be accompanied by a dedicated science communicator, who will break down the technical content into plain language whenever needed. Each completed tier will be rewarded with a certificate of completion. Those who complete all three tiers shall receive particular recognition with a more comprehensive certificate. At the end of the live run, the material will continue to be freely accessible from the ECMWF website.

Register now for the MOOC at <https://www.ecmwf.int/mlwc-mooc>: discover cutting-edge techniques, interact with leading experts and explore real-world practical applications in weather and climate!

ECMWF contributes to APEC Climate Center's BSISO prediction system

A-Young Lim, Hyung-Jin Kim, Bong-Geun Song (all APEC Climate Center, Republic of Korea)

The Boreal Summer Intraseasonal Oscillation (BSISO) is one of the dominant sources of subseasonal convective variability in the tropics. It exhibits northward/northeastward propagation over the Indian summer monsoon region, and northward/northwestward propagation over the western North Pacific and East Asia region. As a consequence, BSISO strongly influences the onset of the Asian summer monsoon, the active and break phase of the monsoon, and middle latitude teleconnections. BSISO predictions thus help monsoon-affected countries to prepare for hazards like heavy rainfall, droughts, and heatwaves. ECMWF has contributed to APCC (APEC Climate Center) BSISO products since they

started in 2013 by providing real-time extended-range weather forecast data.

Collaboration on BSISO predictions

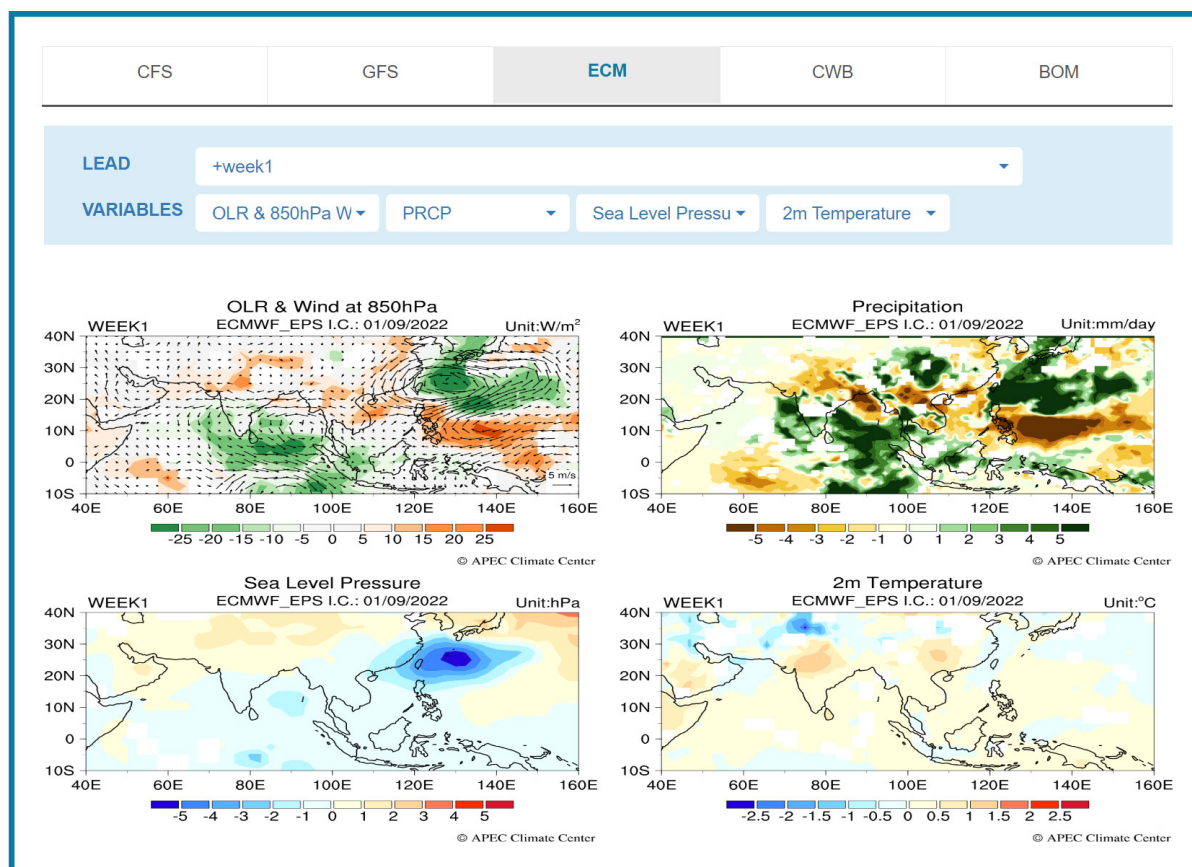
APCC, the only institution that regularly releases BSISO information from various sources, has provided BSISO monitoring and forecast products every day from May to October since 2013. With the goal of improving our ability to understand and predict BSISO, APCC's BSISO activity was initiated in close cooperation with the Madden Julian Oscillation Task Force of the Working Group on Numerical Experimentation under the World Climate Research Programme.

BSISO forecasts are produced using

five subseasonal-to-seasonal (S2S) prediction models from four operational centres worldwide. These are Australia's Bureau of Meteorology (BOM), the US National Centers for Environmental Prediction (NCEP – CFS and GFS models), ECMWF, and Chinese Taipei's (Taiwan's) Central Weather Bureau (CWB). Continuous support and the remarkable skill of ECMWF's model play an important role in making BSISO predictions more informative.

Products available

APCC's BSISO products include phase diagrams, index-based information, anomalies for monitoring and prediction, and simple verifications. In addition to these



BSISO impact anomaly using ECMWF's S2S prediction model. The BSISO impact anomaly shows the expected local impacts of BSISO activity for up to the next three weeks over the Asian monsoon region. These are reconstructed fields based on the composite of the historical BSISO index (1991–2020) in combination with ECMWF model predictions.

conventional services, APCC has recently started to provide BSISO impact anomalies for ten important variables up to three weeks ahead, including temperature, precipitation, wind, and other relevant variables. These impact anomalies are reconstructed based on the composite of the historical BSISO index in combination with model predictions. They are expected to capture the local impacts of BSISO activity over the Asian monsoon region.

In 2022, APCC improved BSISO information by reflecting recent climate characteristics and their influence. To do this, the climatological reference period has been updated from 1981–2010 to 1991–2020 for all BSISO information, following the World Meteorological Organization’s recommendation to use the most recent 30-year period as the climatological standard. Graphical representations on the BSISO website were also improved to increase readability and visibility.

We expect that this improvement will help enhance user experience.

User support

BSISO monitoring, forecasts, and fundamental verification are posted on the APCC website (<https://www.apcc21.org/>). Users are also able to download digitised data, such as most recent and historical monitoring and forecast indices, through APCC’s Climate Information toolKit (<https://cliks.apcc21.org/dataset/bsiso>).

Institute	Model	Ensemble size	Forecast period (day)	Update frequency	Resolution	First operational run from APCC
NCEP	Climate Forecast System (CFS)	4	40	Once a day	T126 L64	July 2013
	Global Forecast System (GFS)	1	16	Once a day	T574, T190 L64	July 2013
BOM	Australian Community Climate and Earth-System Simulator Seasonal 2 (ACCESS-S2)	33	30	Once a day	N216 L85	August 2021
ECMWF	ECMWF Ensemble Prediction System	51	32	Twice per week	T639, T319 L62	July 2013
CWB	Taiwan CWB 1-Tier Global Atmosphere-Ocean Coupled Forecast System (TCWB1T1.1)	3	40	Every 5 days	T119 L40	May 2020

Participating S2S prediction models. A brief summary of five S2S prediction models participating in the APCC BSISO prediction system.

Developing an agriculture application for Africa

Kaah Menang (ECMWF and National Department of Meteorology, Cameroon)

In September 2021, I began a 12-month World Meteorological Organization (WMO) Fellowship at ECMWF’s headquarters in Reading, UK. The aim of the fellowship was to develop an application for Africa within the Climate Data Store (CDS) team of the Copernicus Climate Change Service (C3S).

Soon after completing my PhD in atmospheric radiation at the University of Reading in 2012, I was recruited as a meteorologist at the National Department of Meteorology, Cameroon. In Cameroon, my main responsibility is to produce marine weather forecasts, but I also contribute to producing forecasts for other sectors. In addition, I carry out research on climate extremes and atmospheric radiation. Undertaking

the WMO Fellowship at ECMWF was a unique opportunity for me to enhance my professional skills as well as to contribute to the development of weather and climate forecasting in Cameroon and Africa.

An agriculture application for Africa

During the fellowship, I developed an application in the CDS toolbox (cds.climate.copernicus.eu/cdsapp#!/toolbox) that supports agriculture in Africa. This was motivated by the fact that agriculture is a major economic activity in Africa. For example, in Cameroon it accounts for over 70% of the labour force and 40% of foreign earnings. Past and future information on climate variability and

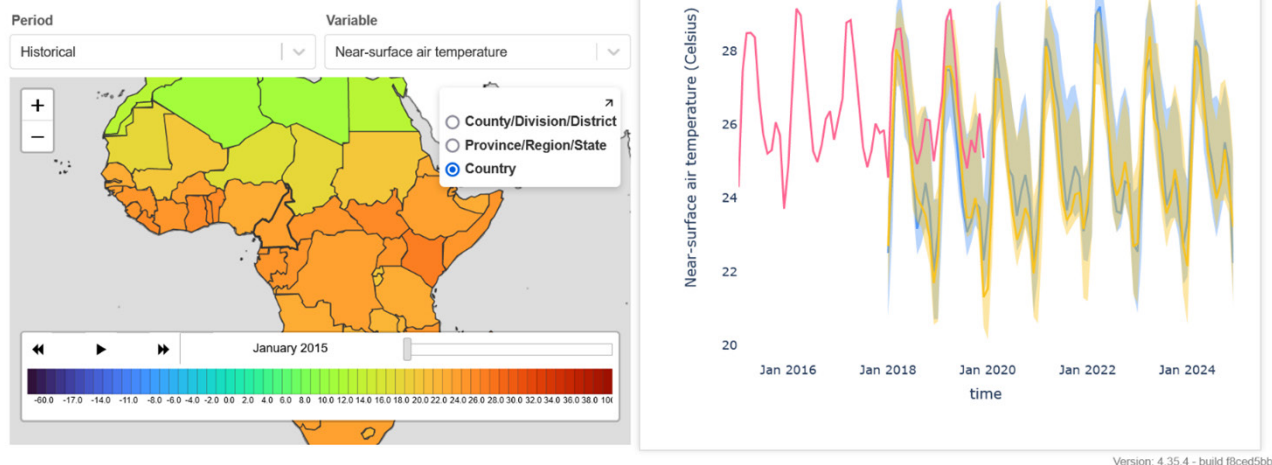
change is crucial for sustainable agriculture in Africa.

The application is called ‘African Agroclimatic Indicators’. It is based on ECMWF’s ERA5 climate reanalysis data and on climate projection data from the World Climate Research Programme’s Coupled Model Intercomparison Project Phase 6 (CMIP6). It is targeted at agro-industries that cultivate biennial and perennial crops. This application will provide long-term past and future climatic information required by these industries for their farming practices and crop yield projections. The application is also aimed at assisting policymakers in decision-making.

African Agroclimatic Indicators

This application presents the trends of some relevant agroclimatic indicators across Africa in the last several decades. The application can also be used to assess how these indicators will evolve in the coming decades for two CMIP6 climate projections scenarios: SSP2-4.5 (with medium greenhouse gas emission) and SSP5-8.5 (with high greenhouse gas emission). This application is thus aimed at assisting agricultural practitioners and policy-makers in decision-making.

Select the period, scenario(if applicable) and variable of interest from the drop down menu to visualise on the map. Chose any of the three administrative levels shown on the map for this visualisation. Click on any area of chosen administrative level to view a time series plot of the variable. Hover over the time series plot to read off values of the variable.



Version: 4.35.4 - build f8ced5bb

Example of the maps and charts developed. ERA5 near-surface average air temperature of African countries (left panel) and time series plots of ERA5 and CMIP6 data for Cameroon (right panel).

More precisely, the datasets used are:

- ERA5 monthly averaged data on single levels from 1959 to the present, and
- monthly averaged CMIP6 projections of 12 models at single levels and two Shared Socioeconomic Pathway (SSP) scenarios (medium and high greenhouse gas emissions), covering the period 2006–2100.

The indicators selected for this application are average 2-metre temperature, total precipitation, soil moisture (upper level), evapotranspiration and 10-metre wind.

These indicators, from both reanalysis (historical) and projection data, can be visualised on an active map of Africa averaged over three administrative levels: country, provincial/state-wide/regional and divisional/county/district levels. A time series of both the ERA5 and CMIP6 data for any of the indicators can be obtained by clicking on the chosen administrative level. The figure shows an example of these visualisations for a few selected years (the final application will include all the years stated above). The shading on the CMIP6 time series shows the percentile range of values in the projection data.

Other activities during the Fellowship

During this fellowship, I took the opportunity to further develop my professional knowledge and skills. In addition to studying some of ECMWF's e-learning courses (available at learning.ecmwf.int), I attended two ECMWF training courses ('Use and Interpretation of ECMWF products' and 'A Hands-on Introduction to NWP models'), and the Using ECMWF's Forecasts (UEF2022) event. Using the CDS toolbox to develop African Agroclimatic Indicators also helped to improve my Python programming skills.

Future work and recommendation

The African Agroclimatic Indicators application is in its final stages of development and will be made available to the public on the CDS in the next few months.

As discussed above, this application was developed for the agro-industries that specialise in biennial and perennial crops. It is equally important to develop another application using seasonal forecasts for annual crop farming (which is practiced by small-scale or subsistence farmers as well as by

some industrial farmers).

The CDS toolbox, which has wrappers in Python, is a very powerful tool for downloading climate data and developing applications based on these data. With basic knowledge of Python programming at the start of my fellowship, I found that it was a steep learning curve to get to grips with the toolbox. At the moment, it would appear that the toolbox can only be effectively used by advanced Python programmers. In the future, I would recommend that the toolbox is re-developed to make it more of a direct Python and interactive interface (e.g. in Jupyter notebooks or a similar environment). This interactive mode will undoubtedly assist users with less advanced Python programming skills to gain confidence in using even the more advanced features of the CDS toolbox.

Work is under way in the CDS team at ECMWF to re-develop the CDS toolbox to make it more user accessible. I believe my experiences during the fellowship have been very useful to demonstrate the development required.

More information on the WMO Fellowship scheme can be found here: <https://www.ecmwf.int/en/learning/training/wmo-fellowship-scheme>.

Fire can now be found on MARS

Francesca Di Giuseppe, Chris Barnard, Pedro Maciel, Shahram Najm, Sébastien Villaume

ECMWF is the computational centre for fire danger forecasts under the umbrella of the EU’s Copernicus Emergency Management Service (CEMS). On 8 August 2022, it started archiving its operational fire danger forecast model output into ECMWF’s Meteorological Archival and Retrieval System (MARS). This marked the end of a long project to define new GRIB2 templates and data models for storing post-processed data in MARS with traceability back to the forcing model. We benefited from what was created back in 2018 to ingest hydrological data of the European Flood Awareness System, which is also operated by ECMWF as part of CEMS, into MARS. Even though hydrological fields have been in MARS since then, we still faced several challenges due to the specific nature of the fire danger forecast dataset. At the end of this journey, fire danger forecasts can now be found in MARS. This achievement marks a milestone in ECMWF’s contribution to CEMS operations as it will ensure a service more robust in terms of archiving, more flexible in terms of product generation, and more easily accessible to all users. Ultimately, the development will bring the service closer to FAIR Data Principles (Findable, Accessible,

Interoperable, and Reusable), published in Scientific Data in 2016, which set guiding principles to support the reusability of digital assets.

The development steps needed

Outputs from the Global ECMWF Fire Forecast (GEFF) model are produced daily to provide fire danger indices to the European Forest Fire Information System (EFFIS). To ease data access, it was decided a few years ago to make the model output available through MARS. In 2020, GEFF v4 was released with a GRIB I/O interface and the removal of netCDF support. The GRIB interface allows a more structured delivery of data and improves internal workflows in terms of validation of the operational forecasts and experiments.

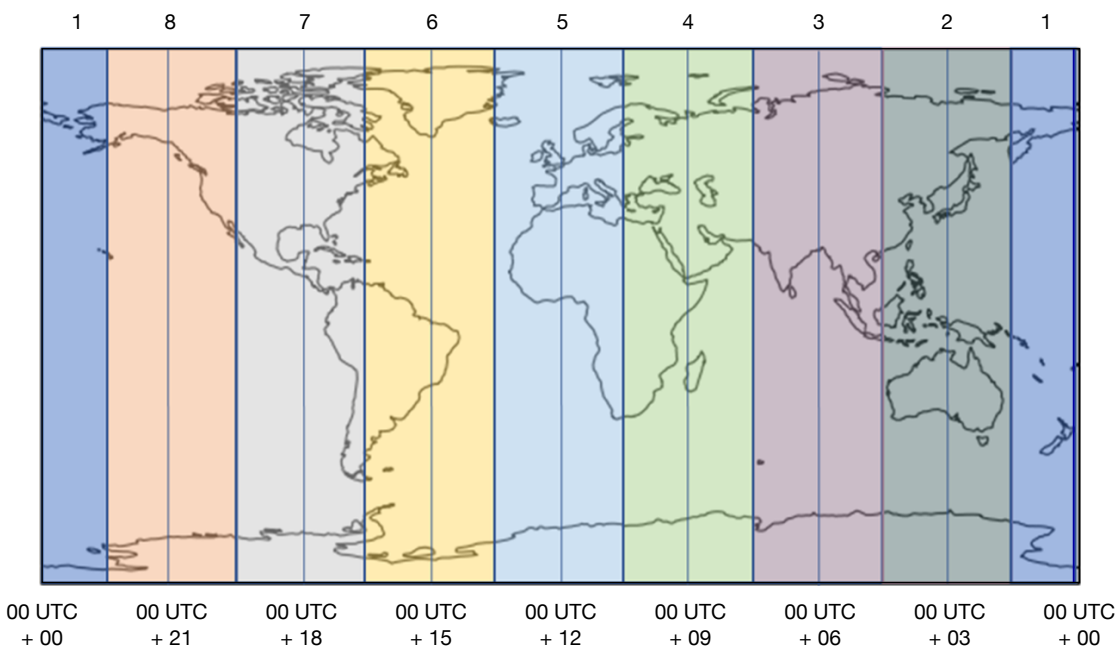
However, it soon became evident that there were several obstacles to including fire-related variables in MARS. Firstly, fire danger forecast model variables had no GRIB code definitions. New parameters were proposed to the World Meteorological Organization (WMO) and were accepted after revision. Secondly, there was no useful template for a forecast expressed in terms of local

time at each location. For our intents and purposes, we needed to define a new template where the important metadata from a collection of forcing time stamps was stored in the GRIB header of the fire danger forecast model output.

Implemented solution for local time template

A model integration at any time in UTC will simulate the atmospheric conditions at a different local time depending on the location. However, fire danger indices should be calculated at local noon, when fire hazard is thought to be at its maximum. To achieve this, a temporal and spatial collage of 24-hour model simulations is performed to produce a snapshot at 12:00 local time. Thus, for example temperature and relative humidity fields are cut into 3-hourly time strips using the closest 3-hour forecast output and then joined together, so that the final field is representative of the conditions around the local noon within the 3-hour resolution available. This is highlighted in the figure.

The implementation of the new local time GRIB2 templates allows the storage of information on the forcing



Concatenation of forecasts to derive a field at 12:00 local time everywhere.
The double line on the right represents the change of date. The stripes correspond to the forecast times specified at the bottom.

data through the key 'localTimeMethod'. Setting the key 'localTimeMethod' enables the use of additional metadata about the forcing data used in creating the field at local time. At the moment, two options are implemented: 'localTimeMethod = 0' generates an output where the nearest forecast or analysis in time is selected to create the local time, and 'localTimeMethod = 1' generates an output by interpolating between the two forecast steps that cross the local time. The template exists in eight versions to account for any combination of deterministic and

ensemble forecasts (two types), instantaneous and accumulated fields (two types), as well as native fields vs post-processed fields (two types).

Benefits and future work

The new templates for local time data were made purposely generic to account for different types of output and applications using numerical weather prediction output.

The template has so far been implemented for GEF operational forecasts. Reanalysis, the extended-range hindcast, and seasonal fire

danger forecasts will soon be added. Other downstream applications, such as the Universal Thermal Comfort Index (UTCI) and drought indicators, such as the standardised precipitation index, will follow suit. The fire danger forecast data is currently archived under the new class 'Global wildfire', abbreviated 'gw' in MARS. One of the immediate benefits of this implementation will be access to non-restricted data by the public through the Copernicus Climate Data Store, which is part of the Copernicus Climate Change Service (C3S) run by ECMWF.

Progress on emulating the radiation scheme via machine learning

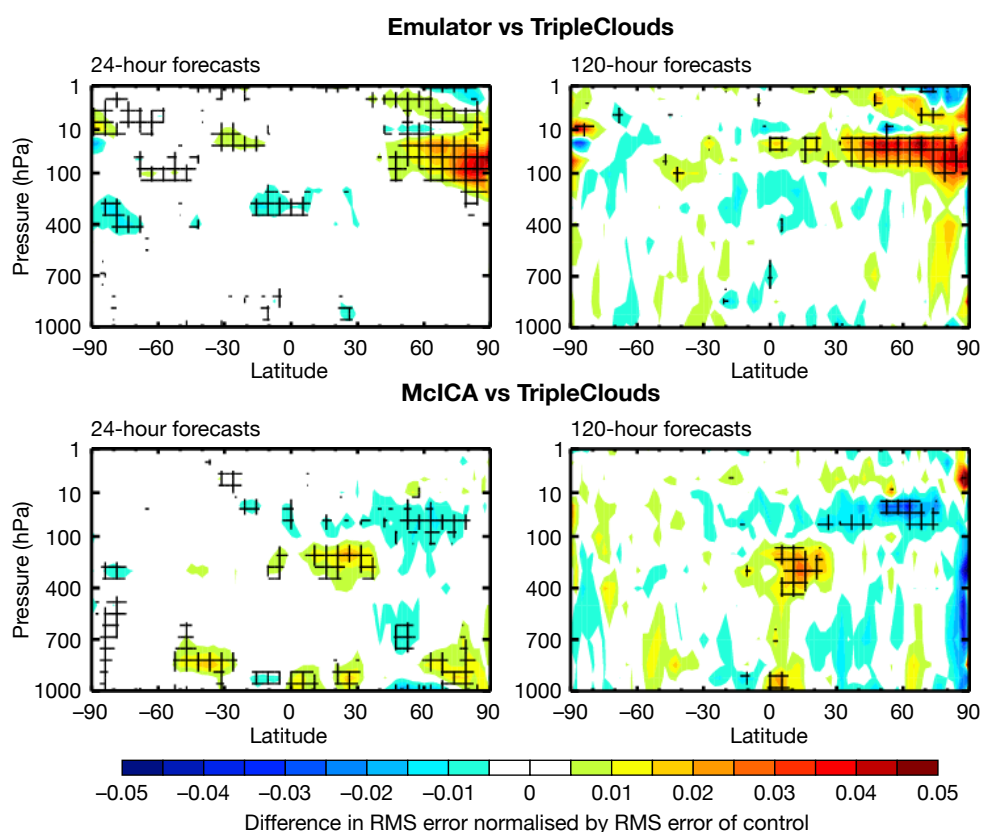
Matthew Chantry, Peter Düben, Robin Hogan (all ECMWF), Peter Ukkonen (Danish Meteorological Institute)

One of the most promising applications of machine learning in weather and climate modelling is the creation of so-called emulators. An original model component is emulated by running the model and storing input and output pairs of the model component to train a

machine learning tool – typically based on deep neural networks. The emulator can then be used within the forecast model. Here we describe progress in emulating the radiation scheme used in numerical weather forecasting.

Why is emulation useful?

For computationally expensive model components, the emulator may be much cheaper than the original component. For instance, emulators could be used to speed up operationally used parametrizations to



Changes in temperature error.

The charts show the change in root-mean-square (RMS) error of temperature if the emulator is used to represent the TripleClouds radiation scheme in 24-hour and 120-hour forecasts (top), and the same change if McICA, the current operational scheme, is compared with TripleClouds (bottom). The experiments were carried out from 1 June to 31 August 2021. Red indicates where the emulator/McICA forecasts are degraded relative to TripleClouds, blue where they are improved. Cross-hatching indicates at least 95% confidence.

save computing time, or to make expensive components such as cloud bin microphysics schemes affordable in an operational forecast. They are also more portable to heterogeneous hardware using existing machine learning libraries. Neural networks can be trained to produce accurate answers at reduced numerical precision, as low as half precision with 16 bits per variable, providing an additional route for performance improvement.

Towards emulating the radiation scheme

As a first step, the emulation of the gravity wave drag parametrization scheme has been considered. Results have shown that a deep-learning emulator can indeed represent the scheme correctly. It can also be used to generate tangent linear and adjoint model code for use within 4D-Var data assimilation. However, to make the approach viable for operational implementation, a more costly component of the Integrated Forecasting System (IFS) – the ecRad radiation scheme – would need to be emulated. Consequently, tests to emulate the ecRad radiation scheme were performed as part of the MAELSTROM EuroHPC-Joint Undertaking project. Datasets have been developed and published for the purpose of emulating radiative transfer within the IFS. These datasets can be used to learn the TripleClouds solver, which will be deployed operationally in the IFS in an upcoming cycle, or the more expensive SPARTACUS solver,

which incorporates 3D cloud effects.

Accurate neural network emulators have been created for the TripleClouds solver. Shortwave and longwave processes are solved separately. For the shortwave process, a series of recurrent neural networks, chosen to mimic the existing algorithm, was created by Peter Ukkonen from the Danish Meteorological Institute and has been adopted here. For the longwave process, a convolutional-based neural network is used. In simulations at about 29 km resolution, when coupling the neural network solution instead of the TripleClouds solver to the IFS, we find good agreement in the troposphere compared to the reference simulation, with only small degradations in the stratosphere above 100 hPa (see the top two figure panels). To contextualise these changes, we also plot the impact using the current solver, McICA, instead of the upcoming TripleClouds solver, which was used to generate the machine learning data (see the bottom two panels).

However, emulators will eventually only be useful in the IFS if they can either be more efficient or portable when compared to the conventional model component; if they emulate a version of the parametrization scheme that would otherwise be too expensive for operational use; or if they can be used to generate tangent linear and adjoint model code for data assimilation.

The Inero library was used to execute

machine learning libraries into the code framework of the IFS. So far, tests have been limited to CPU hardware, and improvements in efficiency are only ~25%. However, in the next step the execution on GPUs will be tested, and significant improvements in efficiency are anticipated here.

Additional tests

Work has also begun to investigate the emulation of expensive versions of the parametrization scheme. A study in collaboration with the University of Reading has shown that neural networks can represent 3D cloud effects as represented in the SPARTACUS solver. Here, emulators were trained to predict the difference between the default TripleClouds radiation scheme and SPARTACUS. Results have been published, but model evaluation has so far only been done in offline simulations, so further work on the coupling and applications within IFS simulations will be required.

Finally, tests using the emulator to generate a tangent linear or adjoint version of the neural network emulator will be performed shortly. While it is still unclear whether simulations will stay stable during the 4D-Var minimisation process, the approach is very promising: it will represent the latest version of the radiation scheme in simulations, while the current data assimilation is still performed with the tangent linear and adjoint model code that was derived for an old version of the radiation scheme.

ECMWF real-time data benefit the wider research community

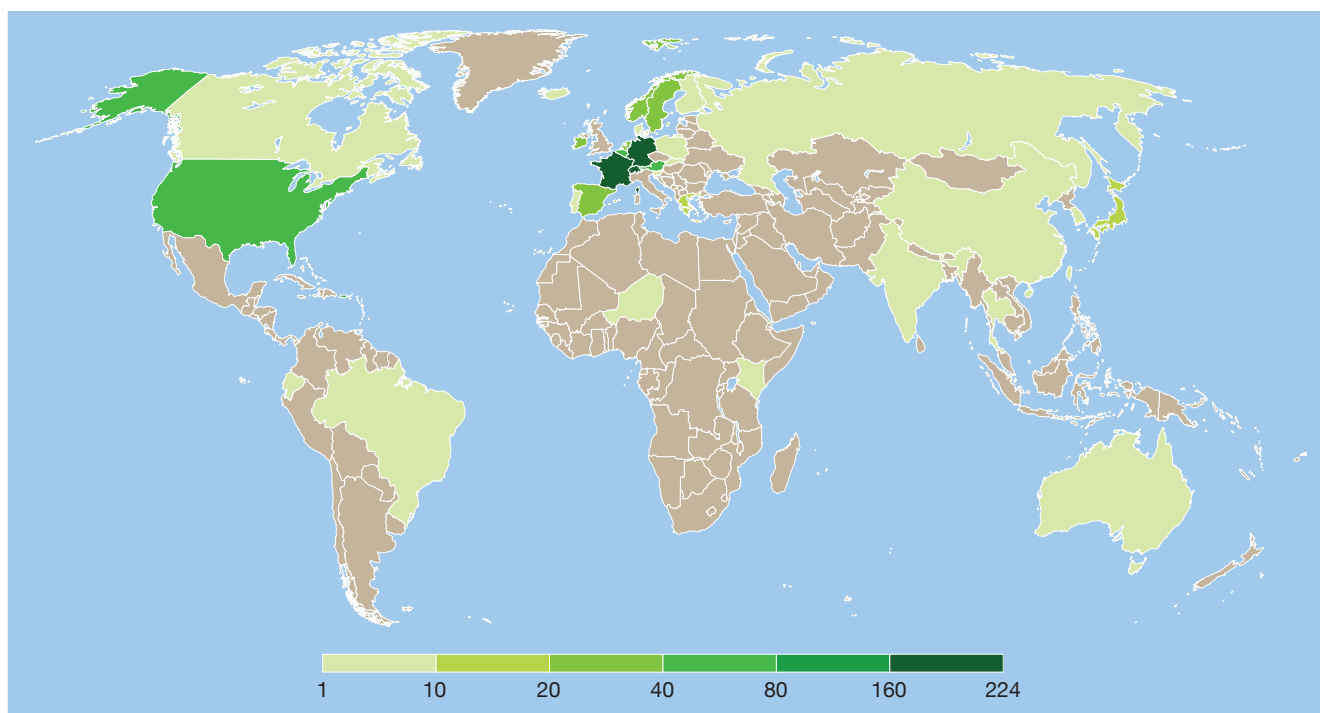
Emma Pidduck, Umberto Modigliani

In addition to providing weather forecast data and services to its Member and Co-operating States, ECMWF serves a much wider community of users with specialist licences and cost-free real-time data. These are provided in accordance with the ECMWF rules governing the distribution of real-time data, first established in 1994. Specialist licences have been created for different categories of

users, including research projects, satellite data providers (e.g. EUMETSAT, the European Space Agency, NASA), national meteorological and hydrological services (NMHS) of World Meteorological Organization (WMO) Members, as well as UN and EU agencies (e.g. the EU's Joint Research Centre).

For research projects, over 1,200

real-time research licences have been approved since 2012. These licences provide free forecast data and access for time-limited, non-commercial research projects (subject to management/Member State approval). The aim of a research licence is to facilitate scientific advancement where the project results are openly accessible to the weather and climate community, while prohibiting



Research licences by country. Geographic distribution of real-time research licences since 2012.

commercial exploitation and the use of data in service provision. In addition, research projects can support long-term collaborations with external organisations, and the exchange of data/results can provide verification and validation of ECMWF forecasts and can contribute to the general development of ECMWF's Integrated Forecasting System (IFS). In addition to these research projects, there are over 7,000 registered users of the MARS archive, each of whom can make use of historical forecast data for their research purposes.

To date, supported research projects have ranged from one month to five years in length, with the average research licence lasting for 1.5 years. Licences have been granted to projects in about 40 countries. ECMWF has recently observed a significant growth in the number of projects relating to artificial intelligence and machine learning, but also in the use of ensemble forecasts for hazard forecasting and long-term forecasts for energy prediction improvement.

Open data to support research

In addition to cost-free data for research projects, ECMWF has a range of open datasets that can be

Examples of work involving ECMWF research licences

The Year of Polar Prediction: <https://www.ecmwf.int/en/newsletter/165/news/yopp-site-inter-comparison-project>

Atmospheric river reconnaissance: <https://www.ecmwf.int/en/newsletter/171/news/ecmwf-participates-atmospheric-river-reconnaissance>

Operational observation targeting: <https://www.ecmwf.int/en/>

[newsletter/172/earth-system-science/using-ecmwf-ensemble-forecasts-operational-observation](https://www.ecmwf.int/en/newsletter/172/earth-system-science/using-ecmwf-ensemble-forecasts-operational-observation)

Mount Everest: <https://journals.ametsoc.org/view/journals/bams/101/11/bamsD190198.xml>

United Nations/European Commission: <https://www.ecmwf.int/en/newsletter/165/news/using-ecmwf-data-humanitarian-support>

used for research projects, including the S2S dataset for seasonal-to-subseasonal research and the TIGGE dataset of global medium-range ensemble forecasts from 13 forecasting centres. Most of the datasets are governed by the Creative Commons CC-BY-4 licence, which permits both commercial and research use, provided data is attributed to the owner. More information can be found in our dataset catalogue: <https://www.ecmwf.int/en/forecasts/datasets>.

Specialist licences

National meteorological and hydrological services of WMO countries can apply for a non-commercial licence at data delivery

charges only. The purpose of the licence is to enable the fulfilment of national governmental obligations related to the protection of life and property. More details can be found here: <https://www.ecmwf.int/en/forecasts/accessing-forecasts/licences-available>.

Real-time data usage by Special Projects

Over the last ten years, ECMWF has supported several hundred Special Projects. Special Projects are 'experiments or investigations of a scientific or technical nature, undertaken by one or more Member States, likely to be of interest to the general scientific community'. Users

within one of ECMWF's Member States may apply for resources as a 'Special Project'. These research activities are supported through the provision of supercomputing resources on ECMWF's high-performance computing facility.

How to obtain a research or specialist licence

ECMWF and all Member and

Co-operating States are permitted to provide real-time data for research purposes, where the research itself is not providing an operational service or used commercially. To access the data, please use the ECMWF Product Requirements Catalogue (<https://apps.ecmwf.int/shopping-cart/>) to configure your data requirements and submit it to your chosen Licensor (Member or

Co-operating State). Note that all research licences are subject to Licensor approval and some delivery fees may apply.

To discuss options, please contact your Member or Co-operating State computing representative (<https://www.ecmwf.int/en/about/contact-us/computing-representatives>) or visit the ECMWF support portal (<https://www.ecmwf.int/en/support>).

Seasonal forecasts of extreme rain in Pakistan

Antje Weisheimer, Magdalena Balmaseda, Tim Stockdale

The floods that inundated large parts of Pakistan during the summer of 2022, causing a huge human and economic disaster, were triggered by torrential and persistent monsoon rainfalls from mid-June through to August. ECMWF's seasonal forecasting system SEAS5, and other seasonal prediction models from the EU's Copernicus Climate Change Service (C3S) implemented by ECMWF, successfully predicted extreme rainfall during the summer over the affected area in forecasts issued at the beginning of May 2022.

Successful seasonal forecasts

SEAS5's forecasts for June–July–August 2022 initialised on 1 May 2022 indicated a strong signal of extreme rainfall over Pakistan. The upper panel of the figure shows the operational forecast product of the probability of being in the highest 20% of the model's climatological distribution as an example. For Pakistan, the probability of rainfall in that highest quintile category was drastically increased to 50–70%. Our confidence in these predictions is a priori high since the SEAS5 historical hindcasts have demonstrated significant skill for high rainfall in this region and season (see <https://bit.ly/3EQ5E2Q>). The location and intensity of the precipitation anomaly was also well predicted at the extended forecast range (up to 46 days ahead) – see the article by Magnusson et al. in this Newsletter.

The good performance of ECMWF's seasonal forecast was no exception amongst other forecasting systems. The C3S multi-model seasonal forecast ensemble also predicted a high probability of extreme rainfall (see the middle panel of the figure). This ensemble includes, in addition to SEAS5, forecast models from the UK Met Office, Météo-France, the Euro-Mediterranean Centre on Climate Change (CMCC) in Italy, the German National Meteorological Service (DWD), the US National Centers for Environmental Prediction (NCEP), the Japan Meteorological Agency (JMA), and Environment and Climate Change Canada (ECCC).

Consistent with these predictions, the South Asian Seasonal Climate Outlook Forum indicated a strongly increased probability of above normal June to September 2022 rainfall for Pakistan, which was followed by high alerts from Pakistan's National Disaster Management Authority.

An estimate of precipitation anomalies from the ERA5 reanalysis is shown in the lower panel of the figure, confirming the agreement in the broad geographical location of the rain with the seasonal forecasts.

An exceptional season

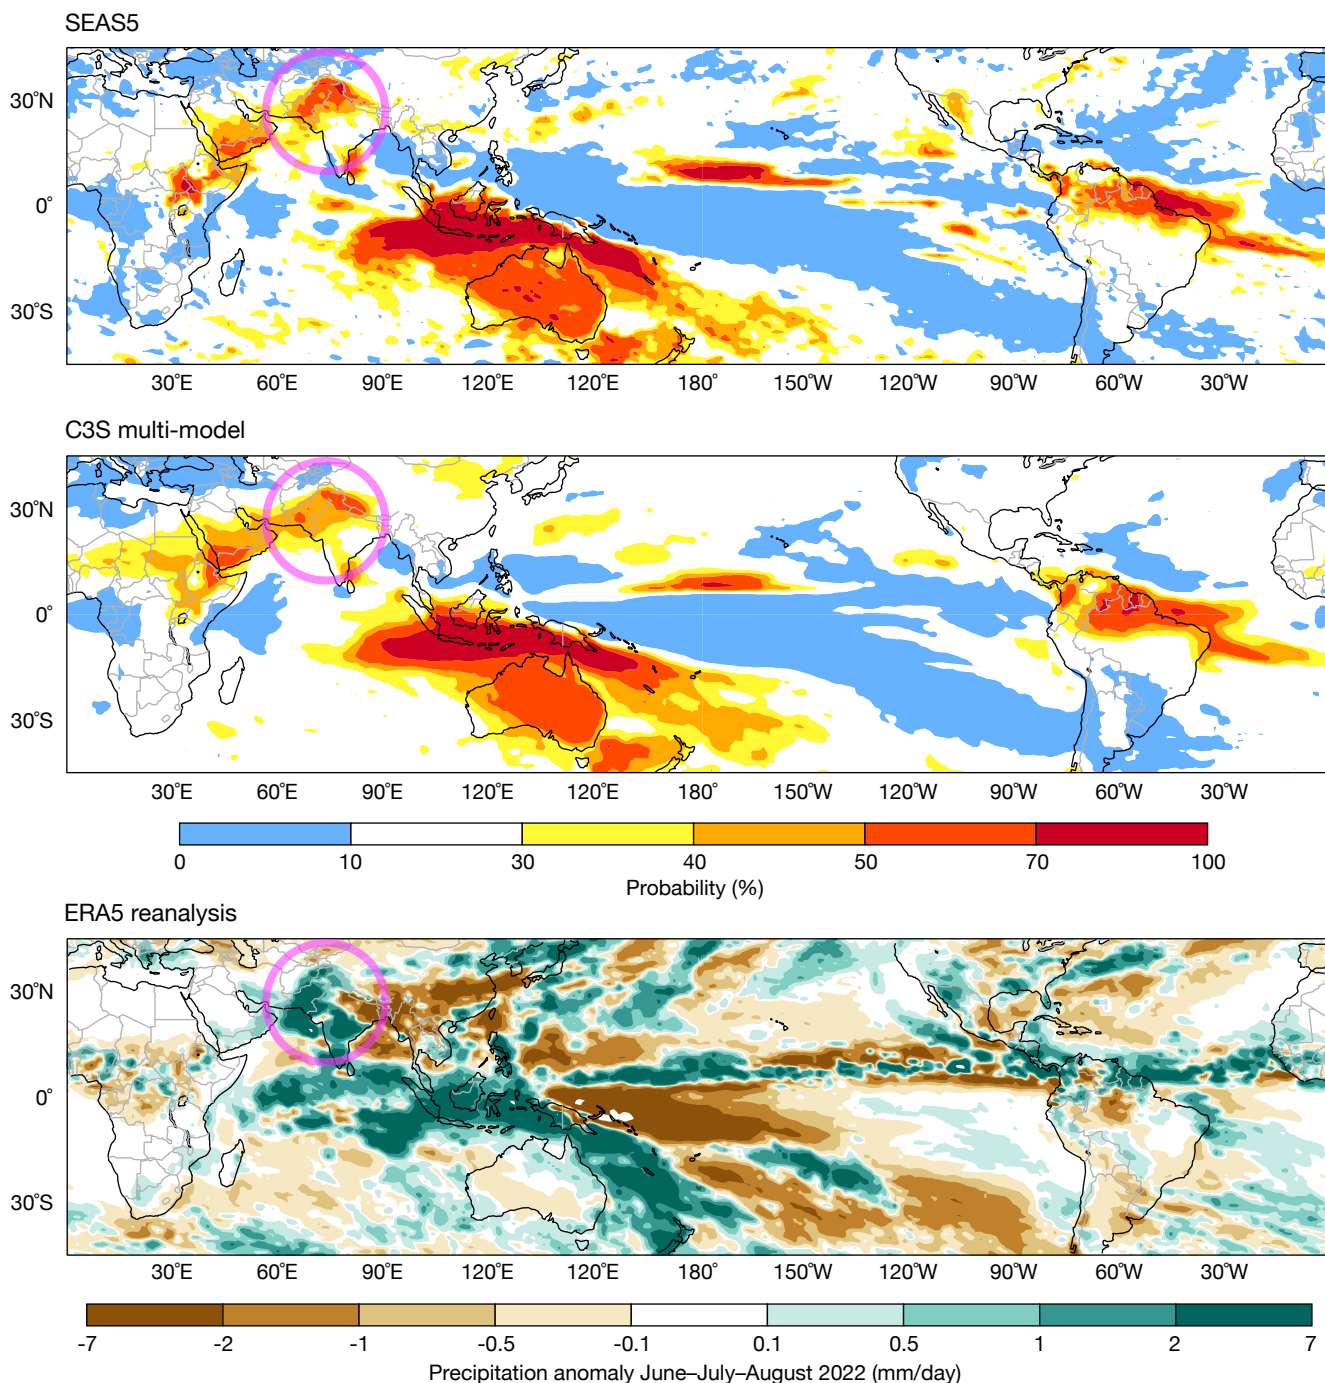
The northern summer season 2022 was characterised by La Niña cold conditions in the tropical Pacific persisting for a third year in a row. The El Niño–Southern Oscillation (ENSO) is known to impact the

South Asian summer monsoon, with La Niña years related to stronger cyclogenesis activity over the Bay of Bengal and higher than normal rainfall over the Indian subcontinent. Pakistan is among the countries that are enormously affected by ENSO.

Furthermore, the tropical Indian Ocean displayed a pronounced positive dipole structure during the summer, known as the Indian Ocean Dipole (IOD), with warmer-than-normal sea-surface temperatures (SSTs) in the eastern parts and colder-than-normal SSTs in the western areas of the basin. A positive IOD is associated with a stronger-than-normal South Asia southwest monsoon.

Both these large-scale tropical sea-surface anomalies were well reproduced in the seasonal forecasts from SEAS5 and C3S (see https://climate.copernicus.eu/charts/c3s_seasonal/) and likely played a vital role in the successful predictions of the Pakistan rains. Other less-well investigated potential factors that influenced the extreme precipitation include tropical Atlantic SSTs and Eurasian land heating.

These successful predictions are an example of the power of seasonal forecasts to provide early warnings and enhance resilience against climate hazards. Building confidence in the physical and computational sciences underpinning the forecasts is an important step towards improving the information value chain.



Precipitation in June–July–August 2022. The top two plots show the seasonal forecast probability of being in the upper 20% of the climatological distribution for SEAS5 (top) and the C3S multi-model ensemble (middle), while the bottom plot shows the ERA5 precipitation anomaly. The climatological period used is 1993–2016. All seasonal forecasts were initialised at the beginning of May 2022. The highlighted region includes Pakistan.

New observations since mid-June 2022

The following new observations have been activated in the operational ECMWF assimilation system since mid-June 2022.

Observations	Main impact	Activation date
Sea-level anomaly from Sentinel-3A and -3B	Ocean analysis	22 June 2022 in the OCEAN5 system

Aeolus positive impact on forecasts with the second reprocessed dataset

Michael Rennie, Sean Healy, Saleh Abdalla, Will McLean, Karen Henry

Aeolus, the European Space Agency's (ESA) novel Earth Explorer satellite, retrieves profiles of horizontal wind information in the troposphere and lower stratosphere via a Doppler wind lidar instrument. The mission is a 'proof of concept', and our understanding of the measurement characteristics and usage has evolved rapidly since launch in 2018.

Direct observations of wind profiles have limited spatial coverage in the global observing system. Therefore, by partly filling this gap, Aeolus provides an important contribution of such information in numerical weather prediction (NWP) and improves the initial conditions of forecasts. We assess the impact of Aeolus horizontal line-of-sight (HLOS) winds from the best dataset available so far – the second reprocessed dataset covering a 15-month period during 2019 and 2020. Aeolus wind processing algorithms have improved during the mission, and hence reprocessing the data improves quality. The largest signal-to-noise ratio of the mission was available during the early part of this period. A good positive impact is shown for a single space-based instrument: improving wind, temperature and humidity forecasts by up to 1–2% in the medium forecast range in the lower stratosphere, with smaller

but still statistically significant improvements at lower altitudes. In 2019, Aeolus also provided up to 5% of the short-range forecast impact partitioned by observation type. Gradual degradation of the signal-to-noise ratio due to instrument issues is responsible for a reduction in the impact of Aeolus in 2022. The operational impact from Aeolus will be an important consideration when assessing a future operational Doppler wind lidar mission to be run by the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT).

Aims of Aeolus and realised applications

Aeolus' payload is the world's first functioning space-based Doppler wind lidar and Europe's first space-based lidar. The aim of the mission is to demonstrate this new technology in space for the benefit of weather forecasting and to improve the understanding of atmospheric dynamics, especially in the tropics, where both mass and wind information are required to constrain the analysis. Although the design life of Aeolus is three years, it completed four years in space on 22 August 2022.

Aeolus orbital coverage and wind information in a 12-hour period is shown in Figure 1. The near vertical 'slices' (tilted off-nadir by ~37 degrees) are the wind component information for 9 February 2020, which had

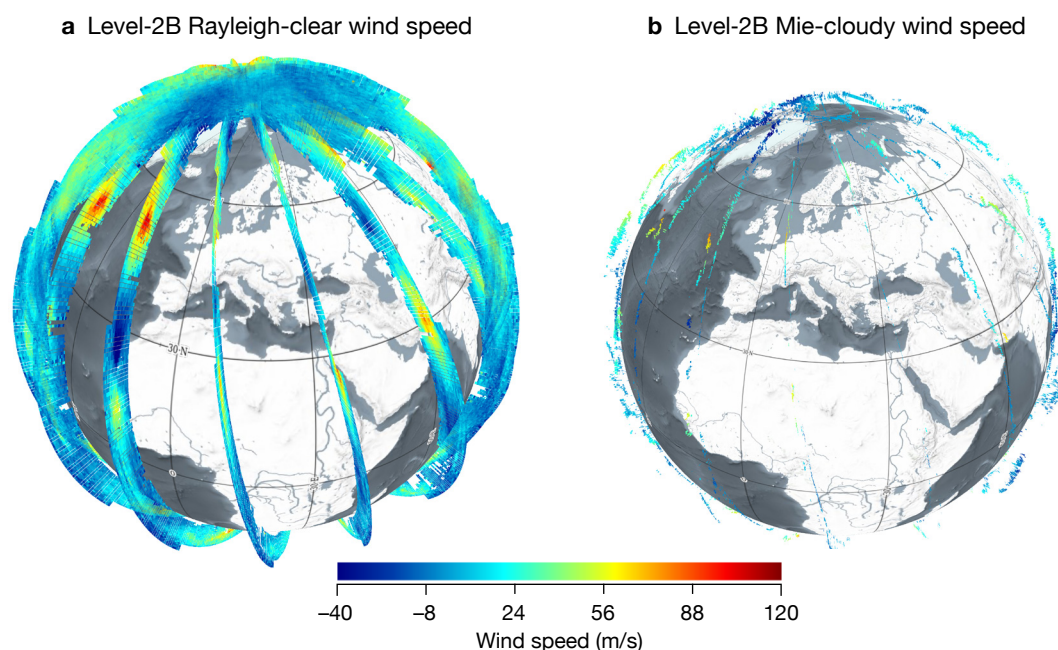


FIGURE 1 Illustration of Aeolus HLOS winds (m/s) over a 12-hour period on 9 February 2020 for (a) Level-2B Rayleigh-clear and (b) Level-2B Mie-cloudy. Lidar 'curtain' style plots from 0–25 km altitude are shown. Produced using the VirES tool.

near record-breaking jet stream winds over the Atlantic Ocean. Most of Aeolus wind coverage is from the Rayleigh-clear winds (backscatter signal from molecules in clear-air), but the Mie-cloudy (particulate backscatter from clouds and aerosol) have significantly smaller random errors and better horizontal resolution than the Rayleigh.

Aeolus winds have been proven to be of sufficient quality and coverage to significantly improve operational NWP forecasts, particularly in the tropics. They have been used operationally by NWP centres for over 2.5 years: pioneered by ECMWF, followed by Germany's DWD (Deutscher Wetterdienst), Météo-France, the UK Met Office and India's NCMRWF (National Centre for Medium Range Weather Forecasting). For details, see Rennie et al., 2021; Pourret et al., 2022; Laroche and St-James, 2022.

When referring to Aeolus winds, we mean the Level-2B processor wind retrievals, which are suitable for use in NWP and scientific research, due to several important extra processing steps relative to the Level-1B data. ECMWF was able to benefit from the data quickly because of its strong involvement with the mission. It was, for example, part of the Mission Advisory Group and L2B processing algorithm development team before and after Aeolus was selected as an ESA Earth Explorer Mission in 1998. The L2B winds are processed as part of the ground segment at ECMWF, and the products are forwarded to ESA and EUMETSAT for distribution around the world.

ECMWF has been strongly collaborating with its ESA

Aeolus Data Innovation and Science Cluster (DISC) partners. These include the German Aerospace Centre (DLR), the software company DoRIT, the Royal Netherlands Meteorological Institute (KNMI), and Météo-France. The DISC has run a quick development cycle on the ground processing software, allowing the data to reach a level of maturity suitable for operational assimilation after 16 months. Meanwhile, the processing algorithms have continued to improve through 15 revisions so far.

Aeolus has important uses beyond NWP. In particular, atmospheric dynamics research has benefited from Aeolus winds to study gravity waves, equatorial waves, sudden stratospheric warming events and quasi-biennial oscillation monitoring, as highlighted at the third anniversary conference (<https://www.aeolus3years.org/detailed-agenda>). Aeolus winds are also proving useful to verify and hence improve other satellite-based wind observation types, such as atmospheric motion vectors.

As the first ever high-spectral resolution lidar in space, Aeolus also provides UV (ultraviolet) optical properties. These are being used in atmospheric composition research, for example on wildfire smoke, Saharan dust, volcanic eruption plumes and atmospheric composition data assimilation (see Box A). They are also being used for cloud research. The unique ability of a Doppler wind lidar to measure dynamics and optical properties should be exploited further in the future – for coupled atmospheric composition, cloud and dynamics forecasts.

a ECMWF's aerosol research activities using Aeolus

In addition to horizontal line-of-sight (HLOS) wind, the backscatter and extinction coefficients from hydrometeors and aerosols can also be retrieved from Aeolus measurements. As part of the Aeolus DISC activities at ECMWF, we assimilate the particle backscatter coefficient in research experiments using the IFS in atmospheric composition mode. This is done in a parallel setup to the one used operationally by the EU-funded Copernicus Atmosphere Monitoring Service (CAMS) run by ECMWF. This work began in late 2021, as a follow-up to previous work by Julie Letertre-Danczak and others. It aims to fully exploit the particle backscatter product of Aeolus and assess its impact on atmospheric composition forecasting.

The Aeolus L2A particle backscatter product has been included in monitoring experiments since the beginning of 2022, augmenting aerosol optical

depth (AOD) measurements in the assimilation, which provide the total aerosol extinction in an atmospheric column, but no vertical profiling. The particle backscatter coefficient is assimilated on top of the operational CAMS configuration, and we are assessing the impact of this additional information on constraining the vertical distribution of aerosols in an atmospheric column.

ECMWF is also evaluating the impact of the assimilation of Aeolus particle backscatter on short-range AOD forecasts, e.g. to improve the AOD forecast from Saharan dust events. We compare our results with the operational CAMS forecast and validate with ground-based AERONET AOD measurements. Further to this, we compare the results with ground-based lidar measurements of particle backscatter and attempt to understand the bias present in the Aeolus particle backscatter.

Reprocessed data and impact experiments

Reprocessing, which refers to the processing of satellite data using a single state-of-the-art processing chain, provides a consistent dataset with superior quality compared to near-real-time data.

Aeolus has two fully redundant laser transmitters, referred to as flight models A and B. Aeolus started with flight model A, but due to decreasing laser pulse energy a decision was made to switch to flight model B in mid-2019. This led to some natural boundaries for the choice of which periods to reprocess.

DISC and ESA planned four reprocessing campaigns. Two campaigns covering the early flight model B period (June 2019 – December 2019 and then June 2019 – October 2020) have already been accomplished, and one covering the early flight model A period (September 2018 – June 2019) is being carried out. The fourth reprocessing is intended to cover the whole mission and is planned for 2023–2024. The first and second sets of reprocessed data significantly improved the quality (particularly systematic errors) with respect to near real-time data, as assessed by ECMWF.

We are currently investigating the NWP impact of the Aeolus L2B winds from the second reprocessing, using both observing system experiments (OSE) and the Forecast Sensitivity to Observation Impact approach. These experiments represent the longest ever OSE assessment for Aeolus. The results are also being compared with a near-real-time data OSE covering late 2021 and 2022, to assess how Aeolus impact has changed.

Trends in Aeolus data quality

Despite the improvement of the ground processing

chain over time, the instrument’s spectrometer measurements have been getting noisier since mid-2019. This has been confirmed by ESA and DISC to result from a gradual reduction in the signal emitted by the Doppler wind lidar. The cause is likely due to so-called laser-induced contamination, reducing transmission through the instrument’s optics.

The declining signal is evident in the statistics of Rayleigh-clear random error estimates for the second reprocessed dataset, shown in Figure 2a. This has been restricted to the tropics to reduce solar background noise effects. It shows that the estimated one standard deviation of error (as estimated from departures from ECMWF short-range ‘background’ forecasts) increased by 40% by October 2020 compared to July 2019 (and by 60% by mid-2022, not shown). The step in the random error estimate around early November 2019 was due to an experiment involving thinner vertical range-bins aimed at maximising collocations with atmospheric motion vector retrievals (part of calibration and validation efforts).

The best atmospheric path signal levels for the mission occurred just after the switch to flight model B laser in July 2019. A stronger signal, which implies higher photon counts, results in smaller random errors for the Rayleigh-clear winds. DISC-DLR monitoring and analysis show that the atmospheric path signal in mid-2022 is only 30% of its July 2019 peak. Observation random errors, which have varied with time but are in the range 4–8 m/s for Rayleigh-clear winds, are larger than the background forecast random errors, which are typically 2 m/s in HLOS space. Therefore, improving the observation random errors translates to improved NWP impact through data assimilation.

The Mie-cloudy L2B HLOS wind random errors (Figure 2b) are less sensitive to the strength of the emitted signal. This is because the Mie backscatter from

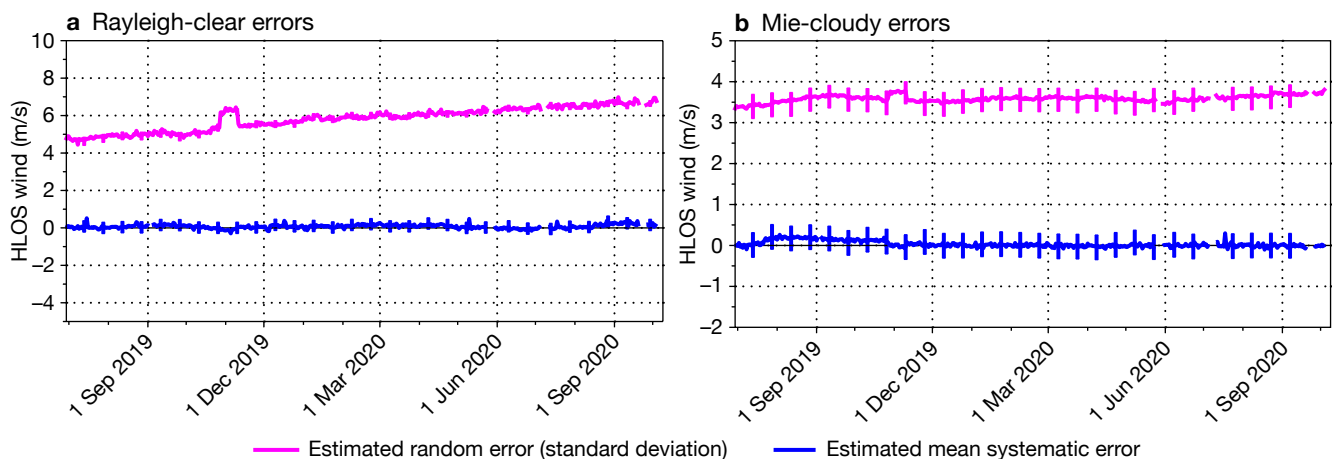


FIGURE 2 Aeolus HLOS wind error statistics derived from observation-minus-background-forecast departures, from the second reprocessed dataset for (a) Rayleigh-clear restricted to low latitudes (± 30 degrees) and (b) Mie-cloudy global. Daily statistics are given for all pressure levels.

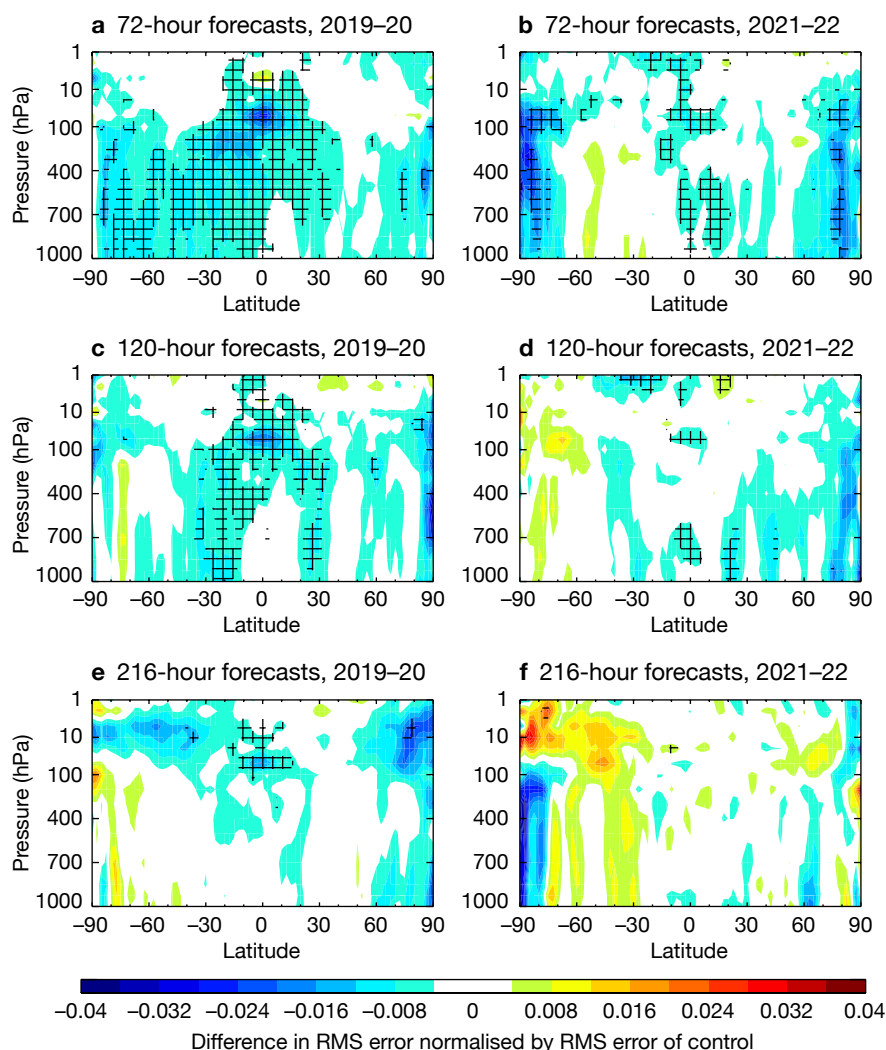


FIGURE 3 Zonal average normalised change in the RMS (root-mean-square) error of vector wind forecasts, compared to operational analyses, from assimilating Aeolus HLOS winds for the period 29 June 2019 to 9 October 2020 using the second reprocessed dataset for (a) 72-hour forecasts, (c) 120-hour forecasts and (e) 216-hour forecasts, and the same for the period 13 December 2021 to 6 August 2022, but using the near-real-time processing dataset, for (b) 72-hour forecasts, (d) 120-hour forecasts, and (f) 216-hour forecasts. Negative values indicate a reduction in error from assimilating Aeolus and positive values an increase in error. Cross-hatching indicates 95% confidence.

particles of ice/water in the clouds is sufficiently strong, whereas the Rayleigh scattering from clear-air molecules is relatively small. However, the number of Mie winds passing quality control checks slightly declines with time, as borderline low signal aerosol backscatter scenes are discarded.

Benefits to forecasts

The impact of Aeolus winds from the second reprocessing on forecast error is highlighted in Figure 3. It is shown along with the near-real-time dataset OSE covering 13 December 2021 to 6 August 2022 for comparison. This shows how Aeolus, with its best signal levels and best processing algorithms, improves the wind forecast by 0.5–2%, with a decent impact maintained into the medium range. The impact on temperature and humidity has similar patterns (not shown). Even the northern hemisphere extratropics geopotential at 500 hPa (~5 km) is improved up to day 4 by ~0.5–1% (not shown). This is despite the widespread presence of high-quality conventional data (radiosondes, aircraft). The largest impact occurs at ~100 hPa in the tropics, particularly the east Pacific Ocean (Figure 4).

This is partly because the tropics have a relatively poor coverage from high-quality radiosonde wind profiles. The wind field in the tropics is also less constrained by temperature information from other satellites. Zonal wind analyses from NWP centres disagree the most at 100 hPa in the tropical east Pacific, an area where Aeolus is making the largest improvements.

Despite the significant impact reduction in 2022 due to increasing Rayleigh-clear noise, particularly in the tropics, there is still a decent positive impact near the poles. It was noticed that the magnitude of the impact reduces as the second reprocessing OSE length was increased, again due to the Rayleigh-clear winds becoming noisier with time.

Due to uncertainty in the verification at short forecast ranges when using the analysis as the truth, it is useful to assess the background forecast fit to other observations, with and without Aeolus. Indeed, the forecast fit to nearly all observation types sensitive to wind, temperature and humidity is improved by assimilating Aeolus. However, the highest impact was found in the tropics, as exemplified using GNSS-RO

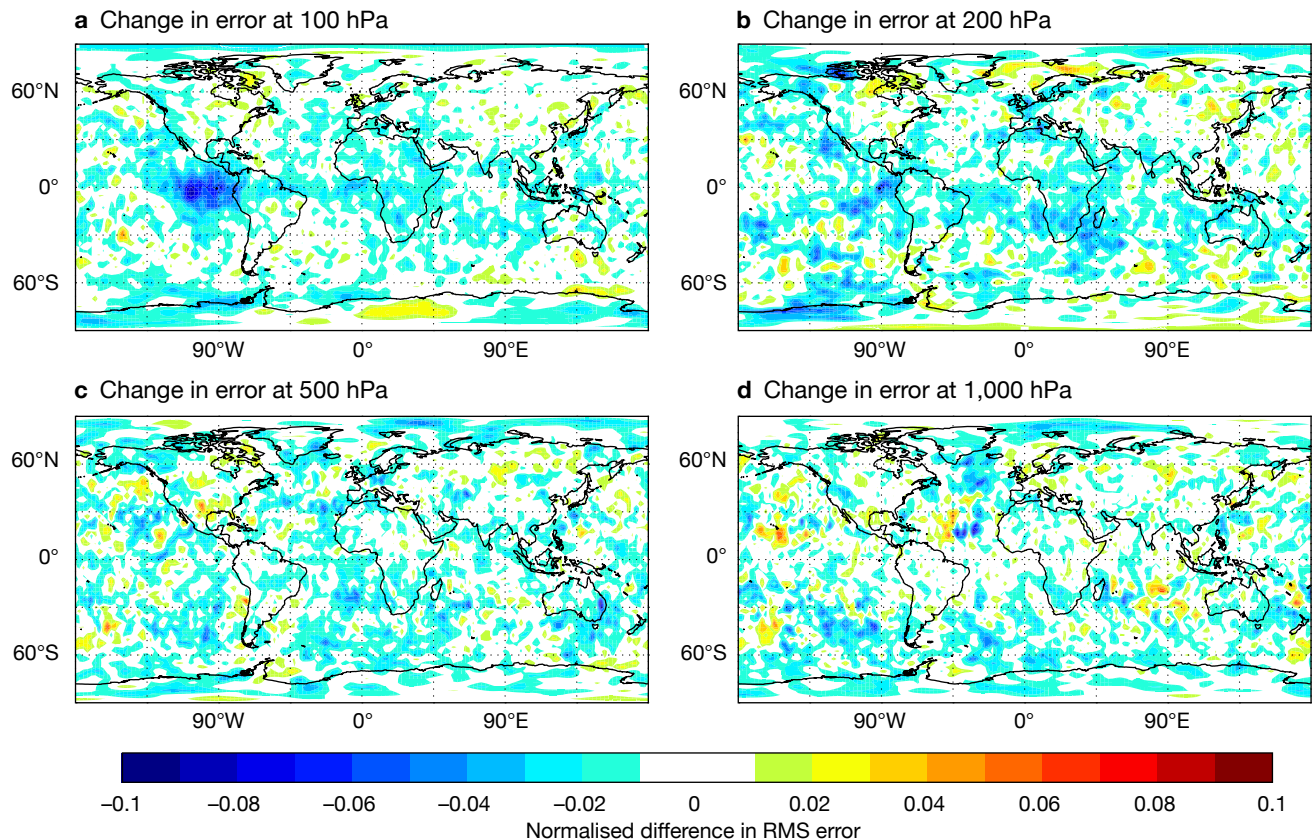


FIGURE 4 Maps of normalised change in the RMS (root-mean-square) error of vector wind forecasts, compared to operational analyses, from assimilating Aeolus (Rayleigh-clear and Mie-cloudy) at a forecast range of three days for the period 29 June 2019 to 9 October 2020 for pressures of (a) 100 hPa, (b) 200 hPa, (c) 500 hPa and (d) 1,000 hPa. Negative values indicate a reduction in error from assimilating Aeolus and positive values an increase in error.

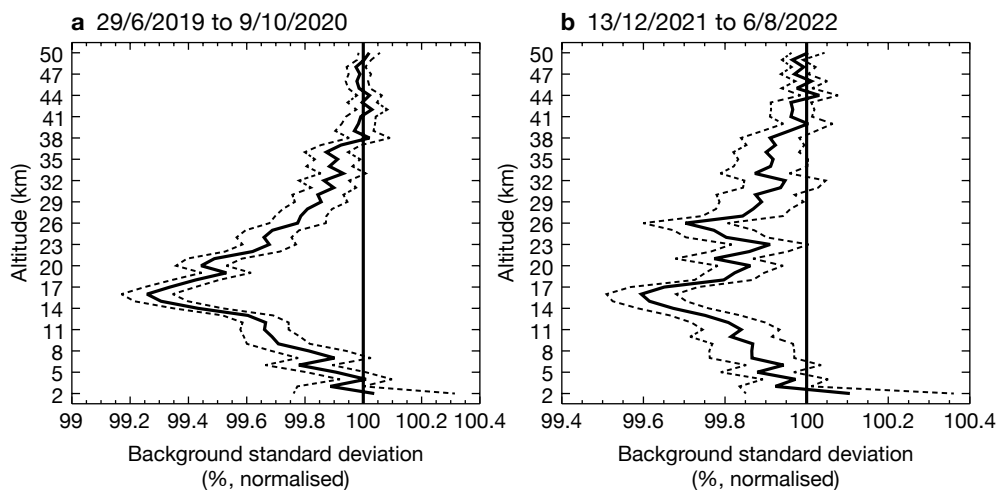


FIGURE 5 The change in standard deviation of GNSS radio occultation observation-minus-background departures in the tropics resulting from assimilating Aeolus HLOS wind observations, normalised so that the control is 100%, for (a) 29 June 2019 to 9 October 2020 and (b) 13 December 2021 to 6 August 2022. Values below (above) 100% show an improved (degraded) fit from assimilating Aeolus. The dotted lines show the 95% confidence range (Student's t-test). Note that the x-axis scales differ.

(Figure 5). The peak short-range forecast impact is found at around 15 km altitude (~150 hPa); the impact was almost three times bigger in the reprocessed dataset than in OSE using near-real-time data in late 2021/2020.

There is an unusually enhanced positive impact for the 2022 period at around 26 km altitude, which is explained by the large sample of good quality Mie-cloudy winds that Aeolus captured from backscatter off

the Hunga-Tonga eruption plume in January 2022 (Aeolus range-bins were raised to 30 km shortly after the eruption, see ECMWF Newsletter 171). The eruption plume apparently brought an unprecedented amount of water vapour to the stratosphere (Millán et al., 2022).

The second reprocessing improved the NWP impact relative to the first reprocessing campaign. This is true particularly in the lower tropical troposphere, as shown in Figure 6, thus justifying the effort on reprocessing.

The improvement seems to result from more Mie-cloudy winds being available for assimilation.

The impact of observations on forecasts can be evaluated with respect to analyses using Forecast Sensitivity to Observation Impact (FSOI). The FSOI time-series for Aeolus shown in Figure 7 is derived from the reprocessed dataset until 9 January 2020, after which it uses the operational FSOI. It concurs with OSE findings that Aeolus impact peaked in the early flight model B period (August–September 2019) due to better data quality, when Aeolus provided ~5% of the relative FSOI. This was of a similar magnitude to some of the highest-performing satellite instruments of the time, such as IASI on EUMETSAT’s Metop satellites. Aeolus

also has similar impact to the radiosonde network (not shown). The impact per observation is strong, too, coming second only to scatterometer winds among satellite-based observations.

The strong impact from one instrument suggests that assimilating more wind profile information is indeed an efficient means of improving forecast skill. Note that Aeolus signal levels were a factor of two less than expected pre-launch, even at its best during the early flight model B period, leading to significantly noisier winds than expected. Therefore, if signal levels could reach Aeolus pre-launch expectations or even better in a follow-on mission, we can expect more NWP impact.

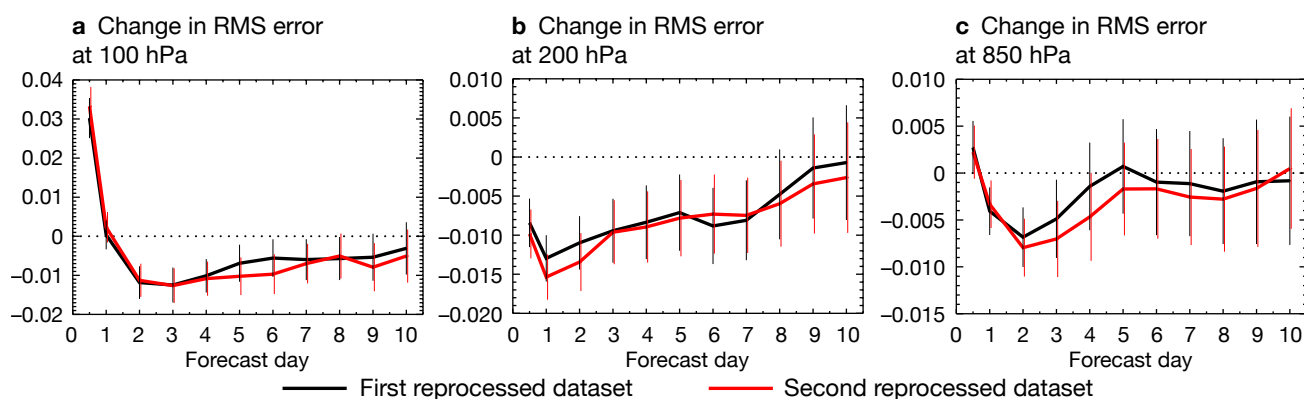


FIGURE 6 Normalised change in the RMS (root-mean-square) error of vector wind forecasts, compared to operational analyses, from assimilating Aeolus for the period 29 June 2019 to 31 December 2019 in the tropics at (a) 100 hPa, (b) 200 hPa and (c) 850 hPa. Negative values indicate a reduction in error from assimilating Aeolus and positive values an increase in error. Since operational analyses are used for verification, the first day or so is untrustworthy.

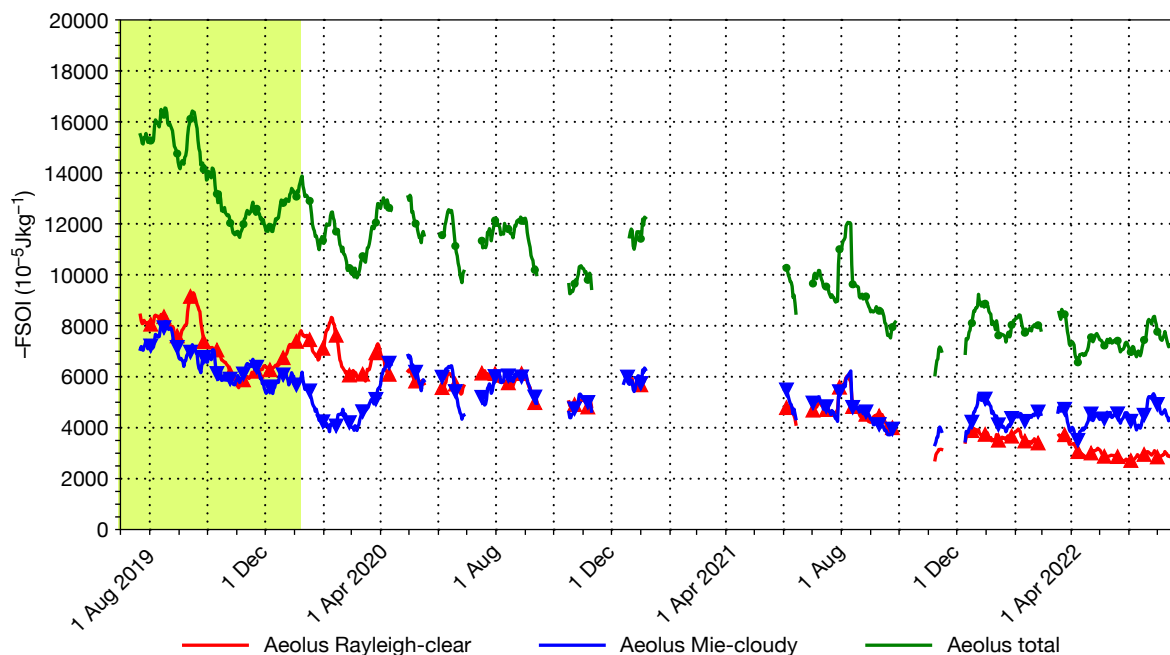


FIGURE 7 Time-series of global, daily summed FSOI for Aeolus HLOS winds, for the flight model B laser period until 21 July 2022. A two-week rolling average has been applied to reduce noise. The y-axis is the negative of the FSOI and therefore positive values represent positive impact. The second reprocessing FSOI is used before 9 January 2020 (light green background) and the operational FSOI thereafter.

Conclusions and outlook

The longest Aeolus OSE performed at ECMWF, using the second reprocessed dataset, provides an assessment of short-to-medium-range forecast impact from Aeolus L2B winds. It gave the largest positive impact seen from Aeolus so far at ECMWF, with statistically significant and good magnitude positive impact on wind, temperature, geopotential and humidity forecasts in the tropical and polar troposphere and lower stratosphere. Root-mean-square error improved by 0.5–2% up to 3–4 days ahead. The positive impact extends to 9–10 days in the lower stratosphere (between about 100 and 50 hPa or 15 and 20 km) in the tropics and the southern hemisphere extratropics. The peak in tropical impact at 100 hPa occurs in the east Pacific – a region typified by large background forecast zonal wind errors. Recent near-real-time data OSE (late 2021/2022) show that Aeolus still provides some positive impact, but noticeably lower than that of the second reprocessed OSE. The reduction in impact is particularly strong in the tropics due to larger Rayleigh-clear wind noise, which in turn is due to the gradual weakening of the emitted signal.

Aeolus impact in the short forecast range, as measured by FSOI for the first and second reprocessing, shows that Aeolus was able to provide one of the best impacts for individual satellite instruments (up to 5% in July–October 2019), similar to the radiosonde network. Operational FSOI still shows positive impact in 2022, but at ~2% it is now less than half of what it was before.

The requirements of the proposed operational follow-on mission being considered by ESA and EUMETSAT exceed those of the Aeolus mission, therefore promising more positive impact.

ESA has plans to switch back to the flight model A laser late in 2022, with the hope of boosting the signal. This

should also help to improve understanding of the transmission losses with flight model B and possibly improve the signal if the losses are restricted to the flight model B optical path. The Aeolus satellite is expected to be switched off and deorbited in the second half of 2023 due to the depletion of its fuel needed to maintain its rather low 320 km altitude orbit, especially with the current enhanced solar activity.

There is scope for further improvements in the Aeolus ground processing and assimilation algorithms, which should further improve NWP impact. This work will be pursued after satellite switch-off as part of the DISC. Such work would also benefit any operational Doppler wind lidar follow-on mission.

Further reading

Laroche, S. & J. St-James, 2022: Impact of the Aeolus Level-2B horizontal line-of-sight winds in the Environment and Climate Change Canada global forecast system. *Quarterly Journal of the Royal Meteorological Society*, **148**(745), 2047–2062. <https://doi.org/10.1002/qj.4300>.

Millán, L., M.L. Santee, A. Lambert, N.J. Livesey, F. Werner, M.J. Schwartz et al., 2022: The Hunga Tonga-Hunga Ha’apai Hydration of the Stratosphere. *Geophysical Research Letters*, **49**, e2022GL099381. <https://doi.org/10.1029/2022GL099381>.

Pourret, V., M. Šavli, J.F. Mahfouf, D. Raspaud, A. Doerenbecher, H. Bénichou et al., 2022: Operational assimilation of Aeolus winds in the Météo-France global NWP model ARPEGE. *Quarterly Journal of the Royal Meteorological Society*. <https://doi.org/10.1002/qj.4329>.

Rennie, M.P., L. Isaksen, F. Weiler, J. de Kloe, T. Kanitz, & O. Reitebuch, 2021: The impact of Aeolus wind retrievals on ECMWF global weather forecasts. *Quarterly Journal of the Royal Meteorological Society*, **147**(740), 3555–3586. <https://doi.org/10.1002/qj.4142>.

The next extended-range configuration for IFS Cycle 48r1

Frédéric Vitart, Magdalena A. Balmaseda, Laura Ferranti, Manuel Fuentes

Extended-range forecasts in ECMWF's Integrated Forecasting System (IFS), which are predictions beyond two weeks but less than a season, became operational in October 2004 to fill the gap between medium-range and seasonal forecasting. Originally, they were produced by running coupled ocean–atmosphere integrations up to day 32 at a TL159 horizontal resolution (about 110 km) every Wednesday at 00 UTC. At that time, the extended-range forecasting system was run separately from the medium-range and seasonal forecasting systems. Since 2004, the configuration of the ECMWF extended-range forecasting system has changed several times. A particularly important change took place in 2008, when the medium-range and extended-range forecasting systems were merged. Extended-range forecasts were then produced by extending medium-range forecasts up to day 32, at a lower atmospheric horizontal resolution beyond day 15. In 2011, the frequency of extended-range real-time forecasts was increased to twice a week, and in 2015 the forecast length was increased from 32 to 46 days. Since 2004, the model resolution of extended-range forecasts has also increased, from about 100 km to about 36 km and from 40 to 137 vertical levels. The current configuration of real-time extended-range forecasts is the following:

- 46-day integrations
- Every Monday and Thursday at 00 UTC
- 51-member ensemble
- Tco639 resolution (about 18 km) up to day 15 and Tco319 (about 36 km) after day 15
- 137 vertical levels.

Re-forecasts are used for calibration of real-time forecasts and for extended-range forecast skill assessment. In the current configuration, they are produced twice a week, starting the same day and month as the real-time forecasts, but for the past 20 years and with an 11-member ensemble.

The availability of the new Atos high-performance

computing facility in Bologna (Italy) provides an opportunity to enhance the ECMWF extended-range forecasting system further. The proposed changes, which will be introduced with IFS Cycle 48r1 in 2023, are the following:

- Daily instead of twice-weekly frequency
- 101 instead of 51 ensemble members
- Extended-range integrations will be run separately from higher-resolution medium-range forecasts
- Tco319 (about 36 km) horizontal resolution from day 0 to 46.

This configuration was supported by users, who expressed their feedback in a survey conducted in 2021. The enhancement in the extended-range configuration complements the strategy followed in the medium range, since it focuses on increasing the ensemble size and forecast frequency rather than increasing the model resolution. This article will discuss the main motivations behind these changes.

Increasing the frequency

Increasing the frequency of extended-range forecasts from twice weekly to daily will significantly improve the quality of extended-range forecasts for days of the week other than Monday and Thursday. For instance, the most recent extended-range forecast available on Sundays is already three days old. Figure 1 shows an example of the benefit of the increased frequency for the prediction of the Madden–Julian Oscillation (MJO) (see the blue line). It shows that, by running extended-range forecasts every day, the predictive skill of a forecast produced on Sunday is significantly higher than the predictive skill of a forecast which is three days old (current configuration). Therefore, the new configuration will provide more flexibility for users to access recently initiated extended-range forecasts.

The daily frequency will also make it possible to create lagged ensembles, where the forecasts of previous days are combined with the most recent forecasts to produce a larger ensemble. The main benefit of a lagged ensemble is that it allows the creation of a much larger ensemble. This enables more skilful forecasts overall, despite the fact that the lagged ensemble includes less

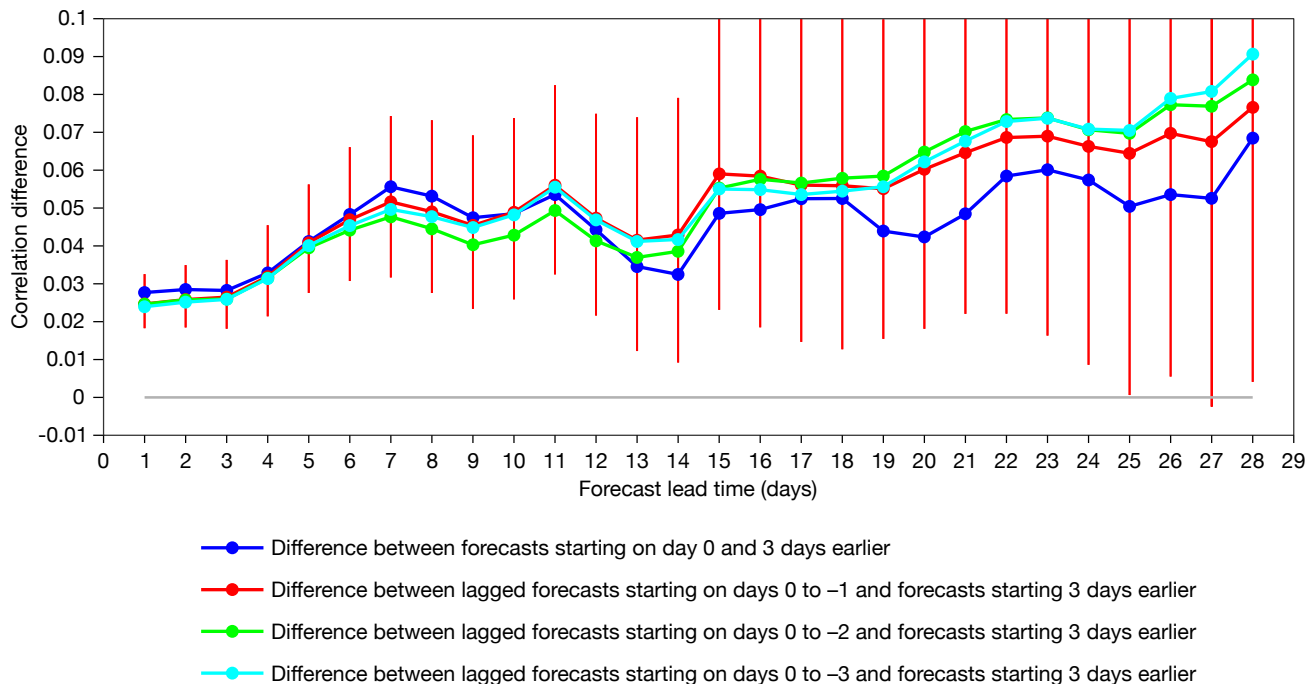


FIGURE 1 Difference of MJO bivariate correlation between 5-member ensemble re-forecasts starting on 1 November and 1 February 1989–2016 and re-forecasts starting 3 days earlier. Positive values indicate an improvement resulting from using forecasts without a 3-day delay. The vertical red lines indicate the 90% level of statistical significance. The red, green and cyan curves represent the difference when using a lagged ensemble with more members (respectively days 0 to -1; days 0 to -2 and days 0 to -3).

skilful ensemble members because they were initialised a few days earlier. However, a lagged ensemble is more difficult to calibrate because it mixes forecasts with different lead times. The benefit of this approach for extended-range forecasts has been documented in Vitart and Takaya (2021), which showed that combining the most recent extended-range forecast with older forecasts provides more skilful forecasts of weeks 2, 3 and 4 than extended-range forecasts produced only from the most recent ensemble. The optimal number of previous forecasts to combine depends on the lead time. For week 2 (days 12–18), the optimal benefits are obtained by combining just the most recent forecast with the forecast produced one day earlier, while for week-4 forecasts, the optimal lagged-ensemble configuration includes forecasts produced up to three days earlier. Including even earlier forecasts would degrade the extended-range forecast skill, despite increasing the ensemble size. Figure 1 shows an example of the benefit of a lagged ensemble for the extended-range prediction of the MJO. According to Figure 1, the benefit of the lagged ensemble (red, green and cyan curves) increases with lead time. It starts to be visible only after about 10 days of forecasts and becomes substantial in weeks 3 and 4.

Better European weather regimes

Probabilities for Euro-Atlantic regimes have been issued for some time and are freely accessible (look for ‘Weather regime probabilities’ under ‘Extended-range

charts’ in <https://www.ecmwf.int/en/forecasts/charts>).

The conceptual model of regimes is useful to condense the forecast information in ensemble forecasting (Grams et al., 2020). By identifying the dominant regime of the day in each ensemble member, we can visualise how the probabilities of regime occurrence evolve in the next 46 days. For example, users can evaluate how long strong westerlies across the Atlantic are likely to persist, or whether a transition into a different flow configuration is more likely to occur. The prediction of persistence and transition from one regime to another can be a challenge (e.g. Ferranti et al., 2015; Grams et al., 2018). On the other hand, because regime occurrence is affected by tropical heat anomalies and by fluctuations of the stratospheric polar vortex, regime predictability in the extended range can rise when such external forcings are in action. It is the sensitivity to these external forcings that makes the regimes ideal targets for extended-range predictions. Since weather regimes are linked to surface weather and weather extremes, through regime probabilities we can also infer the risk of cold spells in winter, heatwaves in summer, and widespread heavy precipitation or thunderstorm activity (e.g. Ferranti et al., 2018).

With the upgrade to Cycle 48r1, extended-range ensemble forecasts will be running daily rather than twice a week. We have evaluated the effect of this increase in forecast frequency on regime predictive skill.

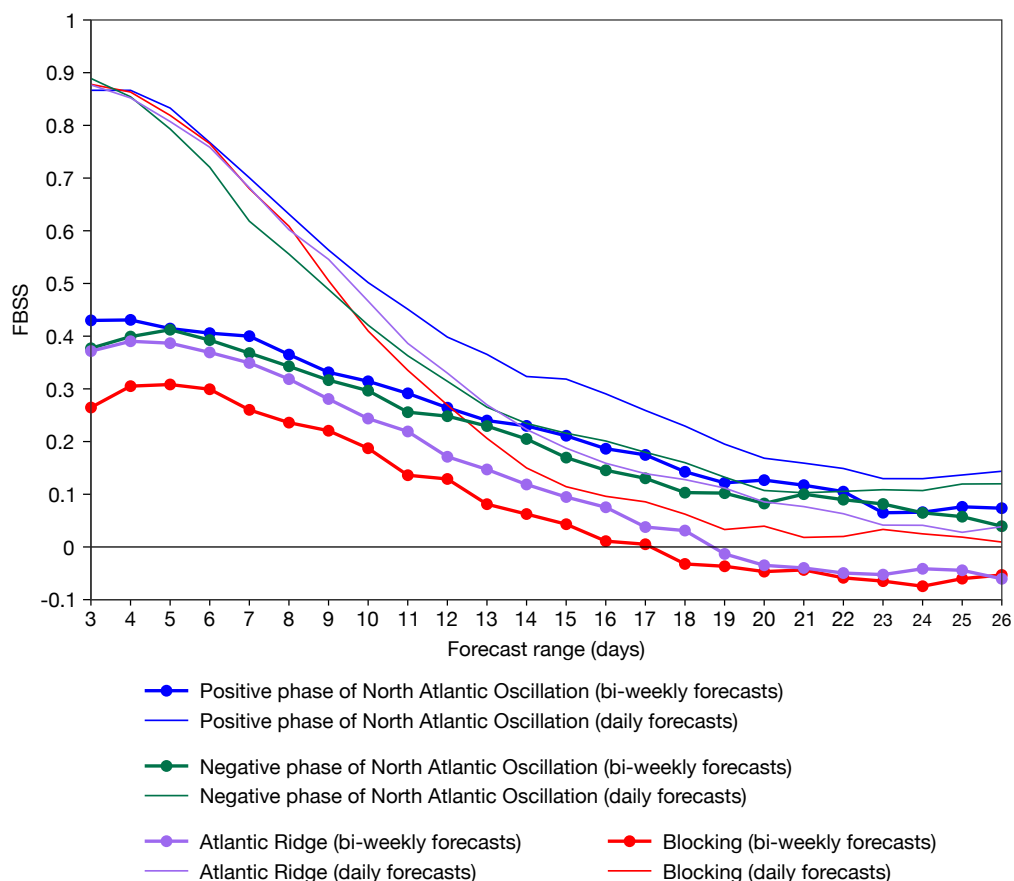


FIGURE 2 FBSS for individual regimes as a function of lead time. Thin lines correspond to the FBSS of daily forecasts, while thick lines correspond to the FBSS of bi-weekly forecasts. A regime is said to occur when the daily projections onto the regime's patterns (defined in terms of geopotential at 500 hPa) exceed the upper third of the climatological distribution. The chart is based on winter forecasts (December to February) for 20 years (2000 to 2019), with a total of about 520 bi-weekly ensemble forecasts and 1,800 daily ensemble forecasts.

Figure 2 compares the Fair Brier Skill Scores (FBSS, after Ferro, 2014) of an experimental set of daily ensemble forecasts with the FBSS of a subset of the same experimental set considering just two forecasts per week. The subset represents the current frequency configuration with forecasts issued every Monday and Thursday. The experimental set covers 20 winters (December to February), and the ensemble has 11 members. The FBSS is based on the classic Brier score with an additional correction term that compensates for skill underestimation due to a relatively small ensemble size. Figure 2 shows that the skill of forecasts issued daily is considerably larger than the skill of forecasts issued twice weekly. The benefit of daily forecasts remains evident well beyond the medium range. For example, at forecast day 15, predictive skill for the positive North Atlantic Oscillation regime increases from 0.22 to 0.35. Even for the most challenging of the regimes, the blocking regime, the skill is enhanced.

Increasing ensemble size

The 51-member ensemble in the current configuration cannot always properly capture small shifts in forecast probabilities when the signal is very small compared to the ensemble spread, which is often the case in high latitudes. Doubling the ensemble size to 101 ensemble members will provide a more accurate representation of the forecast probability distribution function. In particular,

this will ensure a more accurate prediction of the tail of the probabilistic distribution function, which is important for a more accurate prediction of the probability of extreme events. The increase in ensemble size will also help improve probabilistic skill scores, which are dependent on ensemble size (e.g. Weigel et al., 2008). Indeed, Figure 3 shows that the Continuous Ranked Probability Skill Score (CRPSS) is significantly improved at all lead times and for all parameters when increasing the ensemble size from 51 to 101 members, and the improvement is statistically significant. The ensemble size can be increased further by using a lagged ensemble approach as mentioned in the previous section. For example, a 2-day lagged ensemble would include 300 members.

Figure 4 shows an example of the impact of increasing the ensemble size to 101 ensemble members for one of the extended-range forecast products: the probability of tropical cyclone tracks within 300 km over a weekly period, in this case week 4. Figure 4 shows that the areas of high probability of tropical cyclone tracks look broadly similar between the 51- and 101-member ensembles, but there are significant differences: for instance, the 101-member ensemble forecast displays higher probabilities of a tropical cyclone strike over the Mozambique Channel and west and northeast of Australia than the 51-member ensemble. It also shows lower probabilities east of 180°E. In the event, tropical cyclones did strike in the first three locations and did not strike in the fourth.

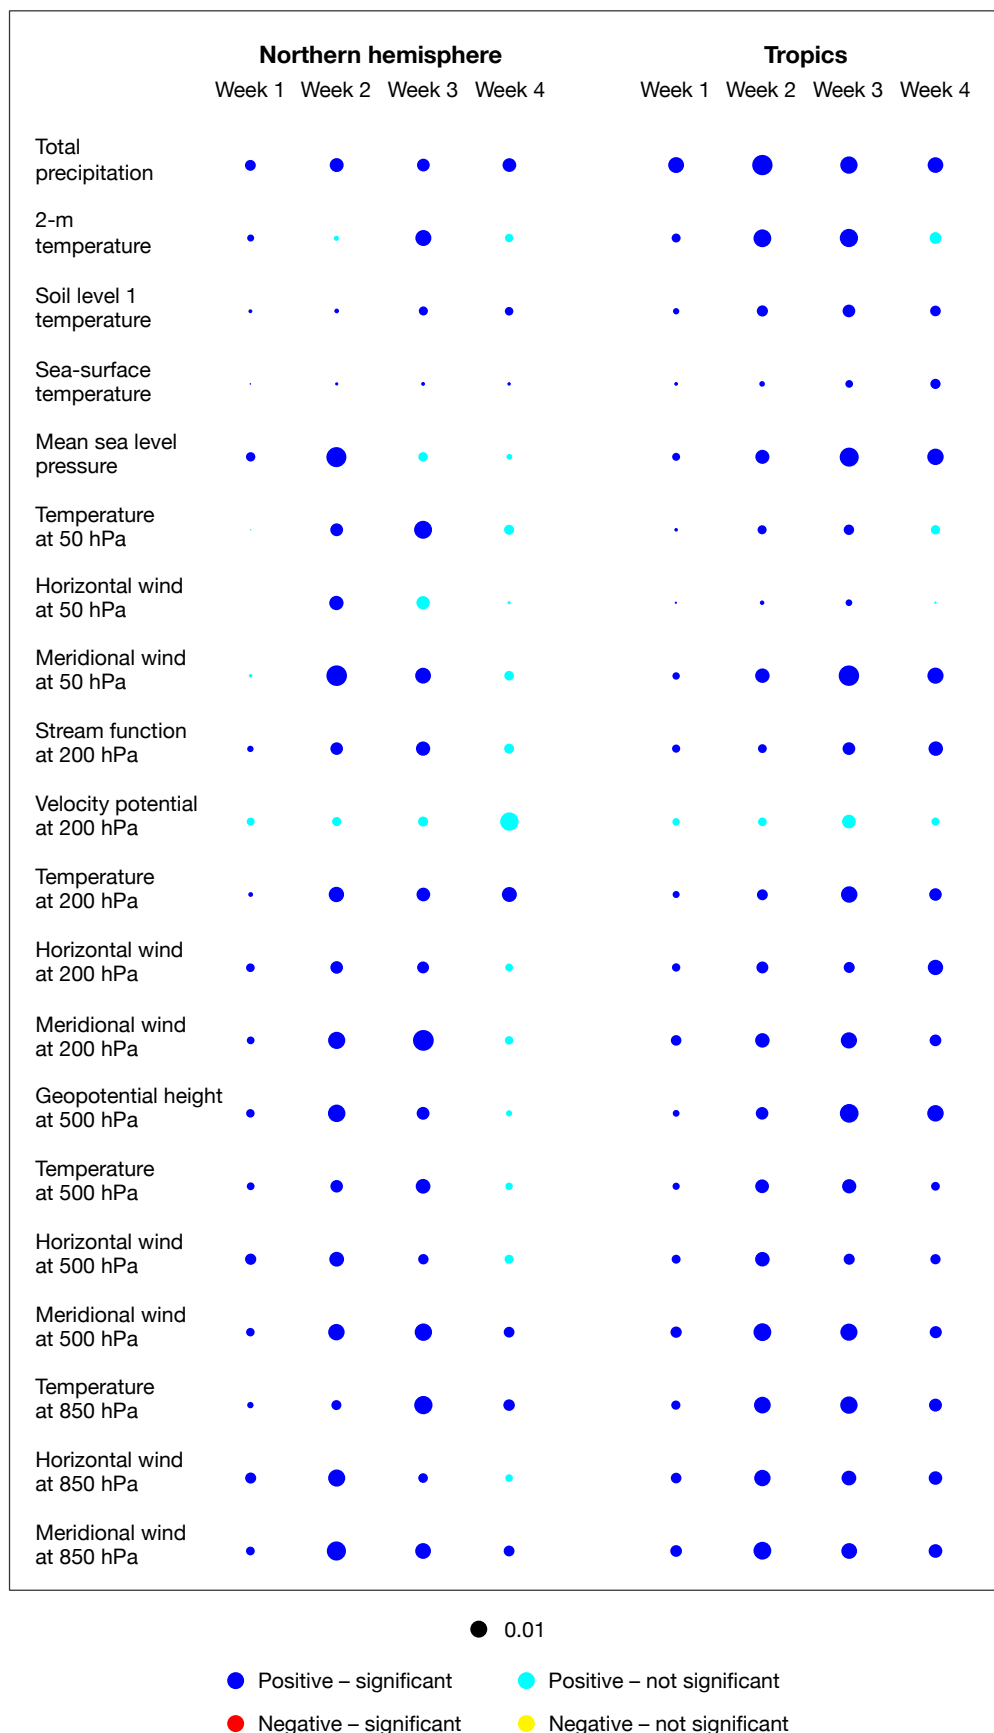


FIGURE 3 Scorecard showing the difference in the Continuous Ranked Probability Skill Score (CRPSS) between the 101- and 51-member ensembles for 20 variables, four lead times (weeks 1 to 4) and two regions (northern hemisphere on the left and tropics on the right). The blue (red) and cyan (yellow) colours indicate an improvement (a degradation), respectively. Blue or red indicate that the difference is statistically significant using a 10,000 bootstrap resampling technique. Re-forecasts were produced the first of each month over the period 1989–2016.

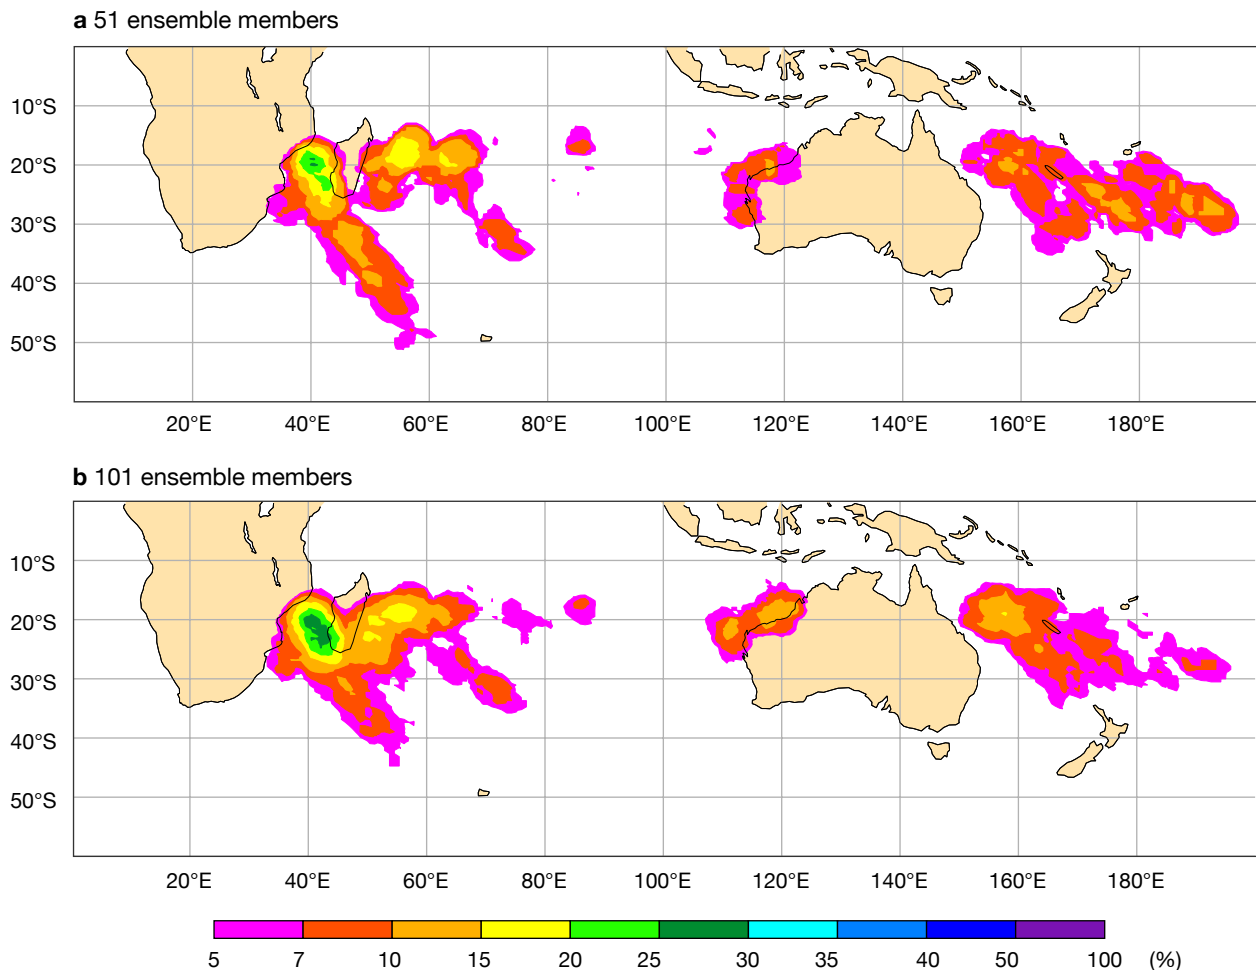


FIGURE 4 Tropical storm strike probability maps for week 4 for the extended-range forecast starting on 7 January 2021, showing (a) the strike probability produced with a 51-member ensemble and (b) the strike probability produced with a 101-member ensemble, both in the real-time forecast configuration.

The larger ensemble size will also provide a better framework for understanding the prediction of a specific extreme event, by increasing the number of possible outcomes linked to specific large-scale patterns.

The benefits associated with daily forecast frequency and with the larger ensemble size of 101 members will be visible in all products. At the time of writing, the list of products available with the cycle upgrade has not yet been finalised. Based on the testing done so far, it is likely that the weekly mean anomalies (Monday to Sunday) and probabilities for the current range of variables will be produced daily. The model climate distribution will be sampled by using three consecutive re-forecast sets, as is done currently. MJO products as well as regime probabilities and tropical cyclone products will also likely be available daily.

Decoupling extended- and medium-range ensembles

The new configuration will provide two ensembles every day at 00 UTC: a 51-member high-resolution

ensemble up to day 15, and a 101-member low-resolution ensemble up to day 46. This will make it possible to create dual-resolution medium-range large ensemble forecasts every day at 00 UTC. For some events which are not too sensitive to the model horizontal resolution, like tropical cyclone tracks, this could lead to more accurate probabilistic forecasts. Running from day 0 with the same resolution during the 46-day integration will remove all the postprocessing issues associated with the change of resolution at day 15 in the current system.

The choice of keeping the horizontal resolution at 36 km for all extended-range predictions was motivated by results from research experiments where the horizontal resolution of extended-range forecasts was increased up to 9 km. These experiments did not show any statistically significant impacts on extended-range forecast skill scores, as shown in Vitart et al. (2019), despite a more realistic representation of the climate and its extremes. This is consistent with previous studies (e.g. Vitart, 2014). They showed that the improvement in ECMWF operational extended-range

forecast skill scores since 2004 could not be associated with the changes in horizontal atmospheric resolution of the operational system. This was, for example, true of the improvement in predictions of the MJO, one of the key sources of extended-range predictability. We note that the impact of the higher resolution during the first 15 days on week 3 or 4 forecast skill scores is not significant, as shown in Figure 34 in Vitart et al., 2019.

There will thus be two sets of re-forecasts in Cycle 48r1: one at a high resolution (Tco1279) up to day 15 to calibrate the high-resolution medium-range ensemble, and one at a low resolution (Tco319) for extended-range forecasts. The main motivation for running extended-range forecasts separately from 15-day medium-range forecasts is to provide more flexibility for the real-time extended-range configuration. For instance, it allows the extended-range forecasting system to have a larger ensemble size than the medium-range ensemble forecasting system. It also allows more flexibility in the design of the re-forecasts. Currently, the cost of the re-forecasts is largely dominated by the first 15 days of integrations at 18 km resolution, while the need for large ensemble re-forecasts is not as strong for the medium range as it is for the extended range. This is due to the fact that, for medium-range forecasts, re-forecasts are used only for calibration. For extended-range forecasts, on the other hand, they are used for calibration and skill assessment, and the optimal re-forecast configuration for calibration is not the same as for skill assessment.

Although decoupled, the medium- and extended-range ensemble systems will remain seamless since, apart from the atmospheric resolution, the models used for extended-range and medium-range forecasting will be identical. This was not the case before 2008. The new dual ensemble balances the resolution enhancements in the medium range with the enhanced probabilistic aspects and frequency in the extended range.

Implementation

The new configuration will be implemented operationally with IFS Cycle 48r1 in 2023. Extended-range forecasts, which are currently archived with the same streams as the medium-range ensemble, will then be coded in GRIB and archived with different streams. These new streams are:

- Extended Ensemble Forecast Hindcast (EEFH)
- Extended Ensemble Forecast Hindcast Statistics (EEHS)
- Extended Ensemble Forecast (EEFO)
- Wave Extended Ensemble Forecast (WEEF)
- Wave Extended Ensemble Forecast Hindcast (WEEH)

- Wave Extended Ensemble Forecast Hindcast Statistics (WEES)

As mentioned above, re-forecasts will also be run separately for the medium range and extended range, since re-forecasts need to be run at the same resolution as the real-time forecasts with which they are associated. Currently, re-forecasts are run twice a week with 11 ensemble members. This will stay the same with Cycle 48r1 for both medium-range and extended-range re-forecast sets. However, it is planned to change this configuration in Cycle 49r1 since the real-time extended-range forecasts are no longer run on a twice-weekly basis. This new configuration has not been finalised at the time of writing but is likely to involve differences between the characteristics of medium-range and extended-range re-forecasts.

Further reading

Ferranti, L., S. Corti & M. Janousek, 2015: Flow-dependent verification of the ECMWF ensemble over the Euro-Atlantic sector. *Q.J.R. Meteorol. Soc.*, **141**, 916–924. <https://doi.org/10.1002/qj.2411>.

Ferranti, L., L. Magnusson, F. Vitart & D.S. Richardson, 2018: How far in advance can we predict changes in large-scale flow leading to severe cold conditions over Europe? *Q.J.R. Meteorol. Soc.*, **144**, 1788–1802. <https://doi.org/10.1002/qj.3341>.

Ferro, C.A.T., 2014: Fair scores for ensemble forecasts. *Q.J.R. Meteorol. Soc.*, **140**, 1917–1923. <https://doi.org/10.1002/qj.2270>.

Grams, C.M., L. Ferranti & L. Magnusson, 2020: How to make use of weather regimes in extended-range predictions for Europe. *ECMWF Newsletter No. 165*, 14–19. <https://doi.org/10.21957/mlk72gj183>.

Grams, C.M., L. Magnusson & E. Madonna, 2018: An atmospheric dynamics perspective on the amplification and propagation of forecast error in numerical weather prediction models: A case study. *Q.J.R. Meteorol. Soc.*, **144**, 2577–2591. <https://doi.org/10.1002/qj.3353>.

Vitart, F., 2014: Evolution of ECMWF sub-seasonal forecast skill scores. *Q.J.R. Meteorol. Soc.*, **140**, 1889–1899. <https://doi.org/10.1002/qj.2256>.

Vitart, F., M.A. Balmaseda, L. Ferranti, A. Benedetti, B. Balan-Sarajini, S. Tietsche et al., 2019: Extended-range prediction. *ECMWF Technical Memorandum No. 854*. <https://www.ecmwf.int/en/elibrary/19286-extended-range-prediction>.

Vitart, F. & Y. Takaya, 2021: Lagged ensembles in sub-seasonal predictions. *Q.J.R. Meteorol. Soc.*, **147**(739), 3227–3242. <https://doi.org/10.1002/qj.4125>.

Weigel, A., D. Baggenstos, M.A. Liniger, F. Vitart & C. Appenzeller, 2008: Probabilistic verification of monthly temperature forecasts. *Mon. Wea. Rev.*, **136**, 5162–5182. <https://doi.org/10.1175/2008MWR2551.1>.

A new way of computing semi-Lagrangian advection in the IFS

Michail Diamantakis, Filip Váňa

Modern weather and climate prediction models, such as the one used in ECMWF's Integrated Forecasting System (IFS), are complex multi-component software systems simulating Earth system processes. At the heart of each Earth system model lies its dynamical core. Its purpose is to solve the equations that govern atmospheric motions. Those equations link all atmospheric dynamical and physical processes which provide 'source terms', in other words the 'fuel' that generates all types of atmospheric motions. Here we present a new, efficient and simple iterative algorithm to compute semi-Lagrangian advection, which deals with the transport of heat, momentum, moisture and atmospheric constituents. It will be introduced in the next upgrade of the IFS planned for 2023.

Advection in the IFS

Although the equations of motion of the dynamical core are continuous, only discrete forms can be solved on a computer. The solution starts from a set of initial conditions, the so-called analysis, and evolves step by step until the final forecast time. Sophisticated numerical methods must be used, which stem from rigorous mathematical analysis, to ensure accurate and stable calculations. At the same time, considerable effort is spent to engineer such methods to work fast on modern supercomputing architectures.

A crucial part of each dynamical core is the advection scheme. This is the part of the model time-stepping scheme which calculates the transport of heat, momentum, moisture and atmospheric constituents around the globe, given the winds at each point in time. The ECMWF dynamical core uses a semi-Lagrangian (SL) advection scheme, which ensures that an accurate solution, with very little numerical dispersion, can be produced stably with long time steps. The latter means that it is computationally efficient, as we can use fewer steps to produce a given forecast. It is also multi-tracer efficient, which means that, when transporting many atmospheric species (gas tracers), its overhead is small compared with other methods.

The SL scheme is combined with a semi-implicit scheme, ensuring that all dynamical core equations are

solved stably and accurately using long time steps. The introduction of the SL technique in the IFS in 1991 was a major enhancement of its capability to deliver efficiently successive resolution upgrades and improvements. This is an old scheme, but because of its success in global atmospheric modelling, it has been continuously developed, enhanced, and improved to match ever-evolving and challenging modelling and computer hardware requirements. What is currently used is an evolution of the original scheme, enhanced to allow accurate and computationally efficient predictions for modern high-resolution requirements and to support atmospheric composition modelling.

With increasing resolution and model complexity and the introduction of reduced precision to speed up the forecast model, we found that components of the advection scheme had become slow and that they required special treatment to preserve the required accuracy. For this reason, part of the advection scheme has been rewritten, using a mathematical formulation that leads to a faster but equally accurate solution.

The semi-Lagrangian advection scheme

Atmospheric motions and energy exchanges alter the transported fields (moisture, air density, temperature, momentum etc.). Any change incurred is the outcome of different adiabatic and diabatic processes (pressure differences, rotation, diffusion, heating, evaporation/condensation, friction etc.), which are calculated separately during each model time step. However, advection is a passive process which moves atmospheric substances (fields) without altering their nature at all. SL advection achieves combining effectively the Lagrangian method with the fixed grid required by the physics. To do this, it uses the concept of backward trajectories: for each model grid point, it must trace backwards in space and time the origin of the Lagrangian trajectories formed by air parcels as they are advected (passively transported) by the wind. Calculation of backward trajectories implies that, at each time step, each parcel will arrive at a grid point, which is also called 'arrival point'. The origin of each trajectory, called the 'departure point', in general lies between grid points. It is computed solving the so-called trajectory differential equation $d\mathbf{r}/dt = \mathbf{V}(\mathbf{r},t)$, where \mathbf{r} denotes the coordinates of an air parcel and $\mathbf{V}(\mathbf{r},t)$ the wind field acting on it, which is variable in space and time.

Hence, when restricted to advection, the time evolution of a transported variable on each grid point is determined by two steps in a SL method: the first step is to find the Lagrangian departure points, and the second is to estimate the value of the transported variable at these locations. The latter is done by interpolating neighbouring grid point values. Having good accuracy in both computational steps is essential for an accurate forecast. Inaccurate departure point calculations in the vertical will generate large temperature errors and negatively impact the flow, while inaccurate horizontal point calculations will introduce phase errors in Rossby waves and forecast the wrong location of cyclones. Inaccurate interpolation will generate substantial amplitude errors in all weather features and will poorly forecast the intensity of weather phenomena.

In this article, we will concentrate on the departure point calculations. These are computed by solving a discretised version of the trajectory equation. This discretised equation is 'implicit' in nature (the unknown departure point appears in the left- and the right-hand side of the equation). It can be solved with a mathematical technique called 'fixed point iteration'. This is an iterative method which produces approximations to the solution. The number of iterations needed for the algorithm to converge to the solution varies considerably in different regions of the atmosphere. In areas where the wind flow is nearly uniform with weak gradients, it converges rapidly. Here, one or two iterations are enough. However, near the jet stream or tropical cyclones, where the horizontal and vertical wind shear is large, convergence is slow, and more iterations are needed. The time step length affects the rate of convergence. The quantity that determines this rate is the Lipschitz number, which is the product of time step and the norm of the wind shear.

The iterative scheme used in the IFS is called SETTLS (Stable Extrapolation Two-Time-Level Scheme). Although this looks like a simple algorithm when the equations are expressed on a Cartesian framework, its simplicity is lost when spherical polar coordinates are used. For example, curvature effects need to be accounted for, and to obtain the latitude and the longitude of each departure point, the inversion of a 3x3 nonlinear system of equations is required. This adds computational expense and complexity in the code such that the final scheme seems to bear little resemblance to the original SETTLS. Furthermore, when the departure point calculations are carried out in reduced precision ('single-precision' rather than 'double-precision' was introduced recently in the IFS to speed up execution), near the poles the accuracy is reduced. It was therefore necessary to keep big parts of this calculation at high precision.

A new algorithm for departure points

Over the last few years, external funding from the EU-funded ESCAPE project helped scientists at ECMWF, in collaboration with external partners from several Member States and the EU, to study in detail the computational performance of fundamental algorithms of the IFS dynamical core as well as physical parametrizations, such as the cloud scheme. Several self-contained software modules were produced, the so-called dwarfs. One of these dwarfs focused on semi-Lagrangian advection. The same method, based on SETTLS, was replicated in the dwarf, but it was implemented differently. Modern data structure and object-oriented programming concepts that were available in ECMWF's Atlas library were used. In practice, this meant that it was easier to implement and test different algorithms for semi-Lagrangian advection.

Among the alternatives explored were iterative algorithms for the departure points, which were computing the semi-Lagrangian trajectories using an equivalent geocentric Cartesian framework, mitigating singularity problems associated with spherical polar coordinates in the polar regions. At the same time, alternative non-iterative algorithms were explored that were cheaper in terms of computation and communication. However, testing suggested that the iterative approach was essential for accuracy. Therefore, a Cartesian iterative algorithm based on SETTLS, which was also simpler, was ported into the operational IFS code environment to be explored further in detail.

At the same time, there were parallel efforts at ECMWF. The plan was to remove the polar singularity existing in spherical coordinates and to improve tangent-linear and adjoint integrations used in the data assimilation system. This was to be done by re-designing parts of the departure point calculations to work with Cartesian coordinates. Such ideas matured after discussions with an ECMWF Member and Co-operating State scientist (Jan Mašek, Czech Hydrometeorological Institute). An option was added in the IFS to compute the departure points with parts of these calculations in a Cartesian coordinate framework. More specifically, the wind vectors were transformed into Cartesian coordinates before they were interpolated at the departure point estimates at each iteration of the SETTLS scheme. This amended algorithm, which was fully blended with the existing scheme, still had parts computed directly in spherical polar coordinates. It was making the overall calculation slightly more complex and expensive. However, it motivated further developments and fusion with the ideas that were discussed in the previous paragraphs.

Having the earlier experience from the Atlas-based preliminary implementations and experiments, it was

clear that a further big step could be taken to simplify and reduce the cost of the computations of SETTLS by rewriting the entire algorithm in Cartesian coordinates. An additional advantage was that the resulting iterative scheme was ‘rotation matrix-free’, unlike the algorithm which it replaced. The latter requires that the wind vectors, after being interpolated at the departure point during iterations, are multiplied by a rotation matrix to account for the change in orientation of the local coordinate system that describes them due to the Earth’s curvature. Furthermore, the calculation of latitude–longitude of the departure points was greatly simplified, no longer needing the inversion of nonlinear system of equations. The improved arithmetic scaling due to measuring distances from the centre of the Earth also meant that there was no need to approximate any terms or to keep any parts of the relevant code in double precision in single-precision runs. Finally, the original departure point calculation code was squeezed to less than 50%, and the associated tangent-linear and adjoint versions of the nonlinear model became much simpler. These are greatly beneficial features from the efficiency and maintenance point of view.

Accelerating the departure point calculation

The Cartesian coordinate algorithm for computing the departure points described so far is much more concise than its predecessor. Nevertheless, it is not faster than the default scheme of the IFS. The reason is that, within a Cartesian geocentric system (see Figure 1), movement on a 2-dimensional spherical surface takes place in 3-dimensional (X, Y, Z) space and requires three wind components to be described (U, V, W). These can be obtained from the wind components (u, v) in spherical coordinates by a simple transformation (see Table 1). On balance, given that the new Cartesian algorithm was simpler, despite needing an extra dimension its cost was similar compared with the original scheme. Could we accelerate these calculations further? Indeed, it turned out to be possible, but before explaining how this was done, we need to highlight some further aspects of these calculations and their dependency with the model time step.

We mentioned earlier that the quantity that determines how many SETTLS iterations are needed to compute the departure points is the Lipschitz number, i.e. the product of time step and the norm of the wind shear. At each resolution upgrade, the model time step is carefully selected to give the highest possible forecast accuracy consuming the least possible high-performance computing resources in the fixed time-critical window of ECMWF operations. For accuracy reasons, the ratio of the time step to grid-spacing must remain roughly the same as the model resolution increases. However, testing that took place before the

upgrade to 9 km horizontal resolution in high-resolution forecasts (HRES) in March 2016 showed undesirable side effects: unphysical stratospheric flow structures were seen over tropical cyclones, as well as large areas of increased upper-air field errors in the winter season above big mountain ranges. The origin of such skill degradations was found in the departure point calculation and the fact that, with successive model improvements, for efficiency reasons, we had pushed the time step to values that resulted in higher Lipschitz numbers. With long time steps, more iterations were needed, and thus iterations were increased from three to five. This increased the cost of the advection scheme but reduced the overall model cost as it allowed a 12.5% longer time step.

Although increasing the number of iterations to five resulted in a net gain in efficiency, the departure point calculation became an expensive part of the model,

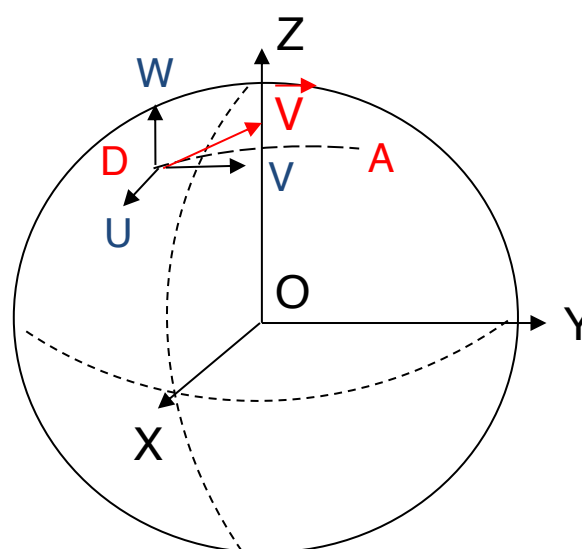


FIGURE 1 Schematic of a semi-Lagrangian trajectory on the sphere from the departure point D to the arrival (grid) point A. The red vector, which is tangential to the sphere, represents movement on the surface of the sphere in a spherical coordinate system, i.e. $\mathbf{V} = \mathbf{V}(t, \lambda, \theta)$, where (λ, θ) is the longitude–latitude and t the time. The same movement can be described using a geocentric Cartesian coordinate system (X, Y, Z) where O is the centre of the Earth and (U, V, W) are the corresponding coordinates of \mathbf{V} .

$X = a \cos \theta \cos \lambda$	$U = -u \sin \lambda - v \cos \lambda \sin \theta$
$Y = a \cos \theta \sin \lambda$	$V = u \cos \lambda - v \sin \lambda \sin \theta$
$Z = a \sin \theta$	$W = v \cos \theta$

TABLE 1 A simple transformation, suitable for a model using the shallow atmosphere approximation, such as the IFS, to obtain the Cartesian coordinates of a sphere from corresponding longitude and latitude coordinates (λ, θ) and to transform a wind vector (u, v) from polar spherical to Cartesian coordinates.

roughly taking up 7% of total CPU time. There were attempts to reduce this cost, for example by combining SETTLS in the vertical with a very economical scheme for computing the longitude and latitude of the departure points, invented by John McGregor (Commonwealth Scientific and Industrial Research Organisation, Australia). In the end, the most practical and economical idea was born in a coffee break with a member of ECMWF's Scientific Advisory Committee, Prof. Eigil Kaas, in 2019. After reviewing options, his suggestion was to test initialising the departure point iterations using the previous time step departure points rather than using the explicit Euler predictor to produce these. Indeed, after some special treatment for the vertical part of the calculation was added, this idea worked very well and was adopted (for details, see Diamantakis & Váňa, 2022).

Verifying the new algorithm

The good accuracy of this new, faster algorithm can be demonstrated with a simple forecast experiment. Figure 2 shows a plot of the 500 hPa geopotential height root-mean-square error (RMSE) difference of 30 forecasts with the old SETTLS implementation in polar spherical coordinates with three iterations from the old operational SETTLS with 5 iterations (left) and of the new SETTLS implementation in Cartesian coordinates with three iterations from the old SETTLS with 5 iterations (right). We notice that a significant increase in RMSE appears over the Himalayas when the number of iterations is reduced in the old scheme, which is completely absent in the new scheme: it can achieve the same accuracy with a 40% reduction in the number of iterations, saving approximately 2–3% of total CPU model time.

There was further detailed testing of this new development in the IFS over different seasons, cycles, forecast tests and resolutions, ranging from standard 10-day forecasts to long climate type runs and from 130 km to 2.5 km resolution. All verification tests confirmed that reducing the number of iterations in the new scheme has no impact on accuracy, with a hint of a slight improvement in cases of strong cross-polar flow. An interesting comparison is shown in Figure 3, where the mean absolute position error and the minimum pressure error (core pressure) are compared for two different sets of consecutive forecasts of the Atlantic 2020 hurricane season, verified against Best Track observations. The first set (cyan line) is produced by standard IFS forecasts in double precision using the old SETTLS scheme, which uses polar spherical coordinates and five iterations. The second set (orange line) is produced by an IFS version which uses the new Cartesian coordinate SETTLS with three iterations, with all model computations in single precision. The result demonstrates the overall quality of the entire single-precision development and the new development: with significantly less computation, the same accuracy is achieved.

Conclusion and further work

In this article, we discussed an alternative formulation for a crucial part of the IFS advection scheme, the algorithm that computes the departure point of the semi-Lagrangian trajectories. This alternative implementation, which will be introduced in IFS Cycle 48r1, brings some important benefits. It needs fewer iterations to find the solution, hence it requires less computation, but it is equally accurate. It is a simpler algorithm that is better suited for reduced-precision calculations as well as for calculations near the poles.

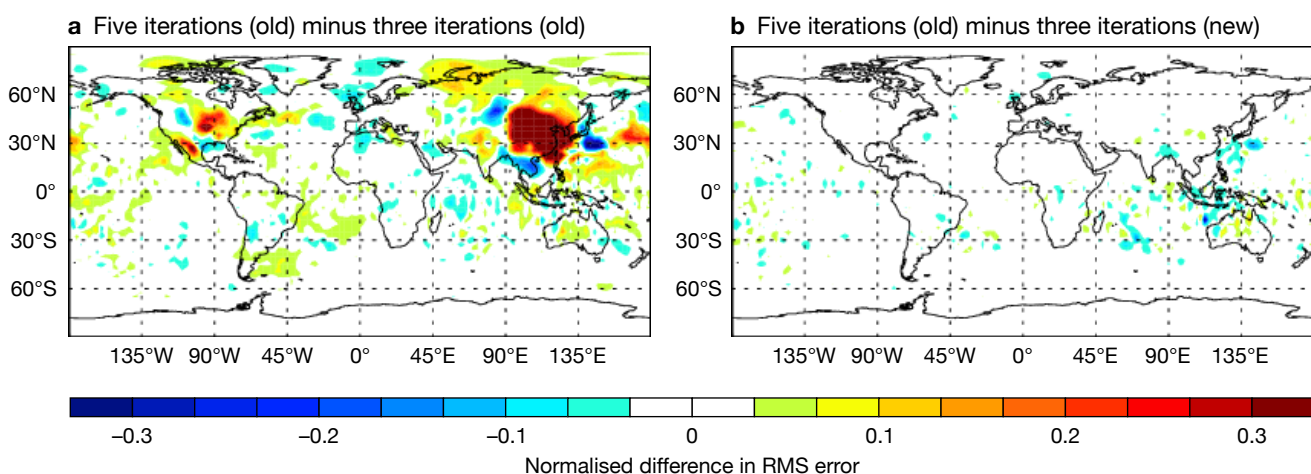


FIGURE 2 The red area in (a) indicates an increase in root-mean-square (RMS) error exceeding 30% downstream of the Himalaya region in the current operational scheme when advection scheme iterations are reduced from five to three. The plotted field is the 500 hPa geopotential height average difference in error at t+48 hrs for January 2020 forecasts. In (b), the same difference from the operational model is plotted, but with the new more efficient version of the advection scheme, which shows that it can produce the same result without loss of skill using three iterations only.

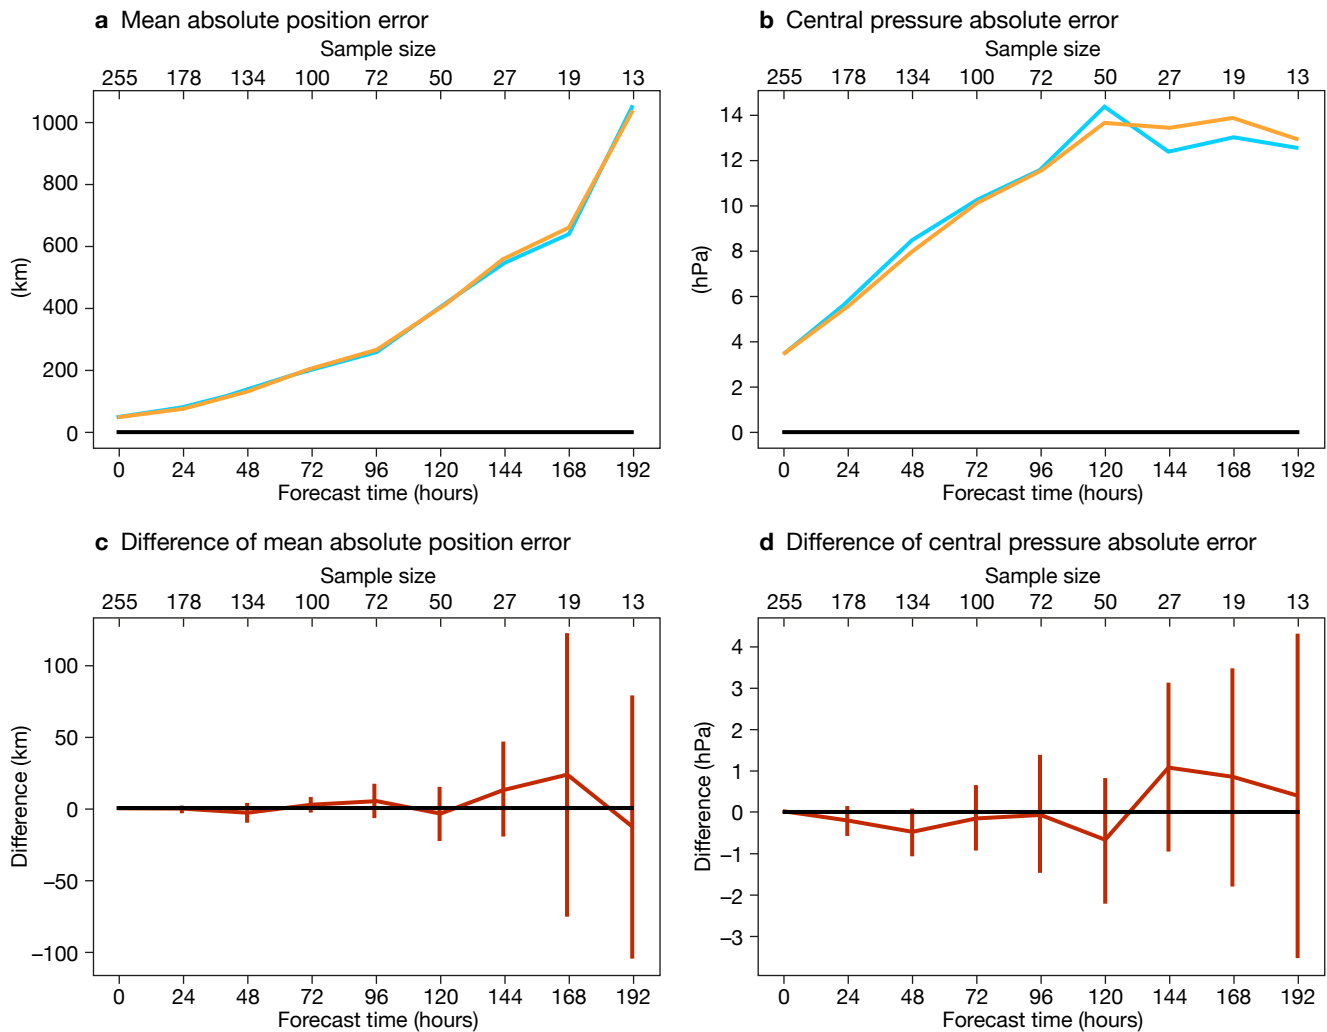


FIGURE 3 Comparing the new Cartesian faster algorithm (orange lines) for computing the semi-Lagrangian trajectory departure points against the original scheme (cyan lines) with five iterations in polar spherical coordinates in the tropical cyclone period of 2020 for (a) the mean absolute position error of tropical cyclones and (b) the central pressure absolute error of tropical cyclones. To show the accuracy of the new approach, the IFS forecasts with the original scheme are in double precision while the forecasts with the new scheme are entirely in single precision. Despite these substantial changes (resulting in an increase of approximately 40% of computational efficiency), equally accurate predictions are obtained. This is shown by the difference between the orange and blue lines for (c) the mean absolute position error and (d) the central pressure absolute error. The vertical bars indicate the range of statistical uncertainty.

This new development could widen the possibilities for future improvements. For example, it could aid efforts to revise and optimise the parallel communication strategy for the semi-Lagrangian advection scheme. This could be done by implementing alternative ‘thin halo’ communication on-demand approaches, which will reduce the memory and communication cost in future high-resolution, massively parallel weather forecast simulations. Furthermore, it could facilitate work to develop a stochastic departure point framework. The aim is to capture flow-dependent uncertainty from the dynamics in regions where it is known that the numerical algorithms can produce larger errors due to the complexity of the flow. The simplicity of the new scheme offers an advantage in implementing such approaches.

Further reading

Diamantakis, M. & F. Váňa, 2022: A fast converging and concise algorithm for computing the departure points in semi-Lagrangian weather and climate models. *Q. J. R. Meteorol. Soc.*, **148**, 670–684.

Diamantakis, M. & L. Magnusson, 2016: Sensitivity of the ECMWF model to semi-Lagrangian departure point iterations. *Mon. Wea. Rev.*, **144**, 3233–3250.

M. Hortal, 2002: The development and testing of a new two-time-level semi-Lagrangian scheme (SETTLS) in the ECMWF forecast model. *Q. J. R. Meteorol. Soc.*, **128**, 1671–1687.

Using genetic algorithms to tune the Copernicus Climate and Atmosphere Data Stores

Baudouin Raoult

On 23 March 2022, the Copernicus Climate Data Store (CDS) and Atmosphere Data Store (ADS) were successfully migrated from Reading (UK) to ECMWF's new data centre in Bologna (Italy). The CDS and ADS are part of the EU-funded Copernicus Climate Change Service (C3S) and Copernicus Atmosphere Monitoring Service (CAMS) run by ECMWF. They are deployed on a private, on-premises cloud infrastructure, running OpenStack. The service is composed of over 250 virtual machines, running on 30 physical servers, each having 48 CPUs and around 190 gigabytes of memory. This article describes how genetic algorithms are used to find an optimal placement of virtual machines onto physical hosts in order to optimise the service offered by the CDS/ADS. This is an example of the increasing use of machine learning at the Centre, in this case to optimise the performance of a service.

Function of virtual machines

Some of the virtual machines serve the CDS/ADS catalogue and its search engine, others manage the

queue of incoming user requests and broker them to the right data source or compute service. Virtual machines also process the requests and handle the download of results back to users. In addition, tens of virtual machines host a disk-only MARS (Meteorological Archival and Retrieval System) server, serving two petabytes of selected ERA5 climate reanalysis data and CAMS data. All these services communicate with each other over a virtual network, and they may be CPU, I/O or network intensive, depending on their purpose.

Every virtual machine has some memory allocated to it, as well as a number of virtual CPUs, virtual disks and virtual network interfaces. OpenStack will not run more virtual machines on the same physical host than the number that fits in the physical memory available. On the other hand, it will allow for more virtual CPUs to be allocated on the same host than the number of available physical CPUs (by a factor of two). This enables a better use of resources as many virtual machines are mostly idle, waiting for work to do. But when some of the services are very busy, virtual machines that share the same host will 'steal' CPU resources from each other.

a Genetic algorithms

Genetic algorithms (GAs) simulate evolution to find the best solution to a problem. Solutions are represented by a vector of symbols, usually bits, floating-point numbers, integer numbers or characters. The vectors are referred to as 'chromosomes' and their elements are referred to as 'genes'.

The algorithm starts by generating a large number of random solutions (known as the 'population'). For each solution (also called 'individuals'), a 'fitness' is computed; the higher the fitness, the closer it is to the best solution to the problem. A new population is then created by selecting the fittest individuals ('parents'); new individuals are created by 'mating' the parents, i.e. by creating new solutions whose genes are a combination of the genes of the parents,

in a process called 'cross-over'. Finally, a few genes of each individual are randomly changed, in a process called 'mutation'; this allows solutions to be considered that were not in the population. The process is then repeated for several 'generations'. The fittest individual of the last iteration is considered to be the best solution.

Even though GAs explore the space of solutions randomly, they often provide useful results, given a large enough population and a large number of generations, as well as carefully selected cross-over and mutation strategies. GAs are a useful tool when estimating the fitness of a solution is easy to implement. For more details on genetic algorithms, see: en.wikipedia.org/wiki/Genetic_algorithm.

It should also be noted that, when a virtual machine reads from or writes to a virtual disk, this results in network transfers between the compute hosts and the storage hosts that compose an OpenStack installation. This means that all networking and I/O performed by the virtual machines of a physical host will make use of the network interfaces of that host. On the other hand, network transfers between the two collocated virtual machines will be very fast, as the data remains within the same host and no actual networking is taking place.

As a result, to ensure a performant system, virtual machines that are running services that consume a lot of CPUs and I/O should be kept apart, while those that transfer large amounts of data between themselves should be collocated on the same physical hosts.

Figure 1 illustrates a situation where virtual machines are placed on physical hosts more or less randomly. Given the following constraints:

- virtual machines of the same colours should not be placed on the same host
- purple and pink virtual machines must be collocated pairwise
- resource usage must be evenly distributed amongst hosts,

a genetic algorithm should output a solution such as the one given in Figure 2.

Implementation

OpenStack provides a command-line tool to query the state of the system and perform actions on it. With that tool, it is possible to build a map of the system, finding

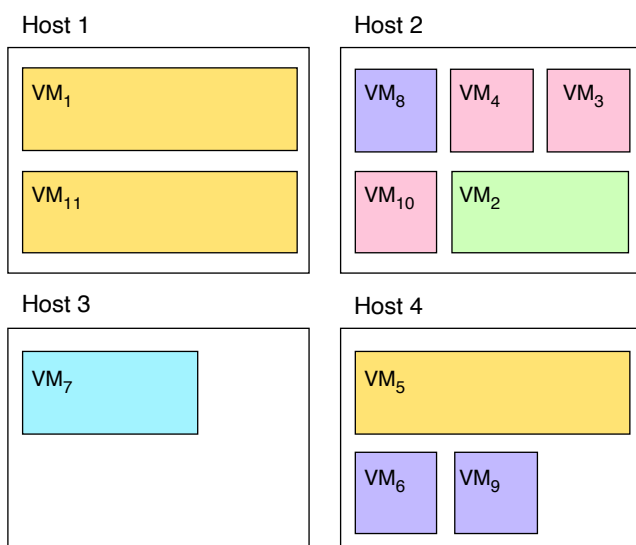


FIGURE 1 Initial placement of virtual machines (VMs) on physical hosts.

out which virtual machines run on which physical hosts, how much memory and CPU they use, etc. Furthermore, the tool also allows an administrator to migrate a virtual machine from one physical host to another, without interrupting the service.

Chromosome encoding

The tool is first used to establish a list of physical hosts as well as characteristics such as the number of CPUs and available memory. Then the current distribution of virtual machines and their resource requirements is mapped, also in terms of CPUs and memory.

The resulting information is encoded in a ‘chromosome’ as follows: the physical hosts are numbered from 1 to N_{hosts} ; similarly, the virtual machines are numbered from 1 to N_{vm} ; we use a vector to represent a possible placement of virtual machines onto hosts: the vector has N_{vm} elements, one per virtual machine, each representing on which physical host the corresponding virtual machine is located (Figure 3).

Using that encoding, the placement in Figure 1 would be [1,2,2,2,4,4,3,2,4,2,1] and the placement in Figure 2 would be [1,2,4,1,4,4,3,1,2,2,3].

Affinity

We define the ‘affinity’ between two virtual machines as the fact of whether they should be located on the same hardware or on different hosts. This information is based on the virtual machine names and is provided in a configuration file.

Fitness function

To compute the fitness of a possible placement, we consider which of the following ‘hard constraints’ are violated:

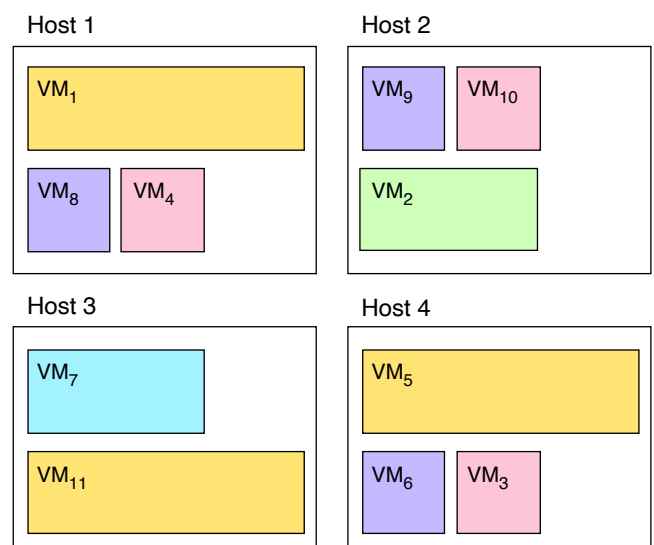


FIGURE 2 Optimal placement of virtual machines (VMs) on physical hosts.

VM ₁	VM ₂	VM ₃	VM ₄	VM ₅	VM ₆	...	VM ₁₉₆	VM ₁₉₇	VM ₁₉₈	VM ₁₉₉	VM ₂₀₀
1	3	1	20	23	3	...	16	9	32	17	1

FIGURE 3 The structure of a chromosome, which has one element per virtual machine (VM), each representing on which physical host the corresponding virtual machine is located.

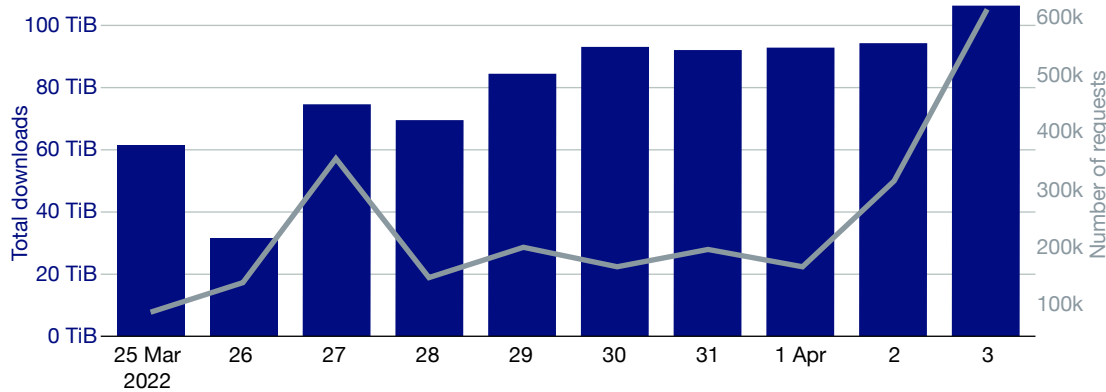


FIGURE 4 A record 106 TiB of data was delivered by the CDS on 3 April 2022.

- the sum of the memory requirement of all virtual machines sharing the same host must not exceed the amount of physical memory of that host
- two virtual machines that should not be collocated must not be on the same host
- two virtual machines that should be collocated must not be on different hosts.

We also consider the following ‘soft constraints’:

- the standard deviation of the number of allocated virtual CPUs on all hosts must be minimal (this ensures that the CPU load is evenly spread within OpenStack)
- the number of virtual machines to be relocated is minimal (although this is not very important).

The fitness of the solution is then defined as:

$$fitness(s) = - (hard_constraints \times large_weight + soft_constraints)$$

Initial population and setup

We find that generating the initial population by including the current placement of virtual machines within the system and its permutations will help the genetic algorithm to converge faster than starting from a purely random population.

The default setup for the genetic algorithm is to run for 100 generations over a population of 1,000 individuals and to use the ‘steady-state’ selection method, a ‘single-point crossover’ strategy, and a random mutation of genes.

Software

The software described in this article is written in Python and relies on the OpenStack command-line tool to interact with the infrastructure. It uses the PyGAD Python library for building the genetic algorithms. The software can be found at <https://github.com/ecmwf-lab/vm-placement>.

Results

Given 30 hosts and 250 virtual machines, the algorithm runs for under a minute on a modern laptop and outputs the list of virtual machines that need to be migrated to different physical hosts. The OpenStack command-line tool is then used to perform the live migration.

The CDS/ADS is a complex system, and its performance depends on many factors, some of which are due to the load generated by the users’ activity. It is therefore difficult to measure which part of any improvement is due to using an optimal placement of the virtual machines. Nevertheless, a few days after tuning the CDS, we reached a record: over 106 TiB delivered in 24 hours, in 600,000 user requests (Figure 4).

A multi-year project to modernise the CDS/ADS started in February this year. The CDS/ADS will continue to be hosted in OpenStack and in addition will make use of Kubernetes, which will improve our control of virtualisation. The software will be enhanced to make use of that finer control.

ECMWF publications

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

Technical Memoranda

- 903 **Lopez, P. & M. Matricardi:** Validation of IFS+RTTOV/MFASIS0.64- μm reflectances against observations from GOES-16, GOES-17, MSG-4 and Himawari-8. *September 2022*
- 902 **Haiden, T., M. Janousek, F. Vitart, Z. Ben-Bouallegue, L. Ferranti, F. Prates et al.:** Evaluation of ECMWF forecasts, including the 2021 upgrade. *September 2022*
- 901 **Hemri, S., T. Hewson, E. Gascon, J. Rajczak, J. Bhend, C. Spririg et al.:** How do ecPoint precipitation forecasts compare with postprocessed multi-model ensemble predictions over Switzerland? *August 2022*
- 900 **Betke, E., T. Quintino, S. Smart & T. Wilhelmsson:** Impact of GRIB compression on weather forecast data and data-handling applications. *August 2022*

EUMETSAT/ECMWF Fellowship Programme Research Reports

- 59 **Duncan, D., N. Bormann & A.J. Geer:** All-sky Assimilation of AMSU-A Window Channels. *August 2022*

ECMWF Calendar 2022/23

2022

Oct 24–28	7th SPARC General Assembly
Oct 31–Nov 4	Sixth WGNE workshop on systematic errors in weather and climate models
Nov 2	Advisory Committee of Co-operating States (Geneva)
Nov 14–17	ECMWF–ESA Workshop on Machine Learning for Earth Observation and Prediction
Dec 1–2	Council
Dec 5	30 Years of Ensemble Forecasting and Symposium for Prof. Tim Palmer
Dec 8–9	WarmWorld Kick-Off (Bonn)

2023

Apr 18–20	Advisory Committee for Data Policy and policy meetings of EUMETSAT and ECOMET
Apr 26	Policy Advisory Committee (virtual)
Apr 26–27	Finance Committee (virtual)
Jun 5–9	Using ECMWF's Forecasts
Jun 21–22	Council
Sep 4–8	Annual Seminar
Oct 4–6	Scientific Advisory Committee
Oct 19–20	Technical Advisory Committee (virtual)
Oct 24–25	Finance Committee
Oct 25	Policy Advisory Committee
Dec 5–6	Council

Contact information

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

Telephone National 0118 949 9000

Telephone International +44 118 949 9000

Fax +44 118 986 9450

ECMWF's public website www.ecmwf.int/

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

For any query, issue or feedback, please contact ECMWF's Service Desk at servicedesk@ecmwf.int.

Please specify whether your query is related to forecast products, computing and archiving services, the installation of a software package, access to ECMWF data, or any other issue. The more precise you are, the more quickly we will be able to deal with your query.



Newsletter | No. 173 | Autumn 2022

European Centre for Medium-Range Weather Forecasts

www.ecmwf.int