

Newsletter

No. 177 | Autumn 2023

Heatwave over southwest Europe

Global temperature records

Machine learning and data
assimilation

What next for Magics visualisation?

Thirty years of the ecgate service

New WMO data sharing



© Copyright 2023

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The ECMWF Newsletter is not peer-reviewed.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

The AIFS is launched



As I write this editorial, we are launching the Artificial Intelligence/Integrated Forecasting System (AIFS). This is our first forecasting system that is entirely based on machine learning, apart from establishing the initial conditions of forecasts. We are making it public in an alpha version: it has a resolution of approximately one degree (111 km), and it makes predictions for wind, temperature, humidity, geopotential and more. This is just a starting point. The AIFS will be regularly updated to increase resolution and to add fields. It will also be developed into an ensemble system, in other words it will produce a set of forecasts that includes information on the uncertainty inherent in any weather prediction.

The development of the AIFS does not mean that we are abandoning our current physics-based forecasting system: an ensemble-based system, which we currently run at a resolution of 9 km in the medium range. But it is a starting point from which we'll be able to progress. Tests of the AIFS root-mean-square error in geopotential height at 500 hPa already show an improvement that grows with lead time compared to our standard Integrated Forecasting System (IFS), based on the same initial conditions.

We decided to launch the AIFS in the wake of several companies' initiatives to produce weather forecasts based on machine-learning methods. These include NVIDIA's FourCastNet, Huawei's Pangu-Weather and Google DeepMind's GraphCast. We have made these systems available on ECMWF's public charts pages, based on our own initial conditions. The AIFS has now been added to those pages.

Machine learning is not just making inroads into forecasting but also into establishing the best possible initial conditions for forecasts. One of the articles in this Newsletter looks into the growing role of machine learning in data assimilation. It takes the view that we will make the furthest progress in the middle ground between model-led approaches and purely observation-driven approaches based on machine learning. The presence of machine learning was also felt in this year's Code for Earth programme. As described in this Newsletter, this ECMWF initiative brings together developers with mentors from the Centre to work on a variety of topics, which this year again included machine learning.

Other articles describe the by far warmest summer conditions ever recorded by the EU's Copernicus Climate Change Service run by ECMWF, and they analyse our forecasts for a heatwave in Europe. This Newsletter also looks at Météo-France's establishment of a Global Broker as part of a new data-sharing solution for the WIS2 community using the European Weather Cloud. In addition, it surveys the history of data visualisation at ECMWF and of the 'ecgate' service to our Member and Co-operating States. The latter article is a reminder that we are here to serve those states, not least with ever-improving forecasts: that rationale has given us the impetus to launch the AIFS now and to develop it further in the years to come.

Florence Rabier
Director-General

Contents

Editorial

The AIFS is launched 1

News

Heatwave over southwest Europe in August 2023..... 2

Global temperature records in the summer of 2023..... 4

One year of operational service in Bologna 6

ECMWF contributes to exascale computing project 7

Simulation of top-of-the-atmosphere visible reflectances with the IFS 8

Monitoring the Global Basic Observation Network 9

New observations July – September 2023..... 10

Ten teams take part in Code for Earth 11

Global seasonal fire danger forecasts available in the Climate Data Store 12

Earth system science

Combining machine learning and data assimilation to estimate sea ice concentration..... 14

Computing

What next for Magics visualisation?..... 22

Thirty years of the ecgate service at ECMWF 27

Running a Global Broker as part of the new WMO data sharing solution..... 32

General

ECMWF publications 37

ECMWF Calendar 2023/24 37

Contact information 37

Heatwave over southwest Europe in August 2023

Linus Magnusson, Claudia Di Napoli

July and August 2023 were in Europe characterised by two major heatwaves in the south, while northern Europe saw a lot of rain with floodings in Norway and Sweden. The July heatwave affected countries around the Mediterranean Sea, with major wildfires in Greece and other countries, while the August heatwave affected south-western Europe (Spain, France, Switzerland). In this article, we will focus on the predictability of the August heatwave. During this heatwave, the city of Lyon in France saw 17 days in a row with temperatures above 30°C (9–25 August), which peaked at 41.4°C on 24 August. In Spain, Bilbao on the north coast set a new record of 44.0°C on 23 August.

Extended-range forecasts

Starting with the extended-range predictions of the peak of the heatwave (22–24 August) for south-eastern France (in a box from 44°N to 46°N and 4°E to 6°E), we find that, as early as 30 days before the peak, the ensemble mean was slightly shifted (1°C) above the

model climate mean (see the forecast evolution plot). As extended-range forecasts can be affected by growing model biases, the plot also includes the model climate for days 30–32, in addition to days 5–7. However, no major differences for this area are present for the model climate distributions.

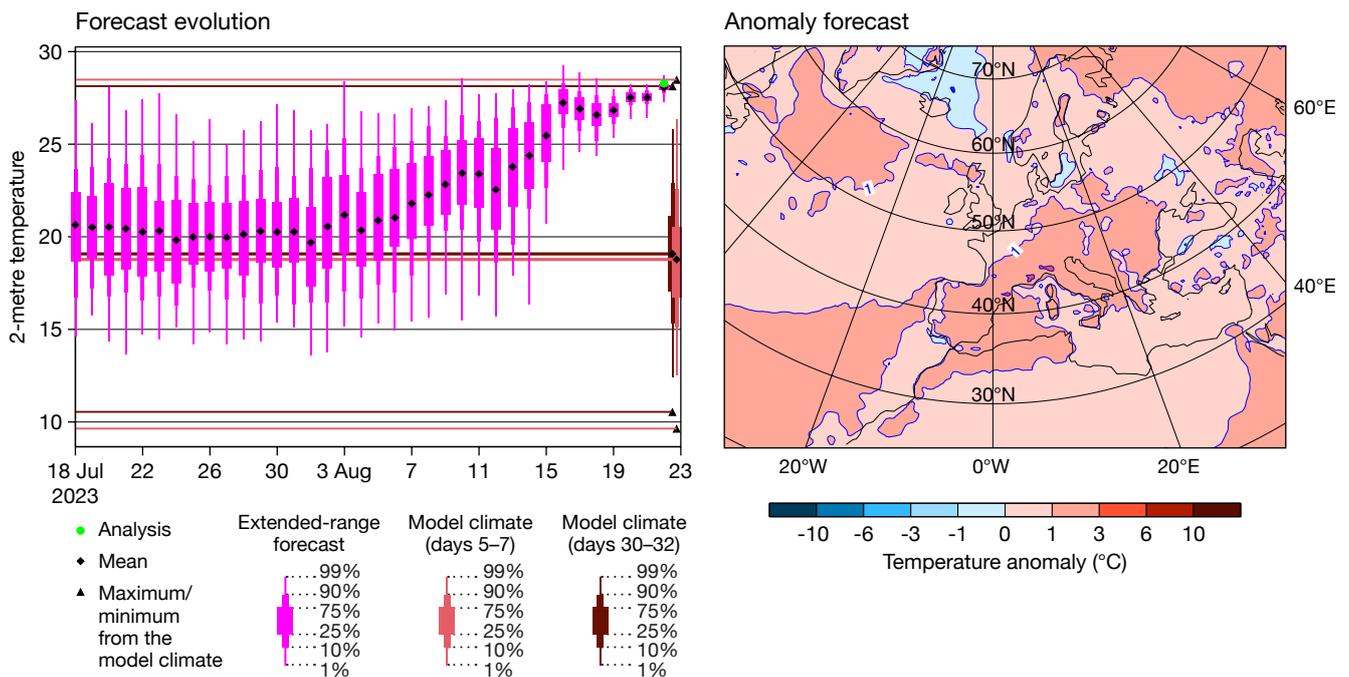
Around 20 days before the peak, the ensemble mean gradually started to be warmer than the model climate mean. This can be seen in the chart showing weekly mean temperature anomalies in the extended-range forecast from 7 August, which were above three degrees over parts of southern France. As the forecast evolution plot shows, the ensemble mean crossed the 99th percentile of the model climate on 16 August, six days before the start of the 3-day evaluation period. The 3-day ensemble mean temperature of the last forecast was close to the model climate maximum and also to the mean of the 166 observation stations in the area. As the model climate is based on re-forecasts from the past 20 years, one should note that it includes the August 2003 heatwave,

which explains the existing extreme of comparable severity in the model climate in the same area.

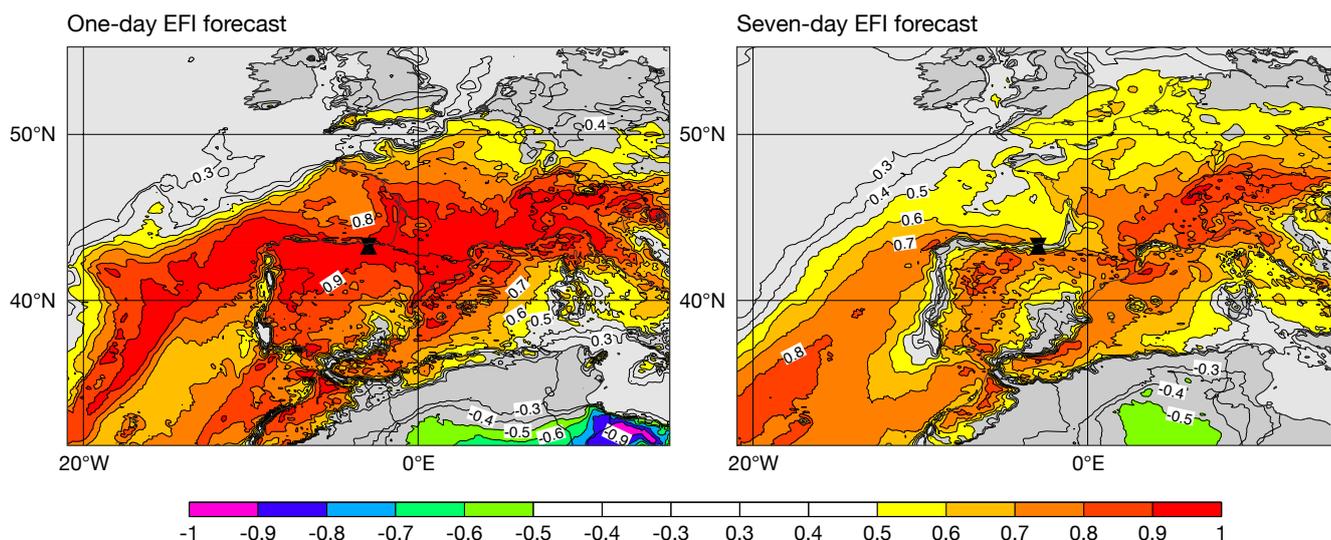
Medium-range forecasts

Looking at the medium-range prediction of the maximum temperature on 23 August, we find values of the Extreme Forecast Index (EFI) of 0.8 and above in the region of south-eastern France and Switzerland, in the forecast from 17 August 00 UTC. However, along the northern coast of Spain and the southern part of the French Atlantic coast, no such extreme was indicated in the EFI forecast from that date, although it is captured in a 1-day forecast (see the EFI charts for 23 August in 1- and 7-day forecasts).

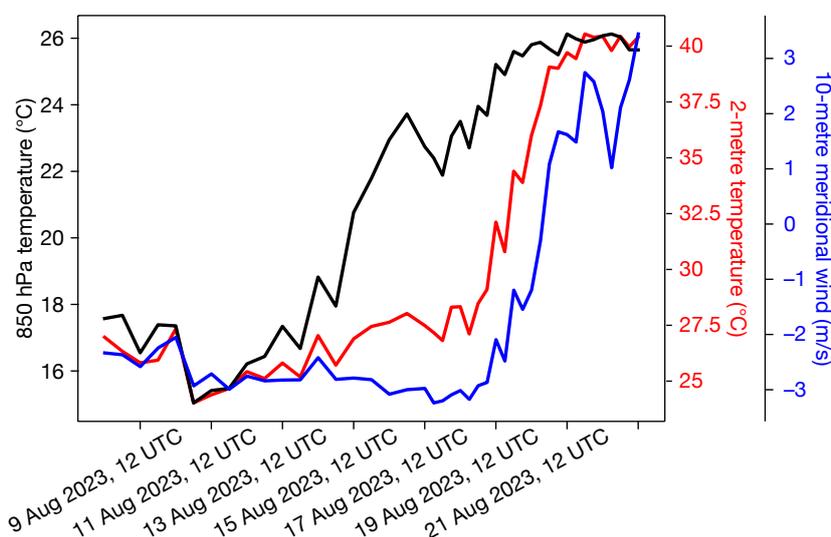
To understand the evolution of the prediction for maximum temperature along the Atlantic coast, we have investigated the prediction of 2-metre temperature (t2m), 850 hPa temperature (T850) and 10-metre meridional wind (v10m) for Bilbao. For the T850 we find a relatively smooth evolution of the signal, which is above 25°C on 20 August 00 UTC (see the figure showing the



Forecast evolution and anomaly forecast. Forecast evolution plot for 3-day (22–24 August 2023) 2-metre temperature over south-eastern France (left) and weekly mean of 2-metre temperature anomaly (21–27 August) in the extended-range ensemble from 7 August (right).



Extreme Forecast Index. The EFI for maximum 2-metre temperature on 23 August 2023 in a 1-day and 7-day forecast. Bilbao in Spain is marked with an hourglass.



Forecast evolution of ensemble means. Forecast evolution plots for the ensemble means of 850 hPa temperature (black), 2-metre temperature (red), and 10-metre meridional wind (v-wind) (blue), valid at 12 UTC on 23 August 2023 for Bilbao, Spain.

evolution of ensemble means). This predictability is in line with what is seen for the 3-day temperature over south-eastern France above. However, for forecasts from before 21 August, most ensemble members predicted a northerly component of the surface winds, leading to onshore winds and relatively mild temperatures in the forecast. The error in the flow observed around 20 August was probably related to a trough east of Portugal that was underestimated in the forecast. The winds from the sea in the forecasts cooled the lower part of the troposphere. Only with the change to a southerly wind direction did the ensemble become really extreme for Bilbao.

Conclusion

To summarise, the predictability of the August heatwave over south-

western Europe was good, with a signal appearing already in the extended range. However, due to errors in the winds along the Atlantic coast of Spain and France, the extreme surface temperature was missed in those locations in medium-range predictions, despite a good signal in 850 hPa temperature.

For heatwaves, one must remember that many aspects play a role for the severity for human health, such as minimum temperatures and the duration of the heatwave. During the August heatwave, temperatures at night (00 UTC) over south-eastern France were up to 8°C above climatology. This was compounded by very high temperatures (> 20°C at 00 UTC) at night for more than eight days in a row. Unusually high temperatures and their persistence in

time may be detrimental to human health as they imply a sustained lack of respite from daytime hot conditions at night. High humidity can also be detrimental because it reduces the rate of sweat evaporation and thus makes it difficult to regulate the body's temperature. To consider the effect of air temperature, humidity and other weather factors from a human health perspective, ECMWF has been producing heat stress indices from the ERA5 reanalysis, with plans to generate operational forecasts too.

For more details on this case, see the ECMWF Severe Event Catalogue: <https://confluence.ecmwf.int/display/FCST/202308+-+Heatwave+-+South-Western+Europe>

Global temperature records in the summer of 2023

Rebecca Emerton, Samantha Burgess, Julien Nicolas, Francesca Guglielmo, David Lavers, Adrian Simmons

The Copernicus Climate Change Service (C3S), implemented by ECMWF on behalf of the European Commission and with funding by the EU, routinely monitors key climate indicators and their evolution. Monitoring data from the ERA5 climate reanalysis over the summer of 2023 has shown just how anomalous global temperatures have been.

Global temperature records

Globally, the average surface air temperature over the northern hemisphere summer season (June–July–August) was by far the warmest in the ERA5 record, which goes back to 1940, with an average temperature of 16.77°C. This is 0.66°C above average, based on the 1991–2020 reference period. July and August 2023 were the two warmest months on record, and are estimated to have been around

1.5°C warmer than pre-industrial levels, based on the 1850–1900 reference period.

In December 2015, the nations of the world adopted the Paris Agreement, under which they would pursue efforts to limit the rise in the climatological global average temperature to 1.5°C above pre-industrial levels. While global average temperatures have exceeded this threshold on several occasions already during 2023, and for some days at a time during other recent years, it is important to stress that the limits set in the Paris Agreement refer to the average temperature of the planet over a 20-year period. Nevertheless, it is important to monitor the frequency with which we exceed the 1.5°C global warming threshold, as the cumulative effects of the exceedances will become

increasingly serious.

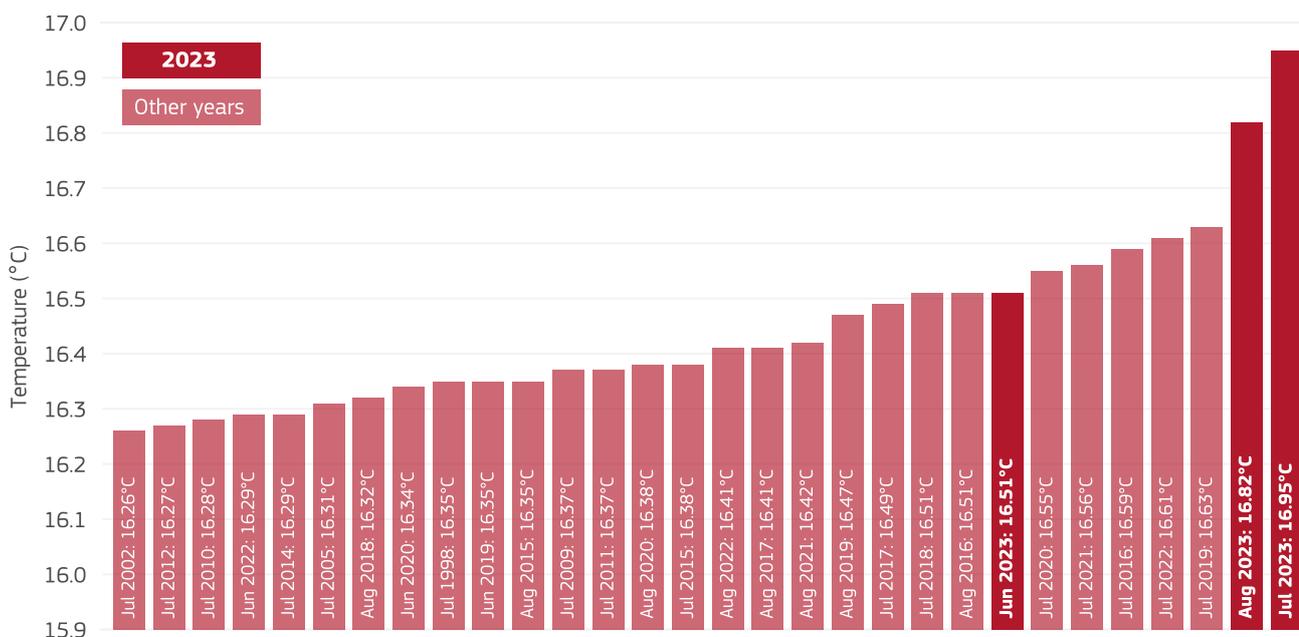
Currently, the warmest year on record is 2016, during which an exceptionally strong El Niño event occurred. The global average temperature anomaly for the first eight months (January to August) of 2023 currently ranks second-warmest on record, only 0.01°C below 2016, with four months of the year remaining and an El Niño event continuing to strengthen in the tropical Pacific.

Sea-surface temperatures

Unprecedented high sea-surface temperatures (SSTs) were an important contributing factor to the record-breaking global temperatures. The global ocean (here referring to the ocean from 60°S to 60°N) has seen SSTs remain at record high levels from April to August. Typically,

THE 30 WARMEST MONTHS ON RECORD GLOBALLY

Data: Globally-averaged surface air temperatures from ERA5 • Credit: C3S/ECMWF

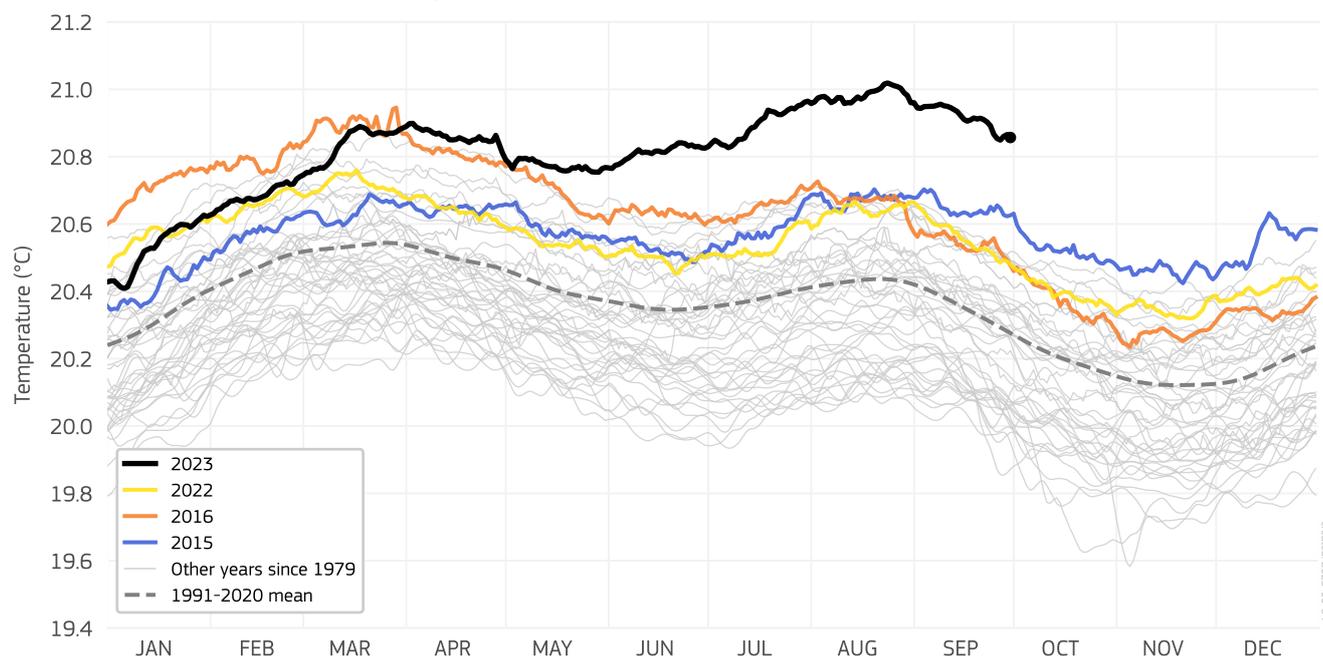


Monthly global average surface air temperatures of the 30 warmest months on record. The months are ranked from lower to higher temperature up to August 2023. Months from summer 2023 are highlighted in bold. Data source: ERA5. Credit: Copernicus Climate Change Service/ECMWF.

DAILY SEA-SURFACE TEMPERATURE

Extrapolar global ocean (60°S–60°N)

Data: ERA5 1979–2023 • Last data: 30 Sep 2023 • Credit: C3S/ECMWF



PROGRAMME OF
THE EUROPEAN UNION



IMPLEMENTED BY
ECMWF

Daily sea-surface temperature. The sea-surface temperature is averaged over the extrapolar global ocean (60°S – 60°N) during 2023 (black line) and for 1979 to 2022 (grey lines, with 2022 shown in yellow, 2015 in blue and 2016 in orange). Data source: ERA5. Credit: Copernicus Climate Change Service/ECMWF.

global average daily SSTs reach their highest level for the year in March. They then begin to fall before a slight increase again in July and August. This year, they saw a sharp rise in March, followed by only a slight dip during April and May. They then continued to rise to reach the highest value on record, 21.02°C, on 23 and 24 August. The previous record, of 20.95°C, was set in March 2016. This year, every day from 31 July to 1 September saw warmer SSTs than the previous record from 2016.

The high SSTs have coincided with the development of El Niño conditions in the tropical Pacific. This naturally occurring climate pattern, which sees warmer-than-average ocean temperatures in the equatorial Pacific, leads to a higher likelihood of unusually warm temperatures across many parts of the planet. However, high SSTs outside of the tropical Pacific basin are also playing an important role. This is particularly true in the North Atlantic (76°W–0°E, 0°N–60°N), which has also seen record-breaking

SSTs this summer. The previous highest daily SST for the North Atlantic was 24.81°C, in September 2022, which was exceeded on 5 August 2023. A new record-high daily North Atlantic SST was reached on 31 August, at 25.19°C.

Regional impacts

During the northern hemisphere summer of 2023, a number of heatwaves were experienced in many regions of the northern hemisphere, including southern Europe, the southern United States, Canada, Japan, Kazakhstan, China, Greenland and more. In contrast, much of northern Europe saw temperatures close to or below average during July and August.

Australia, several South American countries and much of Antarctica also experienced warmer-than-average conditions, and sea ice cover around Antarctica reached record low levels. The Antarctic sea ice extent for August was 12% below average, which is by far the largest negative anomaly for the time of year since satellite observations began.

The unusual behaviour of the sea ice extent around Antarctica began in the spring, after reaching an all-time minimum in February 2023. Typically, Antarctic sea ice reaches its minimum extent for the year in February and then increases to its annual maximum extent in September. This year, the increase has been lower than in previous years, remaining below average from May to August. In the Arctic, sea ice extent was also below average over summer, but above the record minimum from August 2012.

Climate monitoring

More information on the evolution of the climate can be found on the C3S website (<https://climate.copernicus.eu/>), including the C3S Climate Dashboard, monthly climate bulletins reporting on temperatures, hydrological variables and sea ice, regular news articles following key events, and the annual European State of the Climate (<https://climate.copernicus.eu/ESOTC>) report.

One year of operational service in Bologna

Matteo DellAcqua, Manuel Fuentes, Oliver Gorwits, Michael Hawkins, Alfred Hofstadler, Christine Kitchen, Martin Palkovič, Michael Sleight, Martin Suttie

On 18 October 2022, ECMWF successfully migrated all operational forecast services from its data centre in Reading (UK) to its new data centre in Bologna (Italy). On 27 June 2023, it successfully upgraded its operations to Integrated Forecasting System (IFS) Cycle 48r1. This article provides a review of one year of operational service in Bologna.

Resilience and performance

Full handover of the data centre from the Regione Emilia-Romagna to ECMWF took place on 29 April 2022. Since March 2022, when the Copernicus Climate Data Store became operational in Bologna as the first operational service, the data centre has been running smoothly, providing the environment to host ECMWF operational activities.

The mechanical infrastructure has been very reliable and the electrical continuity system successfully covered any voltage losses during the first year of operation. The GARR Consortium, the Italian national research and academic network, provides a resilient 100 Gigabits per second (Gbps) service for dissemination. Alongside GARR, resilient 1 Gbps and ADSL network services are provided by Lepida. These are intended for analyst remote access to the data centre independent of the GARR service as well as for emergency access during incidents.

Each of the data centre's two data halls has an 'IP Fabric' network topology, and the two data stores have a third production network. There is also a Management and Monitoring network supporting all network and server devices, and a limited-scope out-of-band emergency network as a last resort. Application-specific security zones make it possible to carefully manage the access between computer systems. This modern design has already proved its worth during the first year of operations. ECMWF operates a modern network, server, and security infrastructure in Bologna. Some stability issues were experienced in the



The ECMWF data centre in Bologna. The building is a renovated former tobacco factory.

first year of operation, which will be learned from and contribute to continual improvement.

A new HPCF and move of the data handling system

The greater power and cooling capacity provided by the Bologna data centre compared to what was available in Reading made it possible to install a new high-performance computing facility (HPCF) from Atos. Without the larger HPCF, it would not have been possible to deliver IFS Cycle 48r1. Implementing the Atos HPCF has not been without challenges. However, it is now facilitating reliable forecast production and supporting the research and project activities at ECMWF as well as the workload of ECMWF's Member and Co-operating States.

The switch of forecast production and all related services, such as observation acquisition and pre-processing, product generation and dissemination services, to the HPCF in Bologna was carried out on 18 October 2022. ECMWF was running IFS Cycle 47r3 at the time. On 27 June 2023, a major scientific upgrade to Cycle 48r1 was carried

out. The changes included a horizontal resolution increase of the medium-range ensemble (ENS) from 18 to 9 km as well as a major upgrade of the configuration of the extended-range ensemble.

In the autumn of 2022, the Data Handling System was moved from Reading to Bologna: three tape libraries, 17 disk systems, 150 servers, 290 tape drives, and almost 30,000 tapes were transported by lorry. Despite the complexity of the task, it was possible to provide a continuous archive and retrieval service for all critical production components throughout the whole migration process, and the services have continued to perform reliably over the year.

A successful migration

The last year in Bologna has shown that we successfully migrated operations to the new data centre and the new HPCF as well as implementing a major cycle upgrade – altogether a huge undertaking for ECMWF. The migration was used to move from infrastructure that had evolved organically to a design that aims to accommodate the current requirements and support enhancements in the near future without compromising any of our

modern design principles.

Such principles are, for example, improved operational resiliency, better IT security, higher levels of efficiency, and reduced environmental impact, to name just a few. In addition, and to

react to the new situation of having a remote data centre, new software tools and working processes have been introduced. The improvements, delivered through careful design work, improved tool-chains, and efficient

work processes, are clearly visible in the daily results of delivering high-quality forecasts and services to our Member and Co-operating States and other external and internal stakeholders of ECMWF.

ECMWF contributes to exascale computing project

Jenny Wong

In April 2021, the EuroHPC-funded 'IO – Software for Exascale Architectures' (IO–SEA) project began. The three-year project aims to implement solutions for scaling applications to exascale high-performance computing (HPC) systems. It achieves this by designing storage and data access architectures in collaboration with a wide range of use cases with high I/O demands. As the project reaches completion, we summarise ECMWF's contributions in a use case and as leader in the development of DASi, a high-level Data Access and Storage Interface.

What is IO–SEA?

The IO–SEA project (<https://iosea-project.eu/>) is coordinated by the French Alternative Energies and Atomic Energy Commission (CEA) and involves a consortium of ten partners, including Atos and ECMWF. It aims to provide a novel HPC data management and storage platform for exascale computing based on hierarchical storage management (HSM) and on-demand provisioning of storage services. The platform

should efficiently make use of storage tiers ranging from non-volatile memory express (NVMe) to tape-based technologies.

ECMWF's use case

The project operates under a co-design principle, where requirements from I/O-intensive use cases drive the development of the IO–SEA architecture. ECMWF's weather forecasting workflow consumes and produces terabytes of data in a single model run, making it an ideal candidate for evaluating potential solutions to I/O bottlenecks. Solutions developed in the project are benchmarked with a simplified workflow that mimics the competition for I/O resources between the writing of model outputs at each step and the reading of these outputs for product generation.

In the project so far, we have been able to test the workflow on different HPC systems, such as the Jülich Supercomputing Centre's prototype modular supercomputer (DEEP) and the supercomputers operated at IT4Innovations. We have also evaluated the performance

implications of using the Smart Burst Buffer (SBB) ephemeral service developed by Atos. The SBB service consists of nodes equipped with NVMe devices to accelerate data access, and initial tests show positive results in some of the benchmark metrics.

DASi

The Data Access and Storage Interface (DASi, <https://github.com/ecmwf-projects/dasi>) is a layer on top of ECMWF's existing software, which makes that software more accessible to other scientific domains. It achieves this through semantic data management: all access and control of data is performed using scientifically meaningful indexing relevant for the particular application domain. DASi is inspired by the Fields DataBase (FDB) software in operational use at ECMWF. However, unlike the FDB, it can be configured for any scientific domain. The interface enables users to write, retrieve and query data through metadata, as well as to set data policies relating to data lifetime or access frequency.

Increasing the accessibility of



IO–SEA meeting.

In-person meeting of IO–SEA participants on 10 October 2022 in Paris.

ECMWF's existing software fosters collaboration and enables other organisations to contribute and strengthen the community around our software. In this project, many of the partners have contributed to developing ECMWF's software further. Examples include building a GekkoFS and CORTX-Motr backend for the FDB in collaboration with the University of Mainz and Seagate, respectively. CEA has also developed a POSIX interface to DASI, and therefore to the FDB. The solutions we were hoping to

achieve in this project, such as hierarchical storage management, are complicated, and we were able to bring our expertise from the development of the Meteorological Archival and Retrieval System (MARS) and the FDB into guiding the scope and direction of progress.

Outlook

This project has been extremely fruitful in extending ECMWF's FDB object store to support alternative technologies and building a wider community of users and collaborators

for our existing software stack. We have also been able to adapt an emulation of the operational forecast to use novel data node architectures developed by Atos and evaluate their initial impact on the I/O demands of our workflow.

DASI has also been adopted by other use cases in the project across a wide range of domains, including lattice quantum chromodynamics and electron microscopy, with developments continuing until the project ends in spring next year.

Simulation of top-of-the-atmosphere visible reflectances with the IFS

Cristina Lupu, Josef Schröttle, Philippe Lopez

Since 2016, global simulated infrared brightness temperatures have provided an exciting and unique view of the atmosphere predicted by ECMWF's high-resolution Integrated Forecasting System (IFS), up to ten days into the future. More recently, the generation of simulated images from ECMWF forecasts has been further extended to visible reflectances. Reflectance, expressed as a percentage, is the ratio between outgoing and incoming solar shortwave radiation. The new development will not only provide

valuable information about the larger-scale dynamics but will also enable the study of meteorological features that are usually not well detected at infrared frequencies, such as low-level clouds and fog.

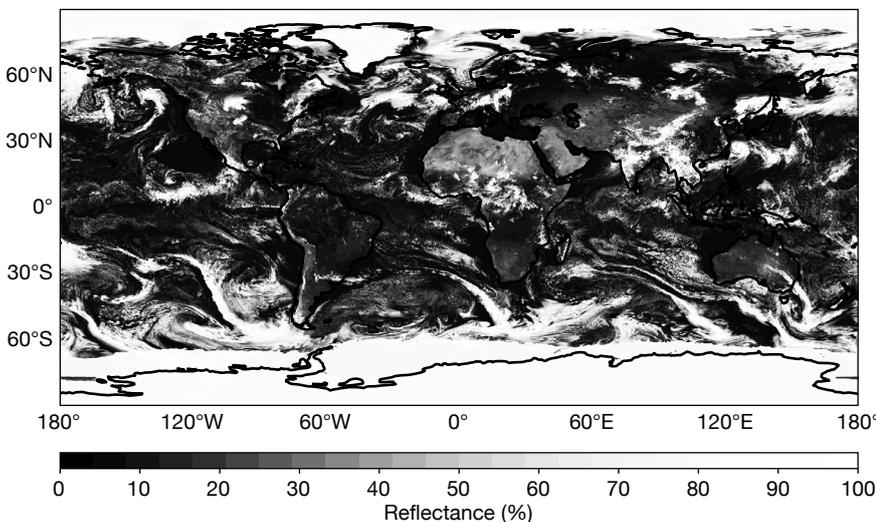
The move to visible reflectances follows recent progress in the NWP SAF (Numerical Weather Prediction Satellite Application Facilities) in radiative transfer modelling at visible wavelengths. It also comes after the integration of an RTTOV-13.0/ MFASIS (Method for FAsT Satellite

Image Synthesis) look-up-table-based observation operator in Cycle 48r1 of ECMWF's IFS. Reflectances that would be seen in a visible channel (for example at 640 nm) can now be computed during the model run at every grid point.

Capturing clouds and surface conditions

The example image shown assumes a nadir view at every model grid point, free from real-life satellite geometry distortions at high latitudes. It enables a unique perspective, to see the entire globe in perpetual daylight at various forecast lead times. Under the assumption that both the sun and satellite are overhead everywhere, the effect of sun glint over sea and lakes is not considered.

A key advantage of simulating reflectance is the ability to capture Earth's clouds at multiple scales and heights throughout the troposphere, as well as variations in surface conditions associated with soil properties, vegetation and snow cover. The figure shows multiscale meteorological patterns: cyclones over the North Atlantic; a synoptic-scale frontal system stretching over Central Europe; bands of deep convective clouds spanning around the equatorial Pacific and Atlantic,



Simulated visible reflectance. Global simulated imagery shows visible reflectance in the 0.64 μm spectral band for a 5-day forecast initialised at 00 UTC on 31 July 2023 from the ECMWF operational forecast.

indicative of the Intertropical Convergence Zone (ITCZ). In addition, the bright Saharan desert contrasts with the darker tropical rainforests of Central Africa, and sea ice around Antarctica and in the Arctic is easily distinguishable from the darker surrounding oceans.

The simulation of top-of-the-atmosphere visible reflectances will soon be an integral part of the operational IFS. It will be available within the standard delivery times of all other ECMWF data and products. This product, available in three-hourly steps up to day two, six-hourly steps up to day five, and twelve-hourly steps up to day ten, will complement the infrared images

already available in ECMWF's catalogue and will be available to the Centre's Member and Co-operating States.

Extreme weather events

The new product will play a crucial role in improving the visualisation of forecasts of high-impact extreme weather events through the complementarity of infrared and visible frequencies. It also provides enhanced capabilities to capture weather patterns over a wide range of spatial and temporal scales.

In addition, the new product will be available at km-scale in the context of the EU's Destination Earth

(DestinE) initiative, where simulated visible reflectances will support the evaluation of extreme weather forecasts. While shorter time intervals uncover the evolution of squall lines or storm tracks in greater detail at the km-scale, Rossby wave patterns in the mid-latitudes or hurricanes in the tropics appear as distinct features in the simulated visible imagery and could help forecasters to better assess model performance.

This work is also relevant to developments towards the assimilation of solar reflectances in ECMWF's 4-dimensional variational data assimilation system (4D-Var).

Monitoring the Global Basic Observation Network

Cristina Prates (ECMWF), Timo Pröscholdt, Luis Nunes, Amro Abouelenein (all WMO), Tanja Kleinert (DWD), Tim Oakley (UK Met Office and WMO)

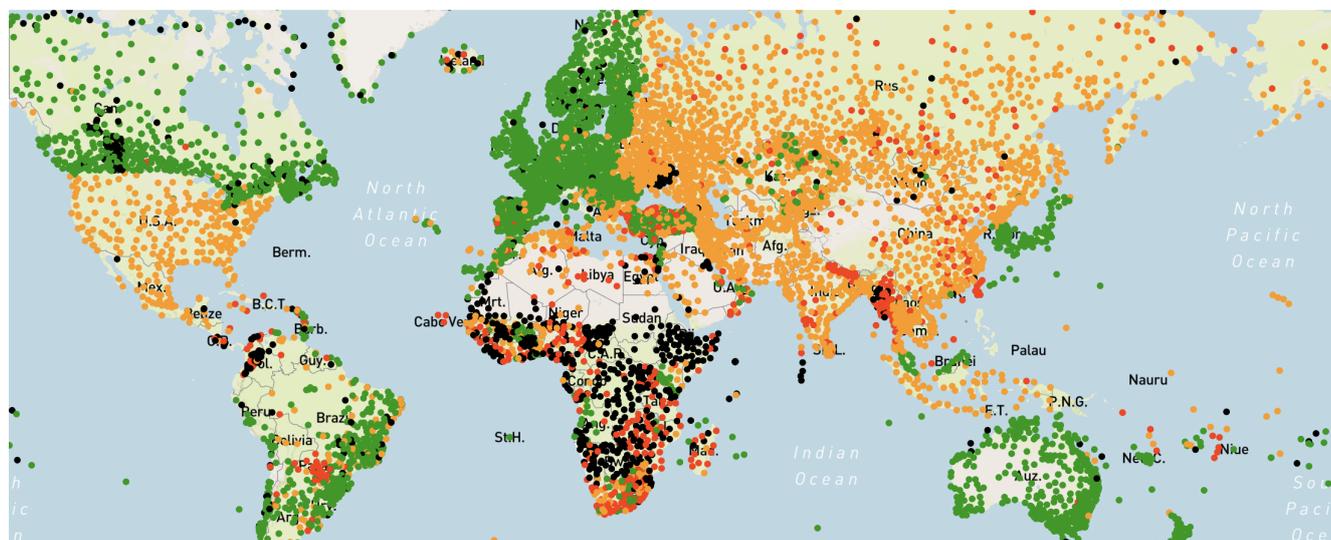
A new module to monitor the Global Basic Observation Network (GBON) was made available in the latest major release in May 2023 (version 1.4.0) of the graphical web-based interface created for the WIGOS (WMO Integrated Global Observing System) Data Quality Monitoring System (WDQMS). This module provides the essential tool to assist the World Meteorological

Organization (WMO) in the implementation and compliance monitoring of GBON that came into force on 1 January 2023. ECMWF is one of four centres supplying near-real-time monitoring information to GBON.

Significance of GBON

GBON is the WMO response to the urgent need to strengthen surface-

based observation capabilities to support numerical weather prediction (NWP) and climate activities. In recent decades, satellites have taken over as the primary source of observations for NWP, offering an enormous quantity of observations with consistent geographical and temporal coverage of the atmosphere. However, some critical variables, such as surface



WDQMS daily monitoring page. The image shows part of the WDQMS webtool daily monitoring page of GBON-affiliated land surface stations, for surface pressure measurements on 11 September 2023. The station's GBON compliance is shown as either compliant (green), reporting below requirements (orange and red), or not reporting at all (black).

pressure or winds, cannot be accurately measured from space. For these, global NWP systems continue to rely on surface-based observations. While in some parts of the world these observations are available in abundance, in others they are scarce or even non-existent. Large geographical gaps in the availability of such ground-based observations have driven the GBON initiative.

In 2019, WMO Members agreed to establish GBON. This international agreement, based on a global optimal design, establishes clear requirements for countries to acquire and internationally exchange surface-based observations of the most important weather and climate variables.

Six land-based variables – atmospheric pressure, temperature, humidity, wind speed, precipitation, and snow depth – have been designated as the most essential for global NWP and climate data reanalysis purposes. Furthermore, WMO Members are required to measure and report these variables hourly in the case of surface land stations and twice per day for upper-air (radiosonde) land stations. These GBON requirements are applied to stations registered in OSCAR/Surface that have been affiliated to GBON by WMO Members.

Role of four NWP centres

The GBON module, like the near-real-time NWP monitoring module in the Wdqms webtool (<https://wdqms.wmo.int/>), also relies on near-real-time monitoring information received from four participating centres that deliver global NWP: the German Meteorological Service (Deutscher

Wetterdienst – DWD), ECMWF, the Japan Meteorological Agency (JMA) and the United States National Centers for Environmental Prediction (NCEP). The GBON module monitors all GBON-affiliated surface land stations and upper-air stations on land. This module assesses GBON station-level compliance using data availability as the primary performance measure. The near-real-time monitoring information provided by the four NWP centres is aggregated by station, variable, and time interval. The calculated total number of reports is then compared against the corresponding thresholds established by the GBON compliance criteria (mentioned above) for the two different station types.

An example of the daily GBON monitoring pages for all GBON-affiliated land surface stations that are required to provide surface pressure (one of the mandatory variables) is shown in the figure. It is evident that some regions of the world are consistently providing surface pressure observations to the NWP centres (shown by green dots). Other parts of the globe, on the other hand, have spatial availability gaps (shown by black dots) and temporal availability gaps (shown by orange and red dots). These gaps in data availability are the result of either missed observations or observations that were made but not shared internationally. It is worth noting that encouraging WMO Members to register their existing stations as GBON and making them compliant with requirements is an important component at this stage of the GBON implementation. The next stage towards achieving a global optimal design of the land-based

network will involve installing new stations to provide sparse areas with better spatial and temporal coverage.

Plans for the future

The implemented GBON module currently only considers availability to monitor GBON station-level compliance. However, the plan is to incorporate quality requirements into the monitoring of GBON stations in the near future. This is extremely important since it supports the WMO's aim to enhance the quality of weather and climate information through the provision of not just more but also better observations. Meanwhile, the Infrastructure Commission is looking into GBON expansion in other domains. The aim is to first extend GBON to hydrology and the cryosphere, for which discussions have just started, followed by climate monitoring, the ocean and atmospheric composition. Also noteworthy is the creation of the Systematic Observations Financing Facility (SOFF), a funding mechanism of the WMO in partnership with the United Nations Development Programme (UNDP) and the United Nations Environment Programme (UNEP), to support the implementation and operation of GBON. This seeks to deliver long-term financial and technical help to countries with the biggest observational gaps, currently to least developed countries and Small Island Developing States. The goal is to enable the generation and international sharing of fundamental weather and climate observations in a sustainable manner.

New observations July – September 2023

The following new observations have been activated in the operational ECMWF assimilation systems during July – September 2023.

Observations	Main impact	Activation date
Ozone retrievals from NOAA-20 OMPS	Ozone analysis in numerical weather prediction	5 July 2023
Radiances from NOAA-21 ATMS	Temperature, humidity, dynamics	23 August 2023
Radio occultation bending angles from PlanetIQ	Temperature in upper troposphere/lower stratosphere	12 September 2023

Ten teams take part in Code for Earth

Athina Trakas

Code for Earth, an ECMWF-run innovation and partnership programme, fosters collaboration and supports advancements in weather and climate research, including in the Copernicus programme and the Destination Earth (DestinE) initiative, which are both EU-funded. Since its first edition in 2018, when it was called the Summer of Weather Code, the programme has brought together talented individuals and developer teams with experienced mentors from ECMWF to work on cutting-edge projects covering a wide range of topics. This year, ten developer teams participated in Code for Earth.

How it works

Each summer, several individuals and developer teams from different backgrounds test, explore and/or develop open-source software solutions supported by ECMWF's mentors. Each team receives a €5,000 stipend after completion. Their projects tackle topics such as data science in weather-, climate- and atmosphere-related challenges, including visualisation, machine learning/artificial intelligence, user support tools and data analysis. By encouraging multidisciplinary collaboration and embracing open-source principles, Code for Earth facilitates the development of cutting-edge solutions and advancements in Earth system sciences.

Since its start, the programme has produced 35+ open-source software developments highly beneficial to activities at ECMWF and in its Member and Co-operating States.

The programme has partnered with Copernicus, Destination Earth and two cloud initiatives: the European Weather Cloud, run by ECMWF and EUMETSAT, and the EU-funded Copernicus Data and Information Access Service WEkEO, implemented by ECMWF, EUMETSAT, Mercator Ocean and the European Environment Agency.

2023 results

This year, ten selected dynamic and enthusiastic developer teams contributed their expertise and

Voices on Code for Earth from mentors

"It is one of the few programmes that bridges the gap between computer science and climate science and enables young people to realise their ideas in this field."

"It's a great programme to get people involved in small research projects!"

problem-solving skills to address real-world challenges. They were supported by mentors from ECMWF and partner organisations who are world experts in their field. ECMWF, in turn, could collaborate with motivated, innovative external teams who bring different skill sets and fresh perspectives to exciting challenges.

- **DeepR: Deep Learning for High-Resolution Reanalysis Data** – DeepR used deep learning techniques, including U-Net, conditional GAN, and diffusion models, to generate regional reanalysis data by downscaling global data from ERA5. A validation framework was established to ensure an accurate representation of physical processes.
- **Atmospheric Composition Dataset Explorer: API for Atmospheric Composition Diagnostics** – This project developed an application programming interface (API) and an application to automate the creation of time series, Hovmöller, and geospatial plots using the Atmosphere Data Store of the Copernicus Atmosphere Monitoring Service (CAMS) run by ECMWF. It facilitated data retrieval, homogenisation, and visualisation of CAMS datasets.
- **Benchmarking Surface Heat Fluxes: Validation of Land Surface Variables** – The LANDVER validation package was expanded to include validation of latent and sensible surface heat fluxes against Eddy-Covariance measurements. This tool provided insights into the impact of soil moisture stress on surface heat fluxes and the functionality of ECMWF's ECLand.
- **Compression of Geospatial Data: Improved Compression Techniques for Diverse Datasets** – This project refined the xbitinfo compression technique to account for varying information densities in datasets, leading to more efficient compression.
- **ChatECMWF: Conversational Search Engine for ECMWF Data** – Using natural language processing and models like ChatGPT, ChatECMWF was developed to help users access ECMWF's weather data through natural language queries.
- **Fire Forecasting: Wildfire Prediction with Machine Learning** – By utilising fire data from ECMWF's Global Fire Assimilation System (GFAS) and meteorological forecasts, a machine-learning framework was developed for European wildfire forecasting, aiming for integration into ECMWF's operational pipeline.
- **TropiDash: Dashboard for Tropical Cyclone Hazards** – TropiDash provided a platform on Jupyter Notebook to visualise key meteorological parameters related to tropical cyclones, aiding hazard comprehension.
- **Sketchbook Earth: Accessible Climate Intelligence Reports** – Moving away from traditional ECMWF tools, Sketchbook Earth used Jupyter notebooks and the Copernicus Climate Data Store (CDS) to present climate data more visually.
- **TesseRugged: Enhancing Reanalysis Data's Spatial Resolution** – TesseRugged worked on improving the spatial resolution of global reanalysis datasets, such as ERA5, through



Group photo. The participants of Code for Earth 2023.

model output strategies and deep learning for downscaling.

- **Diffusion Models on WeatherBench: Machine Learning for Weather Forecasting** – This project used the WeatherBench dataset to explore the potential of diffusion

models in weather forecasting, aiming to advance the field with the publication of code and trained models.

The Final Code for Earth Day at the end of September marked the completion of this year’s edition, when the teams presented their work

and innovative results. The results are accessible on GitHub (<https://github.com/ECMWFCode4Earth>) and might be integrated into operational schemes.

Code for Earth 2024 and beyond

As we look ahead, Code for Earth continues to evolve and adapt to emerging challenges as well as engagements with partner organisations. Constant evolution in IT, combined with new ways of providing information for better-informed users and end-users, creates many potential requirements. We are working on transforming the most relevant ones into Code for Earth challenges each year. For more information and feedback, check our website (<https://codeforearth.ecmwf.int/>) or contact us directly at codeforearth@ecmwf.int.

Global seasonal fire danger forecasts available in the Climate Data Store

Francesca Di Giuseppe, Christopher Barnard, Eduardo Damasio-Da-Costa

On behalf of the EU's Copernicus Emergency Management Service (CEMS), ECMWF has recently widened the fire danger data offering in the Climate Data Store (CDS) to include a set of fire danger forecasts with lead times up to seven months. The CDS is run by the EU's Copernicus Climate Change Service (C3S) implemented by ECMWF. The new data are provided by the Global ECMWF Fire Forecast (GEFF) model, which uses the Centre's seasonal forecasts to drive three different fire danger models developed in Canada, the United States and Australia. The dataset is made openly available for the period 1981 to 2022 and will be updated regularly, providing a resource to assess the predictability of fire weather at the seasonal timescale. The dataset complements the availability of real-time seasonal forecasts provided by the European Forest Fire Information System (EFFIS).

Availability of seasonal fire danger forecast

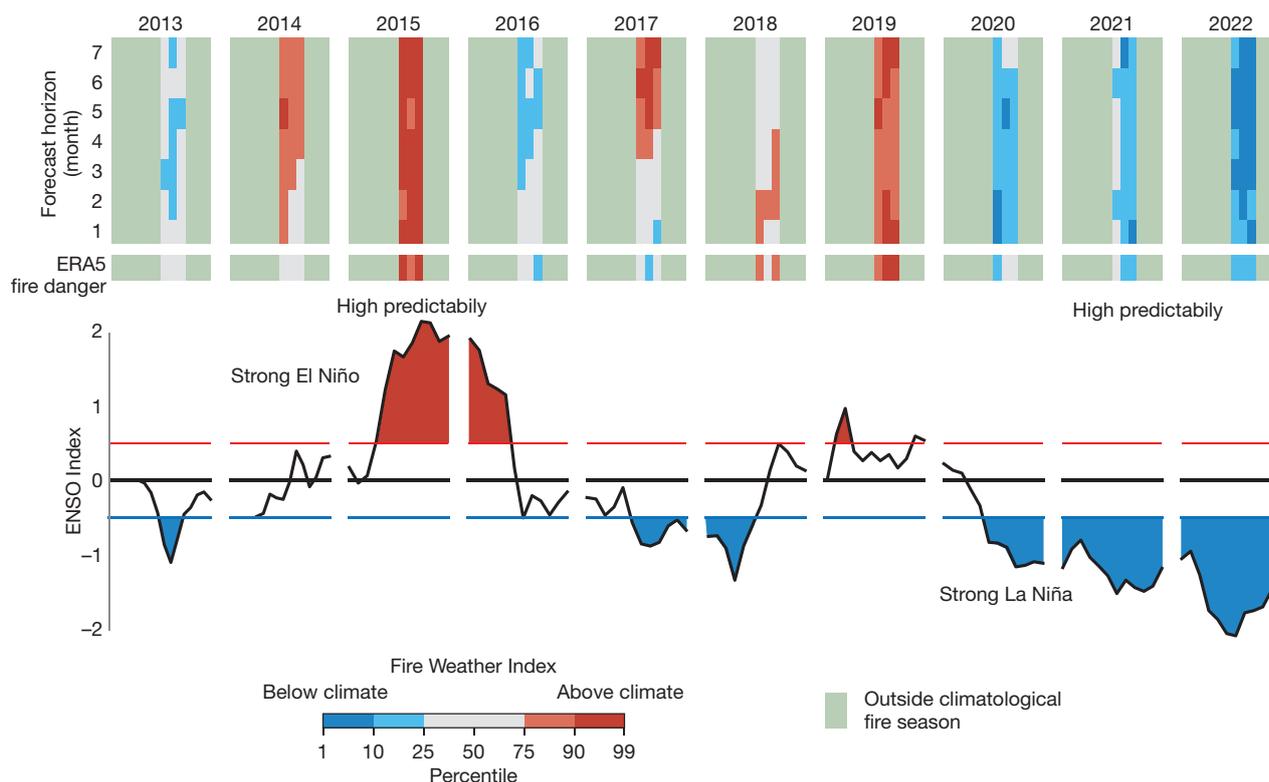
Global fire danger seasonal hindcasts have been generated using ECMWF's SEAS5 seasonal forecasts for the period of 1981 to 2022 as input to GEFF. The GEFF model is open source and available from a public repository under an APACHE2 license. The current version is 4.1. Data are archived in the CDS with several advantages: open access via a user-friendly web interface; bulk access via a convenient API; integration with the CDS toolbox for performing server-side operations; and shared visualisation and data analysis tools.

The fire danger seasonal forecast dataset has global coverage and a spatial resolution of about 0.25 degrees (about 35 km). Natively, data are laid out over an octahedral reduced Gaussian grid (O320) and

archived as GRIB2, a standard format published by the World Meteorological Organization. Users can also request data in NetCDF format, which implies an internal remapping data transformation. Forecasts are issued monthly, on the first day of each month, with a lead time of 216 days (about seven months). The dataset is updated regularly with a delay of a few months compared to real time, which is available through CEMS.

Skill of seasonal fire danger forecast

Forecasting fire danger is key in preventing fires and taking protection measures as it improves the readiness of fire professionals and enables the timely and efficient allocation of resources. A limited number of studies show that, besides well-established fire danger forecasts with lead times of a few days, skilful



Prediction of monthly mean fire danger anomalous conditions between 2013 and 2022 over Indonesia. Months are classified as above or below the 1981–2022 climate mean using percentiles. Anomalies from ERA5 fire danger are compared to SEAS5 fire danger forecasts for increasingly longer lead times to highlight the predictability of anomalous conditions. Months outside the traditional fire season are masked out. They are months in which the mean Fire Weather Index is lower than a third of the year's maximum. The ENSO index helps to identify years of strong positive or negative anomalies, in which El Niño or La Niña conditions are established. These years correspond to periods of high predictability, for which anomalous conditions could be predicted up to seven months before.

predictions of fire danger are possible up to the seasonal timescale for Mediterranean Europe. Local soil moisture anomalies and heatwaves have been identified as an important source of this predictability. Seasonal forecasting of fire weather conditions throughout the world have also been found to correlate with large-scale climate patterns, such as the El Niño–Southern Oscillation (ENSO) and the Indian Ocean Dipole. This shows that fire weather conditions could potentially be predicted months in advance for various seasons and regions.

Looking at the Anomaly Correlation Coefficient (ACC) of the released dataset, verified against the fire danger calculated using ECMWF's ERA5 reanalysis, we found that globally anomalous conditions for fire weather can be predicted with confidence one month ahead. In some regions the prediction can be extended to two months ahead. In most situations beyond this horizon, forecasts do not improve on climatology.

However, an extended predictability window, up to 6–7 months ahead, is possible when anomalous fire weather is the result of large-scale phenomena, such as ENSO. This is a climate pattern characterised by the warming of the surface waters in the central and eastern tropical Pacific Ocean during El Niño conditions. Those conditions often lead to a shift in rainfall patterns, resulting in reduced precipitation in Southeast Asia, including Indonesia. This can create drier-than-normal conditions, especially in peatland areas, making them more susceptible to fires. The conditions established by a strong El Niño exacerbate landscape flammability, but it is human activities that play a significant role in igniting fires. In Indonesia, particularly in the regions of Sumatra and Kalimantan, land clearing practices such as slash-and-burn agriculture, illegal logging, and peatland drainage for agriculture have been responsible for extensive burning in the past. The release of large amounts of smoke and pollutants into the atmosphere has

affected air quality not only in Indonesia but also in neighbouring countries, such as Malaysia and Singapore, generating international health emergencies.

The establishment of a positive or negative ENSO is usually monitored using a multivariate index obtained by extracting the leading combined Empirical Orthogonal Function (EOF) of five different variables over the tropical Pacific basin (30°S–30°N and 100°E–70°W). During a strong positive or negative ENSO, seasonal prediction of fire weather is enhanced up to seven months ahead (see the figure). Efforts to mitigate the impact of fires during ENSO events in Indonesia could therefore benefit from an early warning system at this timescale. This could help to guide land management practices and implement fire prevention and suppression measures before the burning takes place.

The dataset can be accessed at <https://doi.org/10.24381/cds.b9c753f1>.

Combining machine learning and data assimilation to estimate sea ice concentration

Alan Geer

ECMWF aims to analyse the full state of the Earth system from the atmosphere and ocean through to the land surface and cryosphere. This is intended to be achieved through closer coupling of the relevant models and the data assimilation systems that combine the model forecasts with new observations. But aspects of the Earth system, such as sea ice, snow, soil and vegetation, are hard to model from physical first principles, so in practice, the modelling can be quite empirical. Earth system modelling components are often parametrized or fitted based on a limited set of observations from experimental ground stations, and they may struggle to perform in other locations. A ‘model first’ approach has served us well in the atmosphere, where at least the dynamics are mostly well known: here the purpose of observations is to correct the physical trajectory of the model. But the recent explosion in machine learning for Earth system applications has shown us an ‘observation first’ approach. If observation-driven machine learning forecasts

are starting to do better than physical forecasts, it suggests we have not been making good enough use of observations to improve our forecast models. Especially for Earth system applications where models are already partly empirical, there is clearly great potential to let the observations increasingly define these models. However, as demonstrated in this article with the example of sea ice assimilation, the best results are unlikely to come by throwing away physical models entirely, but by carefully combining known physics with empirical components.

Getting more out of satellite observations

Taking a global viewpoint, available observations of the Earth system come mainly from satellites, and primarily from the naturally-generated radiation that is emitted by the Earth. An observational viewpoint starts from understanding what information the measured radiances contain and how to extract that information into useable geophysical variables. If we take microwave observations as an example (Figure 1), we

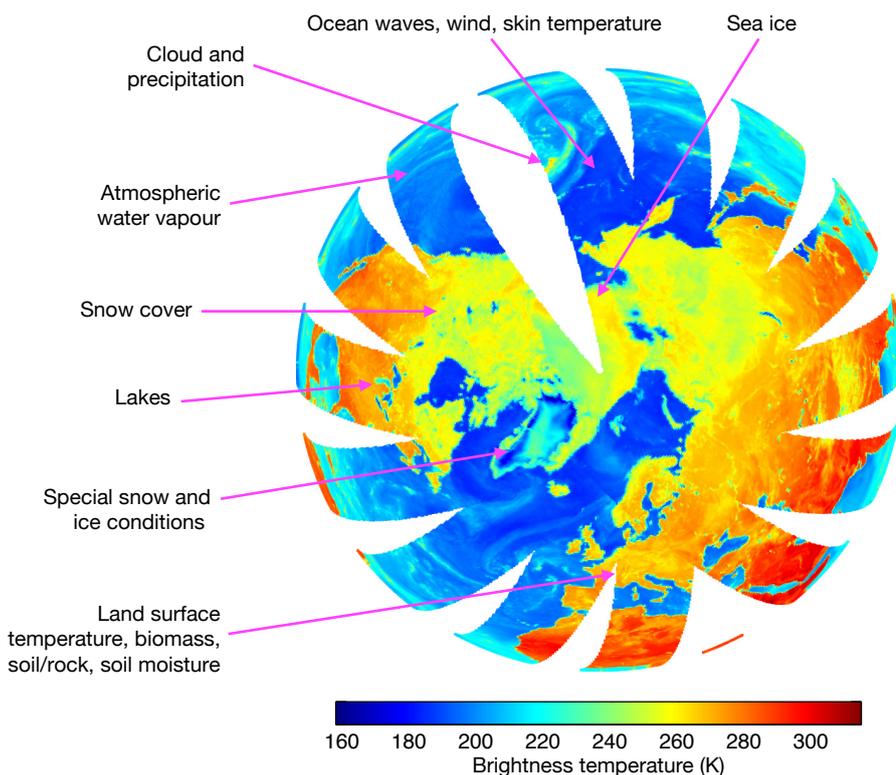


FIGURE 1 Observed brightness temperatures (an alternative way to represent radiances) from 12 hours of microwave imager overpasses of the northern hemisphere on 7 December 2020, labelled with the main geophysical sensitivities in the data. The image is composed of roughly 40 km resolution superobservations based on Advanced Microwave Scanning Radiometer 2 (AMSR2) 19 GHz v-polarized data from the JAXA GCOM-W satellite. (Original data credit: JAXA)

can see by eye the overlapping sensitivities to the atmosphere (clouds, precipitation, water vapour, temperature) and the surface (ocean surface state, sea ice and its snow cover, snow, vegetation and soil moisture). The atmospheric sensitivities are currently used in all-sky radiance assimilation within ECMWF's 4D-Var atmospheric data assimilation. But the surface sensitivities remain mostly unused, in part due to the difficulty of physical modelling. These difficulties concern both the forecast model for the surface state and the translation between the surface state and the observations, which is done by the observation operator (sometimes called an observation model). The goal of making full use of the information contained in these observations is often described as 'all-sky, all-surface assimilation'.

A better use of the information in satellite observations has three benefits:

- we get more information on the atmospheric state in areas where satellite observations were previously discarded, such as sea ice, snow and land surfaces
- we get new information on poorly known variables, such as sea ice concentration, snow cover or soil moisture
- we can use the observations to help constrain, improve and develop better physical and empirical models of the Earth system.

The contention of this article is that we will make the furthest progress in the middle ground between model-led approaches and purely observation-driven approaches. Traditional data assimilation frameworks and newer machine learning approaches alike are best understood as Bayesian problems (see Box C later). What this means is that we aim to maximise the information extracted from current observations while making best use of our prior information, whether that comes from physical laws or from historical observations. This is not only a theoretical argument. A middle-ground, Bayesian-inspired approach is already helping us make rapid progress towards better use of surface-sensitive observations in our operational weather forecasting system. The example here is sea ice, but the same techniques can be used across the Earth system. The question is not whether to use physical or empirical models, but how best to combine them.

The sea ice problem

Sea ice is an attractive first target for approaches that combine machine learning and data assimilation. It is also a perfect illustration of how we can get more out of satellite observations while retaining physical modelling where possible. It covers all three areas where we would expect benefits:

Get more information on the atmosphere: In the current atmospheric 4D-Var, we have to reject any satellite observations with strong sensitivity to sea ice. As a result, there is a particularly severe information desert in southern hemisphere high latitudes. As will be illustrated in this article, putting more observations into this desert results in significant forecast improvements.

Get more information on the surface: Microwave radiance observations are already used in the ECMWF sea ice analysis. The problem is that they are finally introduced into the atmospheric analysis by a slow, roundabout, and sub-optimal route. Sea ice concentration retrievals are inferred using heuristic approaches, then incorporated into a daily analysis at the UK Met Office, then assimilated into ECMWF's OCEAN5 system, then passed to the atmospheric data assimilation with an overall delay of around 48 hours. If we can use our in-house data assimilation tools to infer the sea ice concentration directly from the radiances, we can eliminate the delay and almost certainly do a better job of inferring the state of the sea ice, too.

Use the observations to develop better physical and empirical models of the Earth system: One reason the sea ice analysis is still so far behind the atmosphere is a lack of accurate enough physical models, both to propagate the state forward in time and to model the radiance observations. In particular, the state variables that affect the radiative transfer properties of sea ice include aspects of the microstructure of the ice and snow that are not represented in current physical models.

This final aspect makes the problem just as challenging from a machine learning perspective as for data assimilation. The normal machine learning approach, known as 'supervised learning' (we might also call this 'brute force'), attempts to learn an empirical model given its known inputs and outputs. For the sea ice observation problem, we know the outputs – these are the satellite radiance observations – but the inputs, including the microstructural details of the ice and snow, are basically unknown in the absence of ice core and snow pit measurements. This is a chicken and egg problem, and its solution starts by acknowledging that we need to simultaneously learn the state of the sea ice and the empirical observation model to go from that to the observations. The problem is solved by combining both data assimilation and machine learning in a Bayesian framework with parallels to 'unsupervised' or 'generative' machine learning techniques.

To practically solve this problem, a year's worth of microwave radiance observations were paired with atmospheric profiles from the background (short-range) forecast of the Integrated Forecasting System (IFS). Machine learning tools built on top of Python, Keras and

Tensorflow were used to set up a hybrid data assimilation and machine learning framework to simultaneously learn the sea ice state along with an empirical model for the sea ice surface emissivity. Here the surface emissivity is a convenient way of representing the surface radiative transfer. Unlike standard data assimilation frameworks used for real-time forecasting, this offline approach was able to simultaneously fit a whole year of observations in one go.

Box A and its figure illustrate the resulting empirical-physical model as it applies to the surface radiative transfer for one observation. It takes physical inputs, namely the sea ice concentration, the skin temperature, and the ocean water surface emissivity, and three empirical inputs that represent the microstructural and physical state of the sea ice and

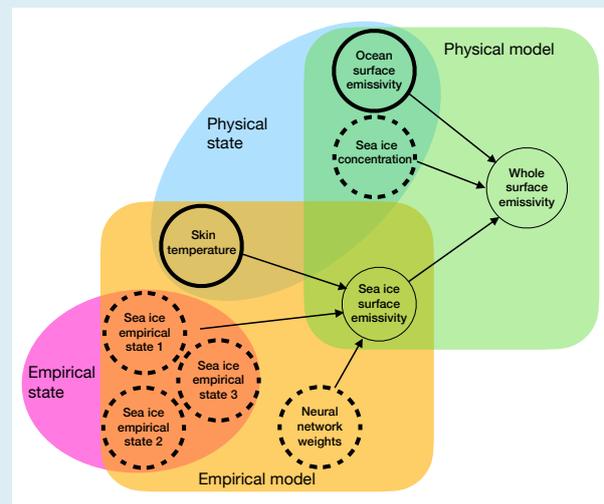
the snow on top of it. It outputs the surface emissivity of the combined ocean and sea ice scene observed by the satellite. Because of the simple physics encoded in the network, it is capable of being used to infer physical variables, namely the sea ice concentration.

The quality of the resulting sea ice concentration analysis is illustrated in Figure 2 during the rapid freeze-over of the Arctic ocean during November 2020. The OCEAN5 sea ice analysis represents the Siberian side of the Arctic ocean as open water, due to the 48-hour delay, whilst the new approach shows that it has already mostly frozen over. The new empirical-physical model is also able to closely replicate microwave observations at frequencies from 10 GHz to 89 GHz, in all seasons and both in the Antarctic and the

a The physical-empirical model

The figure shows the key part of the empirical-physical Bayesian network used for learning the sea ice state and the empirical model for the sea ice surface emissivity. The complete model component describes the mixed-surface emissivity within the field of view of a single satellite observation. Here bold outlines are independent variables, with solid outlines if they are held fixed (obtained from the ECMWF background forecast) and dashed outlines if they are to be estimated within the learning framework. Fine outlines indicate variables that are just dependent functions of other variables. The empirical part of the model estimates the sea ice surface emissivity based on the skin temperature and three empirical variables representing the state of the sea ice and the snow cover on top of it (more of this in Box B). In practice, the empirical model is a simple one-layer linear neural network, though deep nonlinear neural networks have also been trained in this framework but proved unnecessary. The physical part of the model represents the mixed surface emissivity as a linear weighted combination of the ocean surface emissivity and the sea ice surface emissivity. It is this physical emissivity contrast between sea ice and open ocean surfaces that ultimately allows the sea ice concentration to be estimated within the field of view of a satellite observation.

The broader network, not shown here, holds daily maps of the four learned state variables. It offers interpolation operators to go from the maps to the observation time and location, so that in reality it is the maps that are the true trainable variables, not the observation-space variables as shown here. The neural network weights are the same throughout



the year and at all locations, with the aim to create a universally valid empirical model for the sea ice surface emissivity. The mixed surface emissivity is input to a physical radiative transfer model that represents the atmosphere including clouds and precipitation (obtained from the ECMWF background forecast). The aim of the training is to find the maps of surface state and the neural network weights that will generate simulated observations that best fit a year's worth of real observations. The training is done simultaneously on all observations using the variational inverse Bayesian framework common to both machine learning training and 4D-Var data assimilation.

This same network fragment is extracted, along with the trained network weights, to be implemented in operational 4D-Var as part of the observation operator for all-sky, all-surface microwave observations. In 4D-Var, the learnable input variables are estimated at observation locations, but keeping the neural network weights fixed.

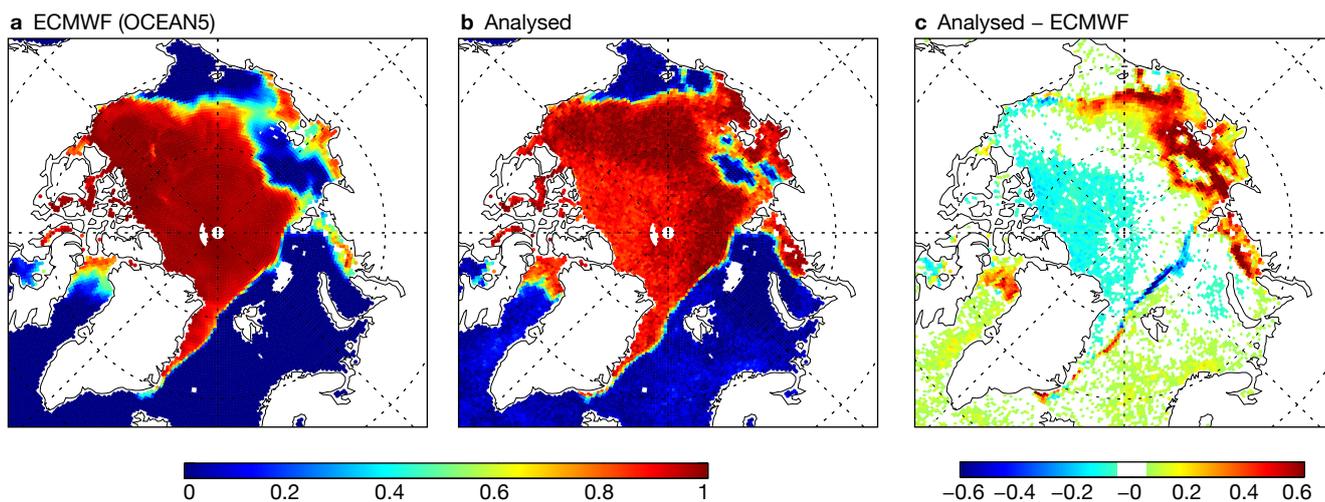


FIGURE 2 Sea ice concentration analyses, showing (a) the OCEAN5 sea ice concentration for 7 November 2020, (b) the analysis from AMSR2 observations for the same date using the new machine-learning/data assimilation framework, and (c) the difference between the two.

Arctic (not shown). The empirical variables representing the detailed microstructural and physical aspects of the sea ice and snow state are explored in Box B.

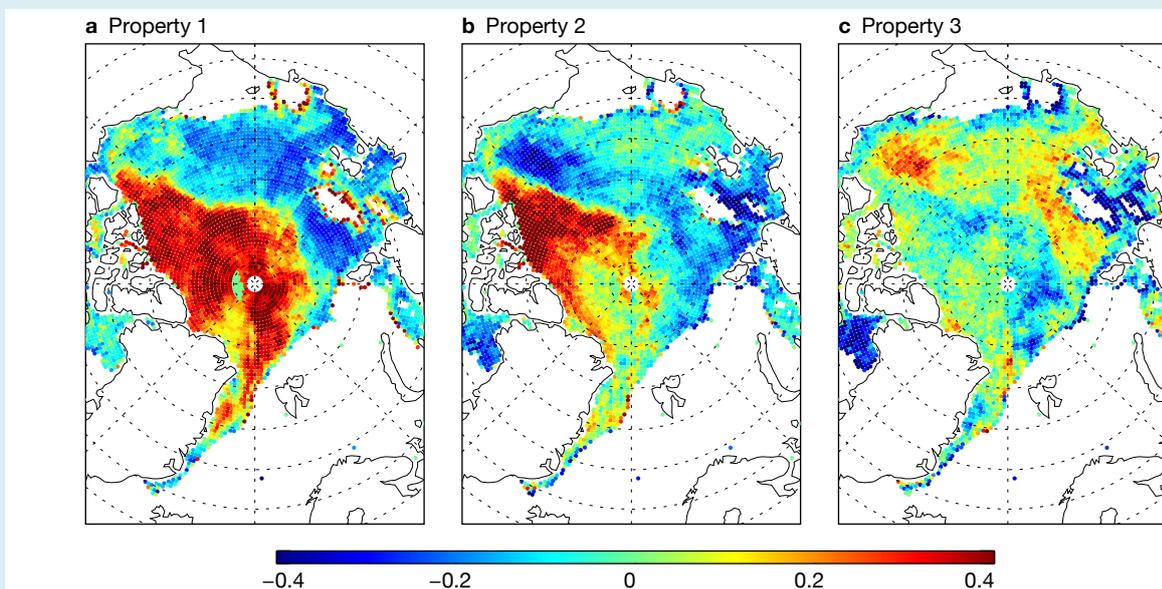
The offline analysis of the sea ice state and observation

model is on one hand an illustration of how the Bayesian approach could be used, along with empirical state variables where required, to solve all our Earth system problems using an observation-first perspective. On a

b The empirical state variables

The simultaneous training of an empirical state alongside an empirical model is what allows us to solve the chicken and egg problem. The figure shows maps of the empirical ice properties on 7 November 2020, the same date as the sea ice maps in Figure 2. Empirical property 1 follows known differences between multi-year and first-year ice. The positive (red) values are found in areas of long-lived ice. The negative values (blue) correspond to areas where

the sea ice is newly formed. New ice contains pockets of brine. This is lost over subsequent months, leaving behind air pockets that strongly scatter microwave radiation. At higher microwave frequencies, the microstructural characteristics of the snow lying on top of the sea ice become most important. The second and third empirical properties help represent this and other aspects of the sea ice and snow properties that affect the radiative transfer.



more practical level, it is also a first step in adding sea ice assimilation to the IFS.

Assimilating sea ice observations in the IFS

From the offline machine learning and data assimilation framework, we have a model (Box A) that can be plugged into the existing atmospheric 4D-Var framework of the IFS. This becomes part of the observation operator alongside the existing RTTOV (Radiative Transfer for TOVS) model for the radiative transfer of the atmosphere including clouds and precipitation. To estimate the four poorly-known model inputs, namely the sea ice concentration and the three empirical variables, 4D-Var uses an observation-space augmented control vector similar to the 'skin temperature sink variable' that has been used in satellite data assimilation at ECMWF for several decades. So, as well as estimating the state of the atmosphere in the usual way, 4D-Var also estimates the sea ice concentration and the three empirical variables at each observation location from microwave imager observations, namely the Advanced Microwave Scanning Radiometer 2 (AMSR2) and the GPM microwave imager (GMI).

An illustration of the quality of the resulting 4D-Var sea ice concentration analyses is provided by Figure 3. During the autumn and winter of 2020, the giant iceberg A-68A was drifting from its source in the Weddel sea and came close to the Southern Ocean island of South Georgia. At the time, A-68A was being tracked visually using radar and visible measurements, but the new sea ice analysis in 4D-Var has made a retrospective track of its movements that agrees well with the contemporary analysis. The figure shows A-68A approaching South Georgia on 4 December. At around 100 km long and 60 km wide, the iceberg is of similar size to the island itself. Allowing for the fact that the observation-space sea ice analysis has a relatively coarse resolution of around 40 km, it shows very similar features to the visual picture from the OLCI sensor on Sentinel-3. By contrast, the OCEAN5 analysis does not represent the iceberg.

Benefit of sea ice observations for atmospheric forecasts

Figure 4 shows the impact of adding microwave observations in sea ice areas on the temperature forecast, based on year-long testing in 2021 and 2022. Notably, this testing period is different to the training

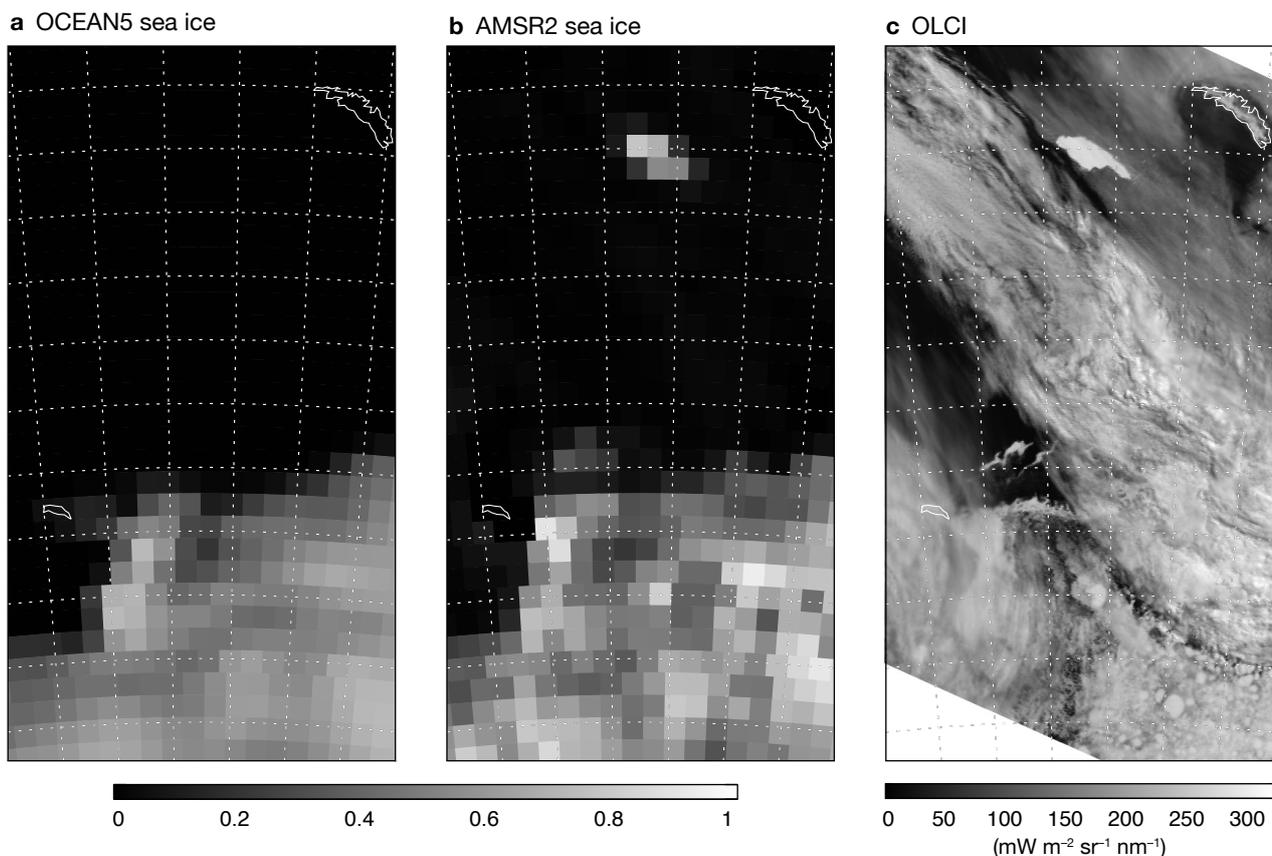


FIGURE 3 For 4 December 2020 around 12 UTC, the panels show (a) the sea ice concentration at AMSR2 observation locations obtained from the ECMWF OCEAN5 analysis, (b) the new sea ice analysis made within 4D-Var at these same locations, and (c) OLCI channel 10 visible radiance observations (Copernicus Sentinel data 2020). The island in the top right is South Georgia and the A-68A iceberg is to its left (west). Towards the bottom of the figure is the main Antarctic sea ice and towards the left part of the domain is one of the South Orkney group of sub-Antarctic islands.

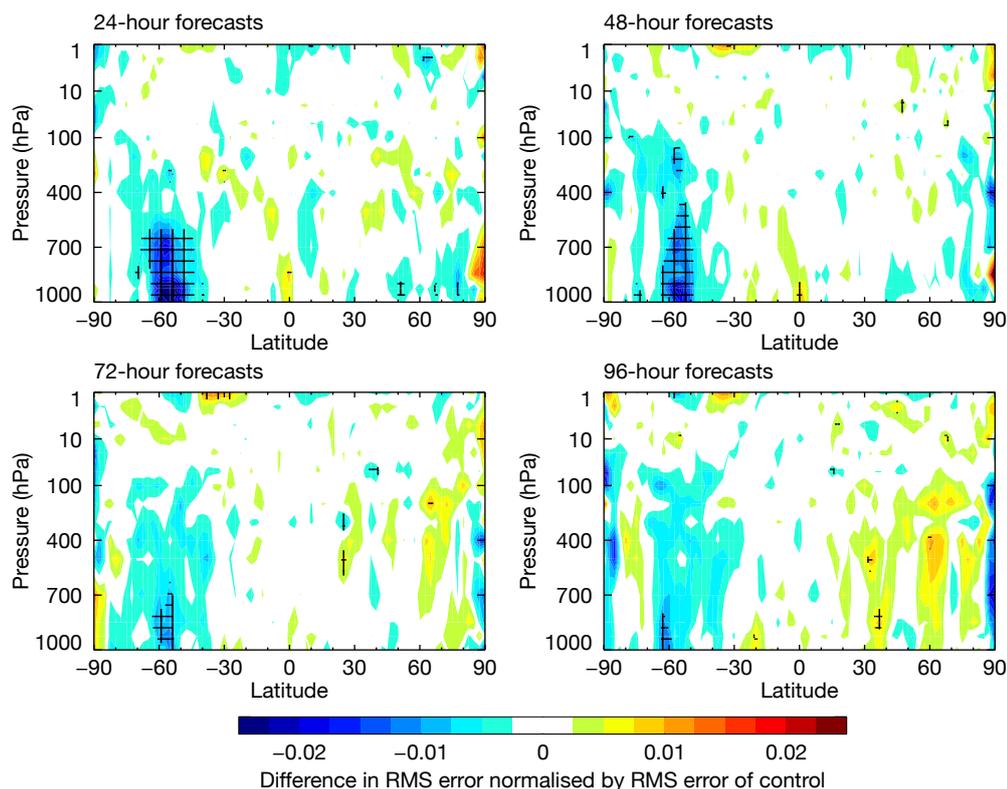


FIGURE 4 Normalised change in temperature forecast root-mean-square (RMS) error, measured against own-analysis, showing the impact of adding assimilation of the AMSR2 and GMI microwave imagers over sea ice and possible sea ice regions. Cross-hatching indicates statistical significance at the 95% confidence level. Blue areas indicate a reduction in the RMS error and hence a beneficial impact from the sea ice assimilation. Based on an entire year of testing in 2021–2022.

period used in the offline machine learning framework, which was 2020 to 2021. The new observations over sea ice areas have little impact on the Arctic forecast, likely because the year-round availability of in-situ measurements helps fill any gaps in the satellite data. But in the Southern Ocean, particularly around 50 to 60 degrees south, there is a statistically significant improvement in the forecast that lasts out to around day 4 and that spans from the surface up to the mid-troposphere. This localised impact is strong enough to also generate a statistically significant improvement in the headline southern hemisphere (20°S to 90°S) forecast scores.

Figure 5 shows the number of observations being assimilated in atmospheric 4D-Var in IFS Cycle 48r1 (control) and experimental IFS Cycle 49r1 (sea ice) from the AMSR2 sensor during June 2022. The added observations in the Arctic are limited to the relatively small area of summer sea ice. By contrast, a much larger area of the Southern Ocean, including both the sea ice and its surrounding areas, gains microwave imager observations for the first time. The new sea ice assimilation has relatively large observation errors and the added ‘safety valve’ of being able to adjust the surface emissivity to fit the observations. A side effect is that it is possible to assimilate for the first time ‘cold air outbreak’ regions over the ocean in the vicinity of the sea ice (Forbes et al., 2016). It appears that most of the atmospheric forecast benefit comes from these new observations, rather than those over the sea ice itself. However, as the modelling improves and observation errors are reduced in future years, it is

expected that observations over sea ice will provide atmospheric information, too.

A further detail of Figure 5 is that the density of observations is up to six times higher in the polar regions than at midlatitudes. This is due to the polar sun-synchronous orbit used by most meteorological satellites, which in this case takes the AMSR2 sensor over part of each polar cap every 100 minutes. Further development of the sea ice and snow analysis could release a huge amount of high temporal frequency data from polar-orbiting satellites in polar regions.

Future developments

The sea ice assimilation described in this article is intended to become operational with Cycle 49r1 of the IFS in the middle of 2024. For the moment, the main practical benefit will be improved forecasts in the Southern Ocean. But the next stage of development is to assimilate the sea ice concentration analyses into the ocean data assimilation component using an outer-loop coupling data assimilation framework. This will also assimilate new sea-surface temperature (SST) retrievals from the microwave observations, testing an approach that could ultimately eliminate the need to use external SST and sea ice concentration (SIC) products in the ocean analysis, along with the 48-hour delay this entails. That next stage is in development and will be documented elsewhere.

The techniques that have enabled us to get started with sea ice assimilation are also immediately applicable to

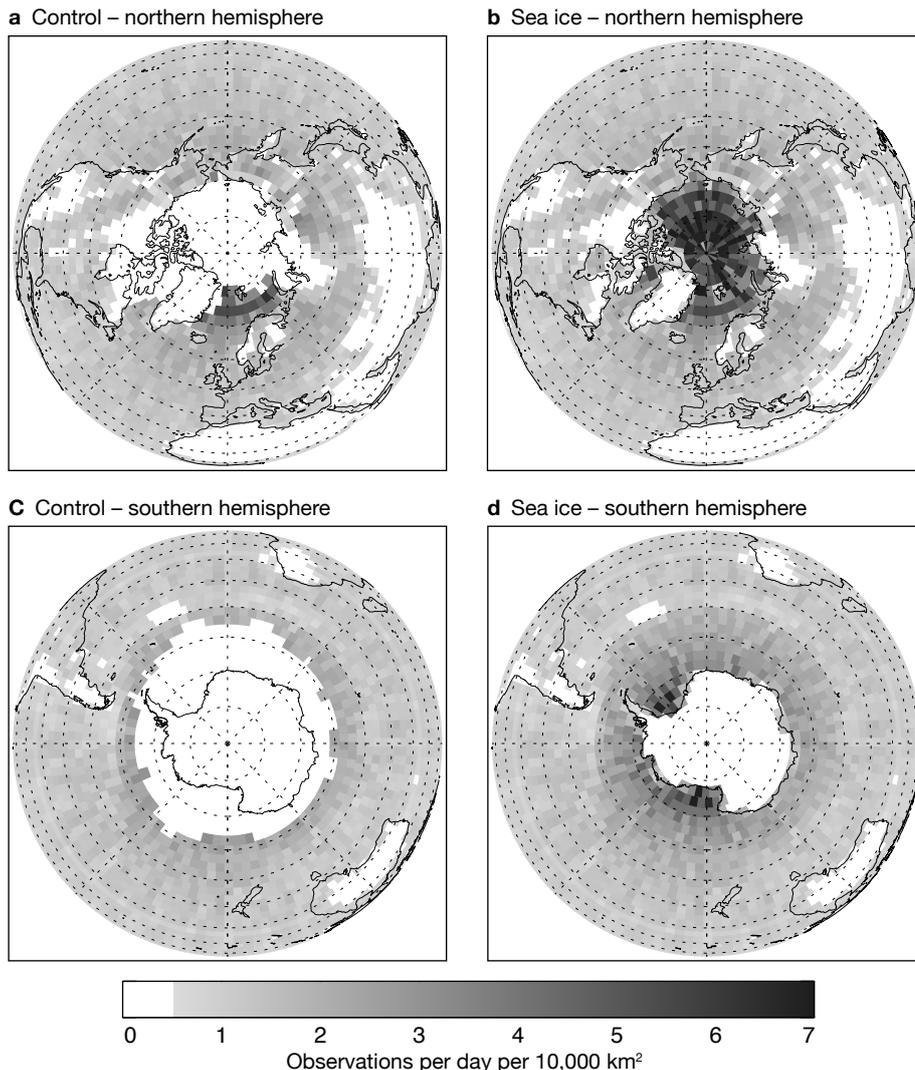


FIGURE 5 Number of observations per day per 10,000 km² assimilated from AMSR2 (in its 37 GHz v-polarized channel) between 1 June and 5 June 2022 in (a) the Cycle 48r1 control configuration in the northern hemisphere, (b) the Cycle 49r1 sea ice configuration in the northern hemisphere, (c) the Cycle 48r1 control configuration in the southern hemisphere, and (d) the Cycle 49r1 sea ice configuration in the southern hemisphere.

snow over land and, with further development, the whole land data assimilation problem, including soil moisture and vegetation. There is scope to extend the empirical-physical modelling to a wider range of microwave frequencies and viewing geometries, aiming to start assimilating not just microwave imaging channels over sea ice (10 to 89 GHz) but also microwave sounding channels, which provide a leading contribution to forecast accuracy over surface types that are easier to assimilate. One aim is to move the surface radiative transfer modelling away from the use of surface emissivity and towards a more physical representation of scattering and absorption within the sea ice and snow pack. And when improved forecast models of sea ice and snow are coupled into atmospheric 4D-Var, the empirical modelling can be re-trained using these additional physical inputs. The long-term aim is to increasingly impose physical constraints and to minimise the use of empirical models and empirical state variables.

For an even wider perspective, we can go back to the Bayesian viewpoint that unifies machine learning and data

assimilation. It is this framework that makes the hybrid sea ice state and model estimation possible, and this framework will lie behind its future extensions into the whole area of Earth system data assimilation.

The Bayesian approach treats both prior knowledge and observations in a theoretically correct way. This means we can obtain the best possible information on the current state of the Earth system and its modelling components (see Box C). Prior knowledge can be encoded using the physical equations of the Earth system that we already know. But we can also encode the areas of the Earth system that we know less well using empirical model components that can be learned from observations.

If we could completely implement the Bayesian network framework, it would help us balance the information gleaned from the limited field measurements that underpin many current Earth system parametrizations with the vast and continuing amount of information coming from satellite observations. It would help us derive much more sophisticated physical parametrizations that would work more universally across the globe. The role of the scientist would move away from finding heuristic or regression

C

The Bayesian approach

Bayes' theorem comes out of fundamental rules of statistics and shows how to combine prior information with new observations to gain an improved knowledge of the world. Bayes' theorem is formulated in terms of 'prior' and 'posterior' probability distribution functions and is hard to solve directly for high-dimensional problems like the Earth system. Variational data assimilation turns Bayes' theorem into a tractable method for ingesting observations by using the assumption of Gaussian error distributions. But the applications of Bayes' theorem are far wider, both philosophical and practical. One key extension is the 'Bayesian network', which enables us to break up huge probabilistic problems into simpler components if we know (or can learn) the statistical dependencies between variables. The hybrid empirical-physical model in Box A is also an illustration of the Bayesian network describing sea ice and ocean surface

radiative transfer, where the arrows represent the dependencies, and the circles represent statistical variables. Again, it is solved by making simplifying assumptions and by doing a variational minimisation to fit new observations, just like 4D-Var or indeed most neural network training. Within these networks, known physics can be represented using fixed physical equations. Unknown physics, such as the sea ice surface emissivity, can be represented using an empirical model such as a neural network. Typical 'brute force' machine learning throws away all the prior physical knowledge and attempts to learn from the new observations alone. Unless these observations contain all knowledge, then by Bayes' theorem, brute force machine learning cannot possibly provide as much posterior knowledge of the world as we can gain by using a learning process that includes prior knowledge.

models based on limited field observations. Instead, the job would be to encode the physics we already know, to correctly describe our confidence in this knowledge, and to make sure that Earth observations are used as completely as possible. This is not just an issue for the 'newer' aspects of Earth system studies, such as sea ice, snow or vegetation, but for the 'older' empirical aspects of atmospheric models, such as cloud and precipitation parametrizations and sub-grid scale physics. Although it would be a daunting task to break apart our current systems and re-implement them in a unified Bayesian framework, many of the tools already exist and are being used for machine learning; indeed, they underpin its current success.

In comparison to a brute-force machine learning approach, a careful hybrid of the physics we already know with all the new and evolving knowledge that comes from observations will always give us a better understanding, not just of the Earth system state but also the physical and empirical models that represent its evolution. It is only in a few places in the world, such as at ECMWF, that all the different components of observations, prior physical knowledge, and Bayesian learning methods (including data assimilation and machine learning) can be brought together to generate the highest-quality Earth system analyses and forecasts.

Further reading

de Rosnay, P., P. Browne, E. de Boissésou, D. Fairbairn, Y. Hirahara, K. Ochi et al., 2022: Coupled data assimilation at ECMWF: Current status, challenges and future developments, *Q. J. R. Meteorol. Soc.*, **148**, 2672–2702. <https://doi.org/10.1002/qj.4330>

Forbes, R., A. Geer, K. Lonitz & M. Ahlgrimm, 2016: Reducing systematic errors in cold-air outbreaks, *ECMWF Newsletter No. 146*, 17–22. <https://doi.org/10.21957/s41h7q7l>

Geer, A.J., 2021: Learning earth system models from observations: machine learning or data assimilation? *Phil. Trans. R. Soc. A*, **379**. <https://doi.org/10.1098/rsta.2020.0089>

Geer, A.J., 2023: Simultaneous inference of sea ice state and surface emissivity model using machine learning and data assimilation (in preparation).

Geer, A.J., 2023: Joint estimation of sea ice and atmospheric state from microwave imagers in operational weather forecasting (in preparation).

What next for Magics visualisation?

Sylvie Lamy-Thépaut, Stephan Siemen, James Varndell

Magics (Meteorological Applications Graphics Integrated System), ECMWF's software library specifically designed for data visualisation, has undergone substantial change since its initial release in 1985. This article aims to provide a history of Magics and highlights how Magics has continually adapted to meet the needs of scientists, specifically in terms of generating high-quality weather maps, effectively handling the ever-increasing data resolution, and adopting new standards and technologies. The article concludes with an exploration of current challenges and how Magics will evolve into the visualisation component of ECMWF's ambitious new open-source project, earthkit.

The 1980s – bringing colour to weather maps

The visualisation of forecasts and observational data is vital to be able to comprehend the vast amounts of data involved in developing and operating a forecast model. Therefore, ECMWF invested from the beginning in good visualisation facilities to allow scientists and users alike to work with and develop the forecast model and analyse their results: two-dimensional maps for the forecasts and graphs for some other data. In the 1980s, this meant mainly printed maps, as computer screens were mostly too low-resolution to display large and complex maps. One focus of Magics development was to make extensive use of colours to allow the overlay of different parameters on maps.

From the beginning, developers acknowledged the necessity of creating an easily expandable solution, foreseeing increasing demands and requests from users. However, during the early 80s, expanding a user interface (API) commonly involved adding parameters to subroutines. Recognising the inherent difficulties in maintaining this approach, the Magics team made a deliberate decision to pursue an alternative strategy. They opted to maintain a concise list of subroutines while providing an extensive set of parameters. This design choice empowered users to customise their plots effectively and streamlined the documentation process.

The first plot ever produced by Magics was the metgram, which helped users to see the evolution of key weather variables for particular locations. At that time,

an open-source solution was not viable. Instead, Member States financed the development of software by ECMWF to support their daily work. Magics itself relied on CONICON, a contouring package that required a paid licence. This algorithm was necessary to generate smooth lines, despite the low resolution of the data at that time. Magics was one of the first software packages able to visualise the new emerging World Meteorological Organization (WMO) standards for gridded forecast data and observations, GRIB and BUFR. Figure 1 shows an early example of the output provided by Magics.

In 1993, Magics became the graphical kernel of Metview, the ECMWF workstation for researchers. For the first time, scientists could quickly visualise data on their screen, zoom in on specific areas, interactively select lines to perform further actions, such as visualising cross sections, and see the value of fields below their cursor.

Early 2000s – adaptation to the web & Python

In the early 2000s, it became clear that key technologies around the creation of weather maps had changed. The model resolution was increasing fast, Metview was more widely used, the web and large screens had replaced printers as the main medium to present maps

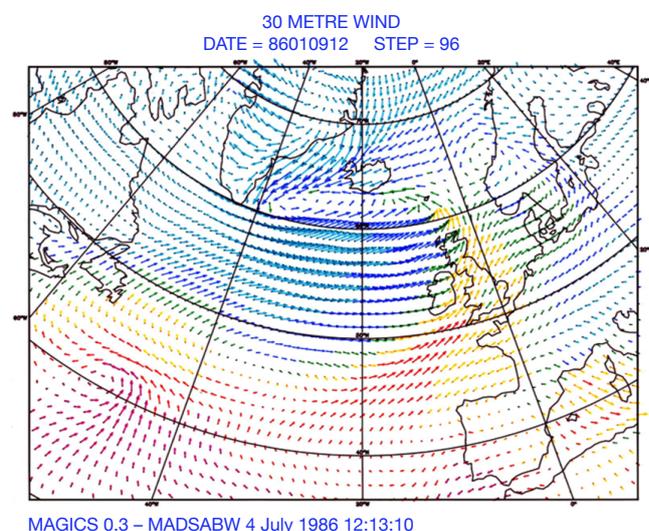


FIGURE 1 A Magics-generated map from 1986, showing the ability to colour 30 m wind arrows by 850 hPa temperature. The map was provided by Jens Daabeck, the first developer of Magics.

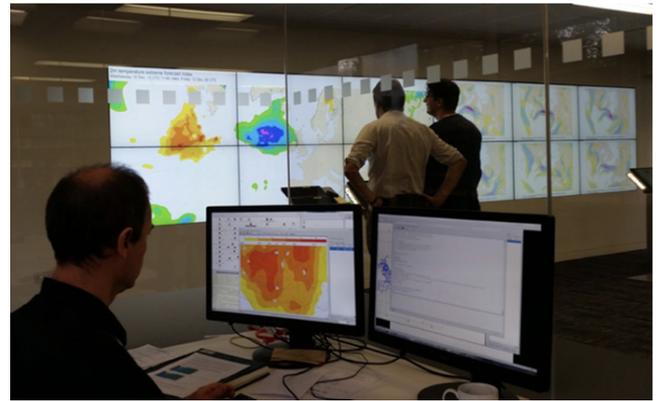


FIGURE 2 The changing medium to work with weather maps. Until the early 2000s, printed maps were the main way to work with forecast data, as can be seen in the left-hand photo of the ECMWF MetOps room in the early 2000s. Nowadays, interactive screen setups have fully replaced printed maps, as shown in the right-hand photo of the ECMWF Weather Room in 2018.

(see Figure 2 for examples of working with printed maps and computer screens). At the same time, users moved their data analysis and visualisation work away from Fortran to Python and web services. Combined with the fact that Magics had been mainly written in Fortran and became harder to maintain, the decision was made to rewrite Magics in C++ and offer additional interfaces to C, Python, XML and JSON to create plots. To smooth the migration, it was decided to ensure backward compatibility with the previous API and preserve the concept of a limited number of action routines enhanced by a comprehensive set of parameters. As a temporary name, 'Magics++' was adopted to differentiate itself from earlier versions. CONICON was abandoned in favour of a freely available algorithm, which opened the door to the journey towards open-source development, making Magics one of the first ECMWF software packages distributed under the open-source Apache Licence in 2005.

The rewrite offered the possibility to incorporate additional features, such as support for NetCDF and ODB (ECMWF's own format to handle observational data), inclusion of political boundaries, and the implementation of thermodynamic diagrams. Moreover, new graphical formats like PNG, PDF or SVG were introduced. Magics was for the first time using an external open-source library, *cairo*, to generate these different formats. A new Python interface was implemented, aligning with the philosophy of the *Metview* macro language favoured by users at that time. Figure 3 shows an example of a weather plot shown by *Metview*.

The new version became the graphical kernel of *ecCharts*, a highly interactive web application for weather forecast data developed in 2010. This time, users gained the ability to zoom in and pan across the latest forecast in their native resolution directly within their web browser. Today, *ecCharts* provides on-demand

visualisation for over 250 parameters, i.e. more than 2 TB of forecast data every day. To enable this, significant optimisation efforts were undertaken, and an extensive list of styles was created. Subsequently, this library of styles was ported back to Magics, enabling the automatic selection of an appropriate visualisation based on metadata. This automatic styling feature is extensively used in *SkinnyWMS*, a lightweight implementation of an Open Geospatial Consortium (OGC) Web Map Service standard. This lightweight implementation can identify GRIB or NetCDF data and allows users to quickly visualise, zoom, and pan through their data. This functionality can be particularly valuable when working in a *JupyterLab* environment.

The 2020s – new challenges

However, the landscape has changed, and today Magics faces new challenges. Meteorology has expanded to encompass a broader community, including the fields of climate and oceanography. To foster collaboration and exchange within this larger community, there is a need for a shorter learning curve to be able to create complex visualisations, and a shift towards a more pythonic approach. Recognising these evolving requirements, the next Magics generation will address these needs by offering an intuitive interface that empowers newcomers to interact more seamlessly and effectively with the rest of the community.

To achieve this, the new Magics will be the visualisation component of *earthkit*, an exciting new open-source project led by ECMWF. As a successor to *Metview*, *earthkit* will provide powerful and convenient Python tools for speeding up weather and climate science workflows by simplifying data access, analysis and visualisation. The geospatial visualisation component of the project, *earthkit-maps*, will provide the extensive range of features offered by Magics through a new Python API, and will be built upon a new backend based on well-established open-source packages like

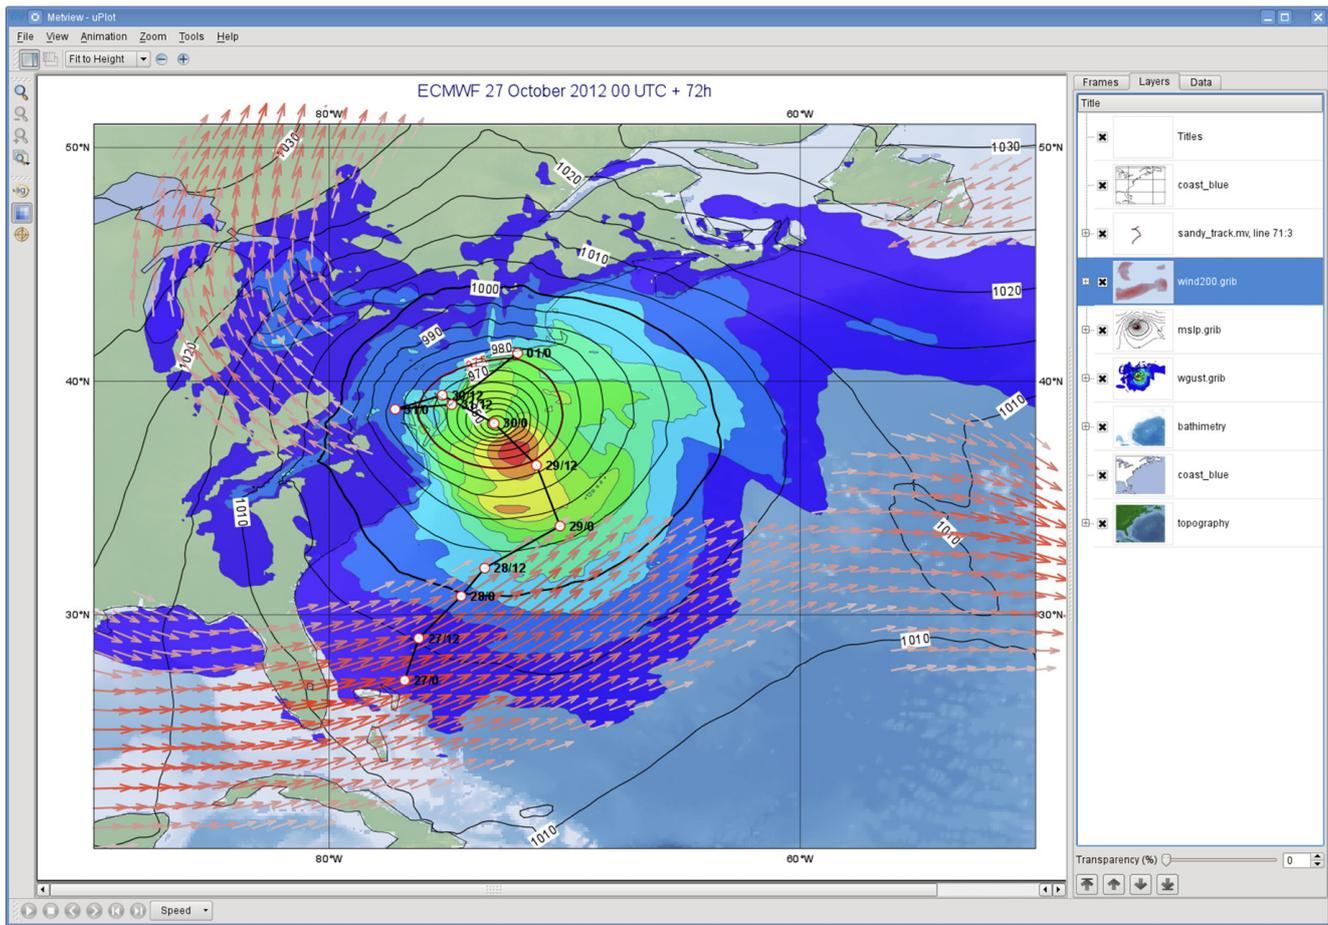


FIGURE 3 Example of Metview capabilities for forecast analysis using Magics, in this case showing Hurricane Sandy in 2012.

matplotlib and Plotly. This shift toward widely adopted tools will ensure compatibility and ease of integration within the data visualisation ecosystem.

Just like the current version of Magics, earthkit-maps aims to provide an extensive range of features to facilitate the visualisation of meteorological and climate data.

Maintaining a delicate balance between speed and high quality, earthkit-maps will continue to deliver efficient rendering without compromising on visual output. To achieve this, earthkit-maps will take advantage of MIR, the ECMWF Meteorological Interpolation and Regridding package, enabling it to select the most suitable interpolation technique. This approach effectively reduces the data size while preserving essential details, resulting in optimised visualisations. By default, the geographical area will fit the specific domain of the data, and an automatic styling mechanism, imported from Magics, will select the most appropriate visualisation techniques based on the data characteristics. This will ensure that the visualisations

are both meaningful and informative. Earthkit-maps will continue to offer a wide range of convenience features such as automatic titles, layouts and labels based on metadata, while also providing full user freedom to customise visualisations with ease.

Based on well-established packages, the transition to earthkit-maps is designed to be smooth for Python users. And support will be provided to existing Magics and Metview users to facilitate their migration process. To assist with this transition, dedicated Jupyter notebooks will be created, reproducing the Magics and Metview example galleries. These notebooks will serve as practical examples, showcasing the capabilities of earthkit-maps. In addition, comprehensive guidelines will be developed to provide step-by-step instructions for Magics and Metview users looking to migrate to earthkit-maps. These guidelines will outline the necessary steps, highlight key differences and improvements, and help users leverage the full potential of earthkit-maps. Figures 4 and 5 provide examples of the code to produce weather maps using the current version of Magics and earthkit-maps, respectively.

```

import metview as mv
YEARS = {
    1993: "Normal conditions",
    1997: "El Niño",
    1998: "La Niña",
}
# getting forecast data from CDS
import cdsapi
c = cdsapi.Client()
filename = "sst_era5_mnth.grib"
c.retrieve(
    "reanalysis-era5-single-levels-
monthly-means",
    {
        "product_type": "monthly_averaged_
reanalysis",
        "variable": "sea_surface_temperature",
        "year": list(YEARS),
        "month": [12],
        "day": "01",
        "time": "00:00",
        "area": [90, -180, -90, 180],
        "grid": [0.25, 0.25],
        "format": "grib",
    },
    filename,
)
# read data from file
data = mv.read(filename)

# define coastlines
coast = mv.mcoast(
    map_coastline_land_shade="on",
    map_coastline_land_shade_
colour="charcoal"
)

# define the view : geographical area: [-20, 100,
20, -60]
view = mv.geoview(
    map_area_definition="corners", area=[-20,
100, 20, -60], coastlines=coast
)

# define a 3x1 layout
page_0 = mv.plot_page(top=0, bottom=30, left=5,
right=95, view=view)
page_1 = mv.plot_page(top=33, bottom=63, left=5,
right=95, view=view)
page_2 = mv.plot_page(top=66, bottom=96, left=5,
right=95, view=view)
dw = mv.plot_superpage(pages=[page_0, page_1,
page_2])

# define the style
style = mv.mcont(
    legend="on",
    contour="off",
    contour_level_selection_
type="interval",
    contour_max_level=32,
    contour_min_level=15,
    contour_interval=1,
    contour_label="off",
    contour_shade="on",
    contour_shade_colour_method="palette",
    contour_shade_method="area_fill",
    contour_shade_palette_
name="colorbrewer_Spectral_17",
)

# define title
titles = []
short_name = "<grib_info key='shortName'/>"
date_info = "<grib_info key='valid-date'
format='%Y %b'/>"
for conditions in YEARS.values():
    titles.append(mv.mtext(text
lines=f"{conditions} - ERA5 {short_name}
monthly mean: {date_info}",
text_font_size=0.35))

# generate plot and save the result in a pdf file.
mv.setoutput(mv.pdf_output(output_name="sst_era5_
elnino_map"))
mv.plot(dw[0], data[0], style, titles[0],
dw[1], data[1], style, titles[1],
dw[2], data[2], style, titles[2])

```

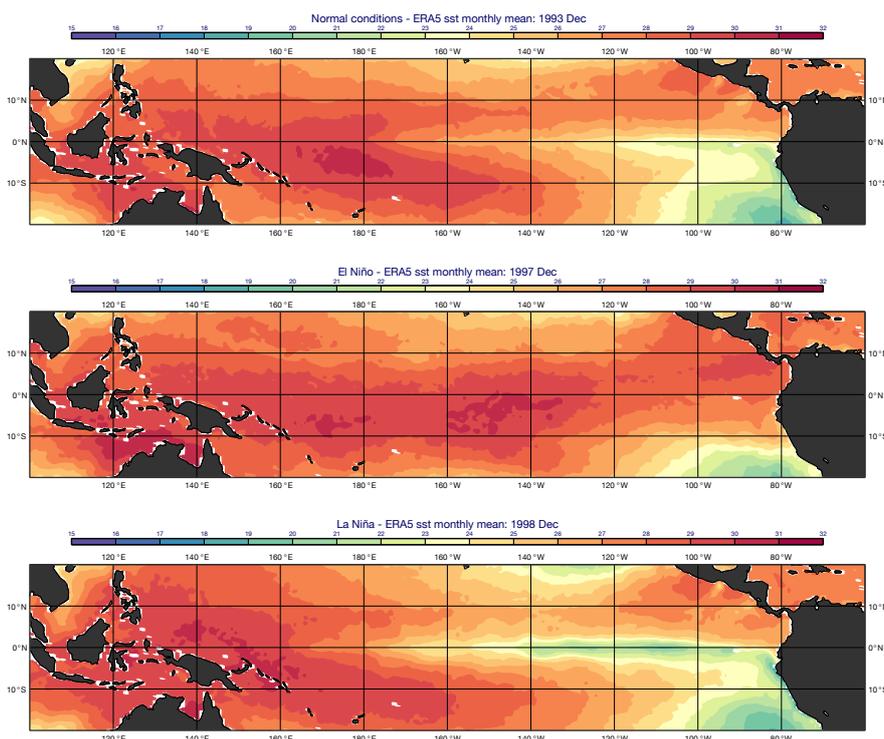


FIGURE 4 Producing three maps using the current version of Magics, showing monthly mean sea-surface temperatures (SST) in the tropical Pacific Ocean in December 1993, December 1997 and December 1998, from ECMWF's ERA5 reanalysis.

```
import earthkit
YEARS = {
    1993: "Normal conditions",
    1997: "El Niño",
    1998: "La Niña",
}

# getting forecast data from CDS
data = earthkit.data.from_source(
    "cds", "reanalysis-era5-single-
    levels-monthly-means",
    {
        "product_type": "monthly_averaged_
        reanalysis",
        "variable": "sea_surface_
        temperature",
        "year": list(YEARS),
        "month": "12",
        "time": "00:00",
        "area": [20, 100, -20, -60],
        "grid": [0.25, 0.25],
    },
)

# define the plot, the geographical area and a 3
rows layout
chart = earthkit.maps.Superplot(
    domain=[100, 300, -20, 20],
    rows=3,
)

# define the style
style = earthkit.maps.styles.Contour(
    colors="Spectral_r",
    levels=range(15, 33),
    units="celsius"
)

# generate plot
chart.plot(data, style=style)

#Add the coastlines
chart.land(color="#555", zorder=10)
chart.gridlines(xlocs=range(-180, 180, 20),
ylocs=range(-20, 20, 10))

#Add the title
for subplot, conditions in zip(chart, YEARS.
values()):
    subplot.title(f"{conditions} -
    ERA5 SST monthly mean: {{time:%B
    %Y}}")

#Add the legend
chart.legend(location="bottom", ticks=range(15,
33))

#Generate the plot and save the result in a pdf
file.
chart.save("era5-el-nino.pdf")
chart.show()
```

Sea surface temperature in December of 1993, 1997 and 1998

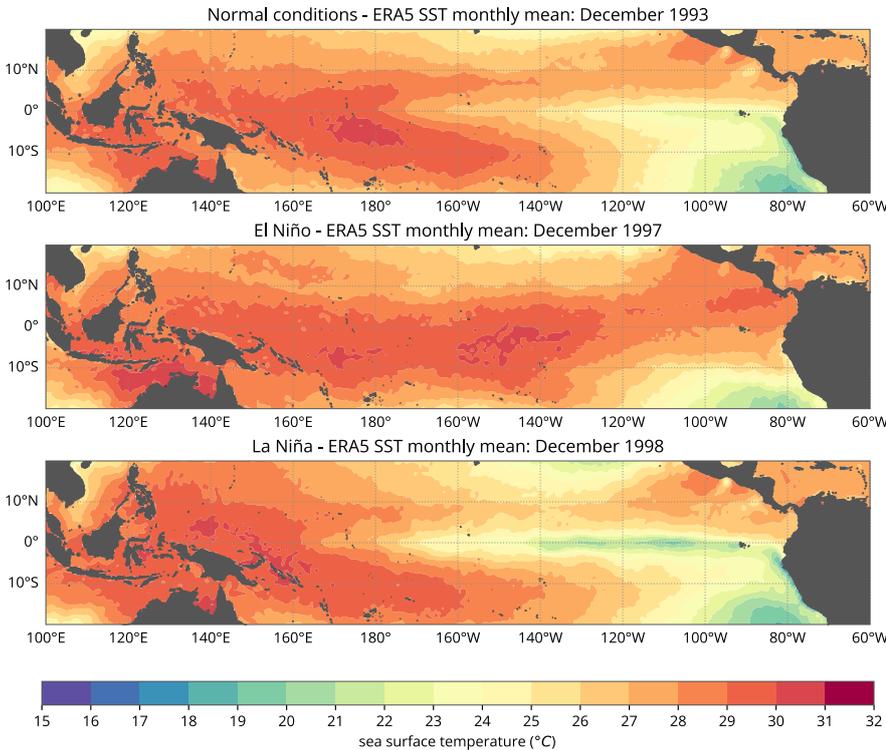


FIGURE 5 Producing three maps using the future earthkit-maps package, showing the same range of monthly mean SST in the tropical Pacific Ocean as in Figure 4.

Conclusion

Magics has evolved from its early days as an expandable tool with a unique approach to parametrization. It has overcome challenges related to data resolution and licensing, and it now faces new challenges in accommodating a larger community

while ensuring a seamless learning curve and compatibility with other open-source packages. With its continued focus on ease of use and automatic data understanding, earthkit-maps will be an asset in the field of data visualisation. It is currently in the beta stages of development with a release target of 2024.

Thirty years of the ecgate service at ECMWF

Paul Dando, Dominique Lucas, Umberto Modigliani

The two leading and most well-known IT systems at ECMWF are the high-performance computing facility (HPCF) and the Data Handling System (DHS). Another system, commonly known as ‘ecgate’, was available to thousands of users in ECMWF Member and Co-operating States for nearly 30 years. It was a general-purpose standalone system and the main access point to ECMWF’s Meteorological Archival and Retrieval System (MARS). It also provided computing resources enabling the post-processing of ECMWF forecast data to produce tailored products addressing specific Member and Co-operating State needs. The move of the ECMWF data centre to Bologna, Italy, and the subsequent decommissioning of the data centre at the Reading headquarters in the UK brought with it an end to the standalone ecgate system. Access to the final version of the system was closed on 15 December 2022. In this article, we look back at the history of the ecgate system, illustrate its usage, and highlight the importance of the service it provided. We also explain how this important managed service, available to Member and Co-operating State users, is now provided by part of the new HPCF in Bologna, called the ECGATE Class Service (ECS).

Key dates in the ecgate service

The ecgate service changed a lot over time. Here we present a list of significant events related to it.

1993 – First new Unix service

The first general Unix service was offered to ECMWF Member State users on 1 September 1993. This service was introduced to replace the interactive service provided on a CDC Cyber NOS/VE system, in line with the global move towards Unix systems, as users became more accustomed to the Unix environment. The first server was an SGI Challenge with two 150 MHz R4400 CPUs and was shared by Member State users and staff from ECMWF’s Operations Department. The main purpose of this first system was as a gateway for users to access services provided by other ECMWF systems. As such, it acted as a front end to the HPCF: it enabled batch job submissions, editing of files and transfer of files to and from the HPCF, Member State systems or the user

archive (ECFILE). Access to MARS or ECMWF’s graphical library (Magics) was unavailable on the initial ecgate system. The system was only accessible from the national meteorological services of our Member States, and users logged in with a password.

1995 – Dedicated Member State Unix system

During the summer of 1995, a new dedicated Unix system called ‘ecgate’ was made available to Member State users. In addition to the previous service, the new system also offered a local batch service. A new Unix MARS client was now also available, enabling Member State users to download and process data from the archive directly on ecgate.

With the increased use of Unix and the Transmission Control Protocol/Internet Protocol (TCP/IP) suite, both at ECMWF and within many Member States, security aspects of remote access to the Centre’s Unix services had become a concern. On the recommendation of ECMWF’s Technical Advisory Committee (TAC), it was decided to implement a strong authentication system to control interactive login access to the ECMWF Unix systems.

The solution chosen was based on the SecurID time-synchronised smart card from Security Dynamics Inc. SecurID cards were credit-card-sized, time-synchronised tokens that displayed a new, PIN-protected, unpredictable authentication code every minute, which users used instead of a password to gain login access to ECMWF Unix systems. Member State Computing Representatives were provided with some administrative control, for example, assigning spare cards or re-setting PIN codes for cards assigned to users in their country. These administrative tasks were performed on ecgate. The introduction of SecurID cards enabled ECMWF to withdraw all password-protected interactive access to the Centre’s Unix systems.

With the opening of this new Unix service to all Member State users and the introduction of the SecurID cards, we also closed direct login access to the HPCF. Member State users were therefore forced to first log in to ecgate. This change enabled us to provide independent Unix systems to Member State users and ECMWF staff.

1996–1999 – Internet access

Following ECMWF Council approval, access to ecgate via the Internet was first opened on demand at the end of

December 1996. This was followed in March 1999 by the lifting of all restrictions on outgoing Internet connections from ecgate. Internet access was further extended to all Member State users in November 2002. These moves opened access to the ECMWF computing environment to a much larger user community, including users from organisations other than national meteorological services.

2001 – Opening access to Co-operating States

In 2001, ecgate access was also granted to users from ECMWF Co-operating States. Until then, these users could only access data from the ECMWF MARS archive using a special MARS client that enabled remote access to the data. Direct access from ecgate enabled them to access and process ECMWF data more efficiently, particularly ensemble forecast (ENS) data, which was becoming more popular then. This was an example of the ‘bring the compute work close to the data’ paradigm for these users.

2002 – ECaccess service

The ECaccess software was developed by ECMWF and made available to ECMWF users. It included many functionalities to facilitate the usage of ECMWF’s computing resources. It had interactive login facilities, interactive file transfer capabilities, remote batch jobs and file handling, and unattended file transfer options. All these ECaccess services were managed on the ecgate server. Amongst others, Météo-France used the ECaccess infrastructure to enable its researchers to remotely submit experiments with its ARPEGE weather prediction model on ECMWF computing systems transparently.

2005 – Member State framework for time-critical applications

With increased user demand for running Member State operational work at ECMWF, and following approval by the ECMWF Council at the end of 2004, a framework for running Member State time-critical applications was set up at ECMWF.

- The first option enabled users to submit batch work associated with some ECaccess ‘events’. These jobs would be started when the ECMWF operational suite reached the relevant event, e.g. the 10-day high-resolution forecast (HRES) data at 12 UTC being available. Most of these jobs would run on ecgate.
- The second option enabled a Member State or consortium to run its own time-critical applications, using the ECMWF workload manager, SMS at the time. These activities were managed on ecgate.

The ECMWF shift staff monitored all these activities.

2018 – ecgate virtual machine to support ecFlow activity

HARMONIE is a mesoscale analysis and forecasting limited-area model used by several ECMWF Member States. Over the years, the HARMONIE research community using ecgate grew to over 100 users. Each user would submit one or more experiments from ecgate. By 2018, HARMONIE used the ECMWF workload manager called ecFlow, with ecFlow servers running on ecgate. To reduce interference of this intense ecFlow activity on ecgate with other ecgate activities, a new additional ecgate service was opened on a virtual machine. This machine managed only ecFlow activities, including HARMONIE experiments and some Member State time-critical applications.

2020–2022 – ecgate system freeze and final termination of service

After 2020 and until its final switch-off on 15 December 2022, ECMWF carried out only essential maintenance on ecgate so that resources could be focused on migrating the data centre to Bologna. A new service to replace ecgate, called the ECGATE Class Service (ECS), has now been implemented in Bologna as part of the Atos HPCF and provides similar functionality, while being better integrated with the HPC managed service.

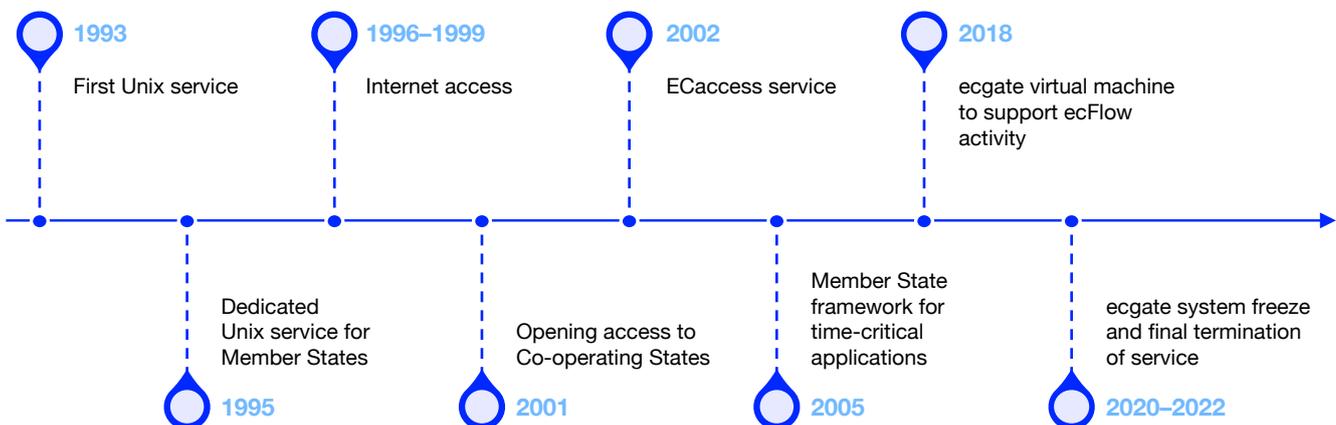


FIGURE 1 Timeline of the evolution of the ecgate service.

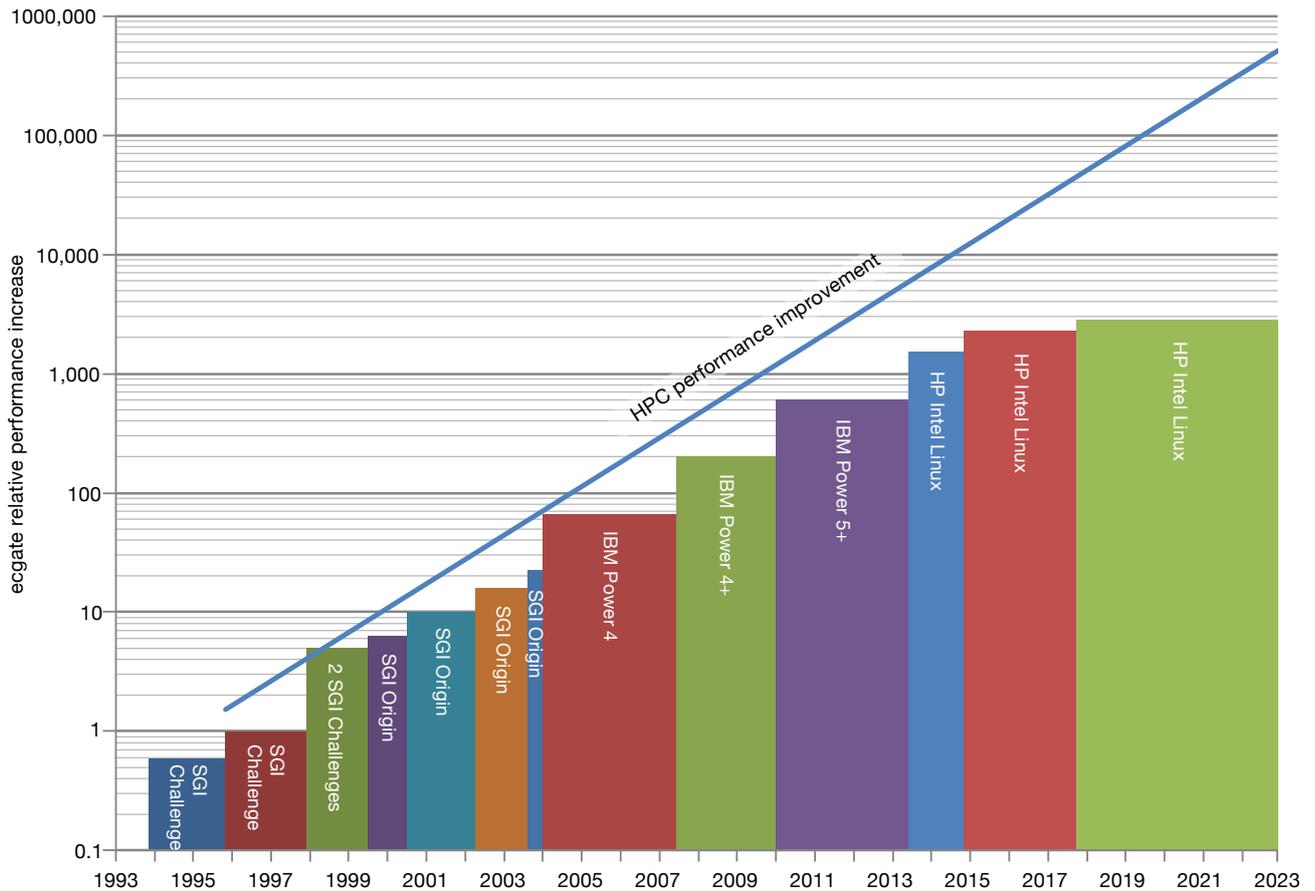


FIGURE 2 ecgate relative performance upgrades compared to the improvement in HPC performance.

A timeline showing these events is presented in Figure 1.

Upgrades of ecgate

The ecgate system was upgraded multiple times, as illustrated in Figure 2. These upgrades included operating system upgrades and hardware upgrades. For example, after 2013 different versions of hardware with the Linux operating system were used. Although the performance increase of ecgate broadly follows the HPC historical growth, the acquisition strategy and process for introducing new ecgate systems were independent of replacing the HPCF.

The most crucial consideration for ecgate was that the system could handle the burst of Member State users' activities between 05:00 and 08:00 UTC and between 17:00 and 20:00 UTC, as illustrated in Figure 3. These two activity peaks corresponded to Member State users' time-critical option 1 jobs, accessing the latest operational HRES and ENS forecast data. These high CPU utilisations were mainly due to the MARS interpolation of real-time data. It was, therefore, essential that the ecgate system could cope with vertical or horizontal resolution increases of ECMWF's numerical weather prediction (NWP) model, implemented as part of upgrades of the Integrated Forecasting System (IFS).

A typical workday on ecgate

By 2020, we had over 3,000 registered Member and Co-operating State users with access to ecgate.

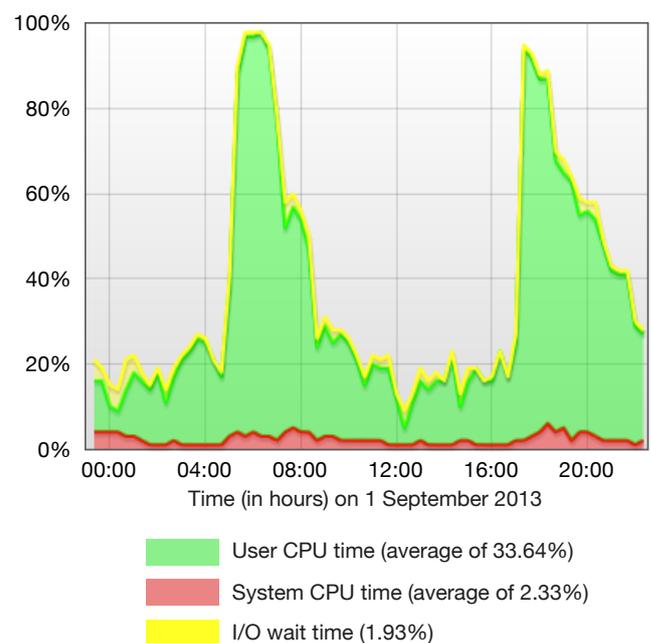


FIGURE 3 CPU utilisation on one node of ecgate on a day in 2013.

Of these, over 300 users were very regularly active. The main activity on ecgate was accessing ECMWF's latest forecast data via the time-critical option 1 service. About 150 users took advantage of this service and submitted over 1,500 jobs daily.

Use of ecgate was also the most common way registered users accessed data in ECMWF's MARS archive, including historical operational data, reanalysis data and, more recently, data from the EU-funded Copernicus Atmosphere Monitoring Service (CAMS) implemented by ECMWF. These user activities would run 24/7, with a turnaround varying according to the load on the system and dependent on the load of the Data Handling System. Delays in these activities would be caught up with during quieter times, at night or during weekends.

Over a hundred users were regularly managing their own work through SMS/ecFlow on ecgate to run HARMONIE, ARPEGE, or even IFS experiments on the HPCF. In addition, many users gained access to ECMWF's graphical applications, Magics and Metview, on ecgate. Many third-party software packages were also available on ecgate and Member State users accessed them daily.

The ecgate service in hindsight

The first ecgate Unix system was set up in 1993 as a front end to the HPCF, offering only basic job and file management facilities. Over the years, ecgate was enhanced to support most user requirements, except the possibility to run computationally intensive (parallel) applications, which would run on the HPCF.

a

“Access to the ecgate service was essential for the Hungarian Meteorological Service. Hungary has been an ECMWF Co-operating State since 1994, installing several ECMWF software packages on local servers. The right to use ecgate enabled the development of critical operational applications, such as custom clustering products for Central Europe. It also supported work that informed several master theses, such as on the dynamical downscaling of ECMWF ENS products with the ALADIN mesoscale limited-area model. These activities fostered the increasing cooperation between Hungary and ECMWF.”

Istvan Ihasz (Computing Representative and Meteorological Contact Point for the Hungarian Meteorological Service, OMSZ)

Having a separate and independent ecgate system gave ECMWF more flexibility in installing software packages needed by our users. For example, graphical packages and other Graphical User Interface applications would only be available on ecgate. The upgrade of ecgate in 2013 to Linux also gave a short-lived advantage compared to the HPCF, as it meant easier installation of third-party software packages and enabled an easier transition in 2014 to the new ECMWF HPCF, which would also run the Linux operating system.

Over the 30 years of its existence, ecgate was

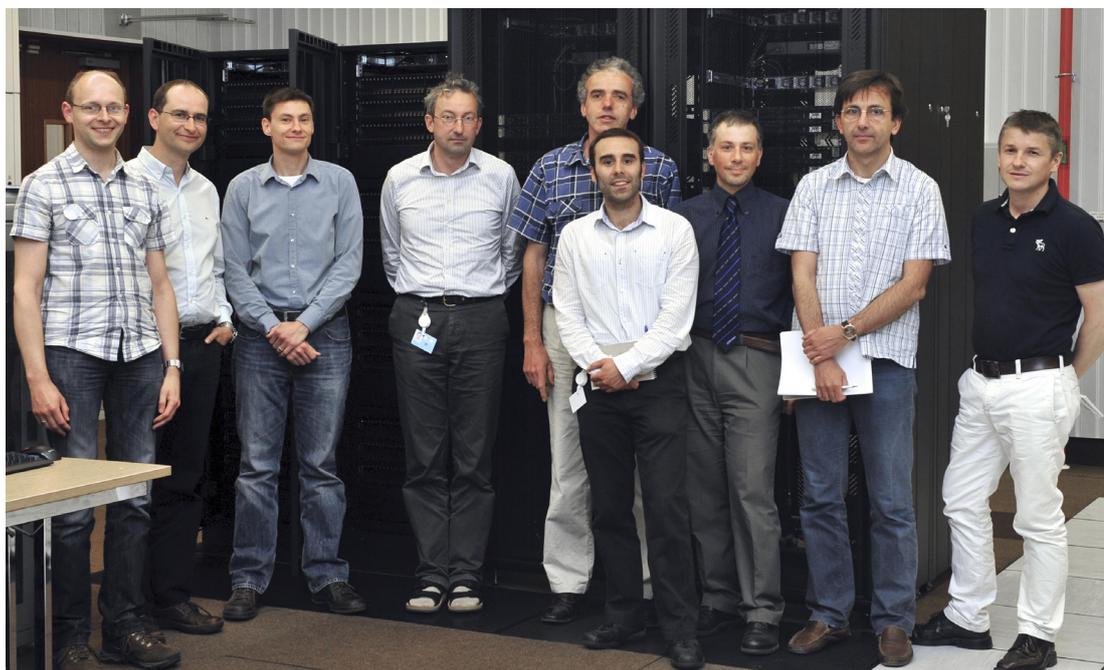


FIGURE 4 The team who provided the ecgate upgrade in 2013.

available to Member State users close to 100% of the time. In recent years, do-it-yourself software installation packages like conda have helped users to maintain and use their Python software stacks. On a few rare occasions, we have not been able to cater for some specialised user requirements on ecgate or, more generally, at ECMWF. This was mainly due to network and security constraints.

Figure 4 shows the ECMWF team working on the upgrade of the ecgate system in 2013, and Box A provides a view from one of our Member and Co-operating States.

Over the years, ECMWF conducted several user satisfaction surveys. We always received very good feedback on the ease of use and reliability of the service provided by ecgate and on its importance for Member and Co-operating State users to make effective use of ECMWF data either directly or via post-processing or limited-area downscaling.

The future

A general-purpose managed computing service is now available on part of the Atos HPCF, called ECS.

It offers similar functionality to that provided previously by the ecgate system and, by using the same HPCF infrastructure, it provides more efficient integration with the full HPC service, simplifying its use, maintenance, and support. We will continue offering such an important service to Member and Co-operating State users for the foreseeable future. The full Atos HPC managed service (<https://www.ecmwf.int/en/computing/our-facilities/supercomputer-facility>) complements it to run NWP models, while the European Weather Cloud (<https://www.europeanweather.cloud>) supports customised applications using ECMWF or EUMETSAT data on custom and personalised virtual platforms. More details on the ECGATE services available in Bologna can be viewed here: <https://confluence.ecmwf.int/display/UDOC/Atos+HPCF+and+ECGATE+services>. These three services complement each other and offer our users a wide range of computing platforms and advanced facilities to process ECMWF data efficiently.

Running a Global Broker as part of the new WMO data sharing solution

Rémy Giraud, David Podeur, Thierry Lacoste (all Météo-France)

The World Meteorological Organization (WMO) is upgrading its Global Telecommunication System (GTS) and the WMO Information System (WIS), which are used to transmit WMO data globally, to an Internet-based service called WIS 2.0 (WIS2). As described in a previous article in this Newsletter (<https://www.ecmwf.int/en/newsletter/176/computing/wis-20-wmo-data-sharing-21st-century>), in order to provide a reliable, efficient service for all WIS Users, the following Global Services have been defined:

- **Global Broker:** Meteorological centres will be responsible to make sure that all messages announcing the availability of new data and metadata can be easily obtained by all users. The Global Broker will provide a subscription service using the MQTT (Message Queuing Telemetry Transport) standard and Free and Open Source Software solution with an additional companion software (specific to WIS 2.0) to ensure uniqueness of messages as well as verifying the correct format of those messages.
- **Global Cache:** In order to provide quick and reliable access to core data as defined by the WMO Unified Data Policy, a copy of this data will be made available by a Global Cache. Storing data from originating WIS2 Nodes, the Global Cache will then make available the core data to all WIS Users.
- **Global Discovery Catalogue:** Each dataset available on WIS 2.0 must be described by a metadata record, using the OGC API - Records standard (soon to be ratified). The Global Discovery Catalogue will provide a discovery and metadata service using Free and Open Source Software, as well as provide quality assessment capabilities in support of continuous improvement of WIS 2.0 metadata.
- **Global Monitoring:** WIS 2.0 being an operational solution, it must be monitored. Each WIS2 Node and Global Service will provide metrics relevant to their operations. Global Monitoring Centres will collect the metrics and make available a visual dashboard presenting those metrics and alert the Centres when an unexpected event occurs in support of corrective action.

This article provides more detailed information on the architecture of WIS2 with a focus on one of the Global Services: the Global Broker and the instance of a Global Broker operated by Météo-France using the European Weather Cloud.

Publish and subscribe (pub/sub) in WIS2

In the weather, climate, hydrology and ocean community, where new data are almost constantly produced, it is key to make this new data available to customers as quickly as possible. Historically, the solution has been based on pushing the data from producers to consumers. This method has many merits, including in particular its simplicity. However, it also has a lot of drawbacks. For example, consumers cannot change easily what they receive, and the successful product distribution by producers depends on the availability of the IT systems of consumers.

When designing WIS2, its architects decided to replace the 'push by default' by 'pull by choice'. The new challenge was therefore to find a way to inform the user that new data are available.

Many messaging systems, such as X/Twitter and WhatsApp, are using a publish/subscribe approach. Consumers subscribe to channels, topics or feeds of interest, and producers publish their information. Consumers are immediately informed about any new information and can access it. WIS2 is following the same design principle. A data producer publishes a message in a topic, and users, by subscribing to the topic of interest, will be informed that new data are available. The message contains an HTTPS link to the resource. Global Services, in particular Global Caches and Global Brokers, are providing the required scalability and redundancy needed for WIS2 operations.

The WIS2 Notification Message

Each WIS2 Node will announce the availability of data using a WIS2 Notification Message. This message, which in the case shown below is from the Swedish Meteorological and Hydrological Institute (SMHI), is published on a local broker managed by SMHI:

```

{
  "id": "019918dc-0419-4bb1-aa70-5de539621087",
  "version": "v04",
  "geometry": {
    "type": "Point",
    "coordinates": [
      23.389945,
      67.20363
    ]
  },
  "properties": {
    "datetime": "2023-07-31T17:14:00Z",
    "hierarchy": "origin/a/wis2/swe/smhi/data/core/weather/surface-based-observations/synop",
    "data_id": "wis2/swe/smhi/data/core/weather/surface-based-observations/synop/WIGOS_0-20000-0-02095_2023-07-31T17:14:00",
    "integrity": {
      "method": "sha512",
      "value": "9cb67d6a6e5735f0166d93e07a19309cfe0824ca7f0b380d23eb74ebb6b59cbd26bec604e3fe1debee7c4b34f839c19dfe09e24b021f93d69ddb3ad8eacf08a2"
    },
    "pubtime": "2023-07-31T17:28:27.204915Z"
  },
  "type": "Feature",
  "links": [
    {
      "href": "https://wis-tst.smhi.se/WIGOS_0-20000-0-02095_2023-07-31T17:14:00.bufr",
      "rel": "canonical",
      "type": "application/x-bufr"
    }
  ]
}

```

The Global Brokers are subscribing to this broker and, after checking its uniqueness, this message is published on the broker part of the Global Broker. WIS Users (including Global Caches) will be informed of the availability of the dataset if they subscribe to the Global Broker. They can download the data if they are of interest to them. Figure 1 summarises the flow of messages in WIS2.

What constitutes a Global Broker?

A Global Broker is made up of two components:

- A highly available, scalable publish/subscribe broker
- Software ensuring deduplication of published messages by the WIS2 Node, the Global Cache and the other Global Brokers.

The broker part

Having decided to rely on a publish/subscribe

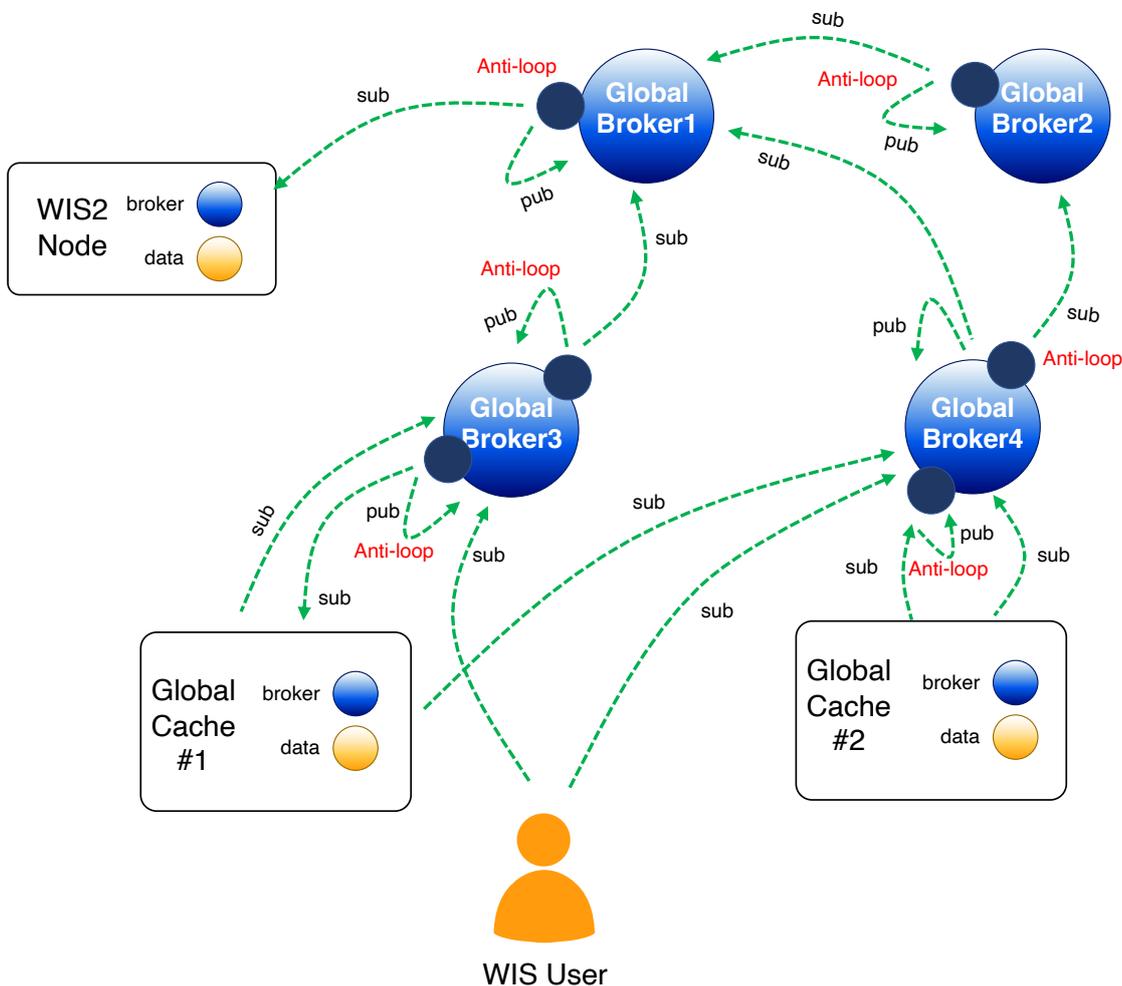


FIGURE 1 The flow of messages in WIS2 to alert a WIS User to the availability of data. The 'anti-loop' functionality is described below.

approach, the next decision while designing WIS2 was to choose the underlying protocols. Considering that WIS2 has decided to leverage open standards whenever possible, the potential protocols considered were:

- AMQP (Advanced Message Queuing Protocol) version 1.0
- Message Queuing Telemetry Transport (MQTT) version 3.1.1 and version 5.0.

After thorough analysis and based on hands-on experience by various WMO Members, WMO experts decided to rely on MQTT 3.1.1 and 5.0 for the publish/subscribe protocol in WIS2. Originally, MQTT 3.1.1 was designed to support the Internet of Things (IoT). It is widely used by many industries to collect data from sensors located in cars or machinery. Nowadays, the two versions of the MQTT protocol coexist, and most of the tools available as Free and Open Source Software support both versions. For example Mosquitto, the broker used in the WIS2-in-a-box reference implementation of a WIS2 Node, is compliant with both versions. A notable exception to this is RabbitMQ, whose currently available releases do not support version 5.0. However, VMWare, the owner of RabbitMQ, has confirmed that an upcoming release of RabbitMQ will support MQTT 5.0. It should be released by the end of 2023.

Some of the other MQTT broker implementations are EMQX (free and enterprise version), HiveMQ (only commercial, licence-based) and VerneMQ (free for personal use and commercial licence otherwise). In order to benefit from support from a commercial company, Météo-France decided to rely on VerneMQ to run its Global Broker. Support has been tremendous and has helped us to have a reliable solution to run the broker. The VerneMQ broker has been deployed on a cluster of virtual machines to provide a redundant and scalable service.

The anti-loop function

In a typical deployment of MQTT, the broker, the publishing part and the subscribing part are managed by the same entity. In particular, it is important to note that there is very little protection to prevent a publisher from flooding the broker with many messages.

First, to avoid potential flooding by publishers, it was decided that no WIS2 data provider would be allowed to publish on a broker that is not managed by the data provider itself.

Second, considering the extent of WIS2, an architecture with a single redundant broker ensuring all of the exchange of messages was discarded, and a design with multiple brokers providing a distributed, redundant, and reliable operation was adopted. Considering that there is

no standard method to copy messages between brokers, so that all messages are available on all brokers, a specific WIS2 solution called the 'anti-loop' has been designed.

The anti-loop tool:

- subscribes to as many brokers as needed; the brokers might be part of a WIS2 Node, a Global Cache or another Global Broker
- publishes to its local broker after having checked that the message has not yet been published to the broker.

In the WIS2 Notification Message example above, the id of the message is used to implement this 'anti-loop' feature.

Météo-France implementation of the anti-loop function

As this part is specific to WIS2, there is no off-the-shelf software that implements this feature. In late 2022, the anti-loop function was developed to be able to deploy the Global Broker feature during the pilot phase of WIS2, starting in early 2023. A flow-based, low-code solution based on the open-source tool Node-RED was used. This has ensured a rapid take-off of the WIS2 pilot phase while providing the required features.

A flow is a succession of nodes providing high-level components, such as MQTT subscription, variable manipulation, MQTT publication, and OpenMetrics (Prometheus) monitoring. It is also possible to develop specialised functions in javascript for features that are not part of a pre-existing node. As part of the Node-RED flow, Redis (another open-source tool) is used to store and then detect the uniqueness of the id providing the core feature of the anti-loop function.

The flow shown in Figure 2 is the most basic implementation of the anti-loop function. It is very compact and extremely easy to understand while providing the advanced features required. After six months of use in the pilot phase, it has been extremely reliable. What started as a test implementation can now be considered as production-ready and mature enough to be used in production for the upcoming phases of WIS2. The code and a docker container providing the anti-loop feature is available on GitHub (<https://github.com/golfvert/WIS2-GlobalBroker-Redundancy>) and Docker Hub (<https://hub.docker.com/r/golfvert/wis2gb>).

The version used today is more advanced and provides additional features. In particular, what started as being a single point of failure is now working in an active/passive mode on a Docker set of virtual machines. The current version of the flow is shown in Figure 3.



FIGURE 2 Node-RED flows are read from left to right. The message is received by the *Subscriber* (MQTT Subscribe) node, connected to the remote WIS2 Node. The message is sent to the downstream nodes to make available statistics (*Prometheus* and orange *Subscriber* nodes) and also to the other stream where the anti-loop function is performed. The *Save msg* node formats the message to query the redis database (*redis get*). The *Check ID* is checking whether the message has already been seen or not. Out of this node, the upper connection is used when the message was not received before. *Prometheus* and orange *Publisher* are providing statistics. *Prepare pub* and *Publisher* (MQTT Publish) are used to publish the received message to the broker of the Global Broker. Then *Save* and *redis set* are used to store the message on the redis database so that further messages with the same id will be discarded.

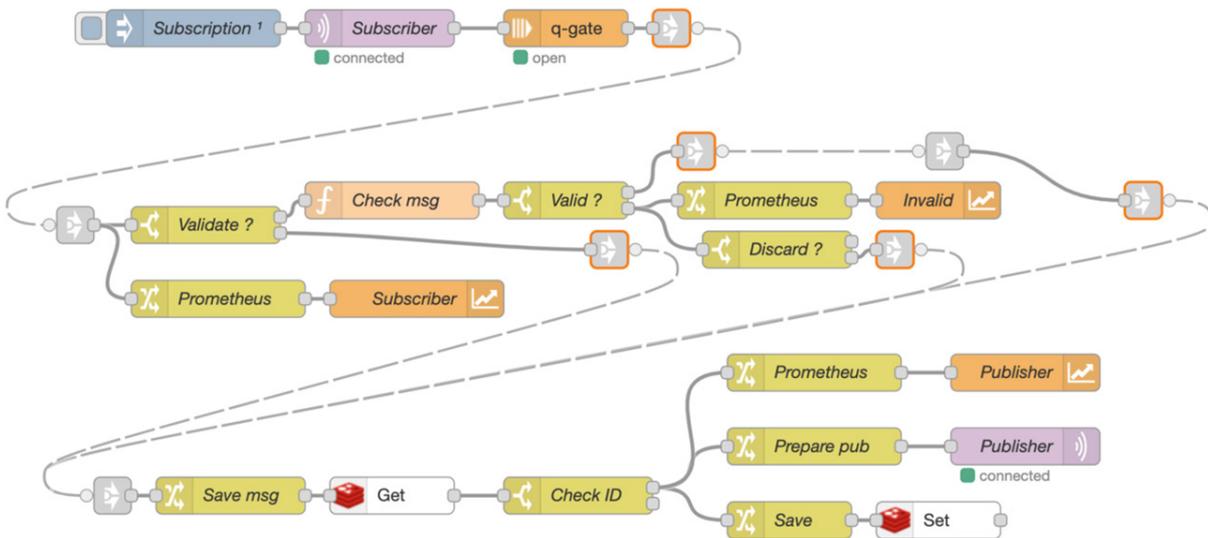
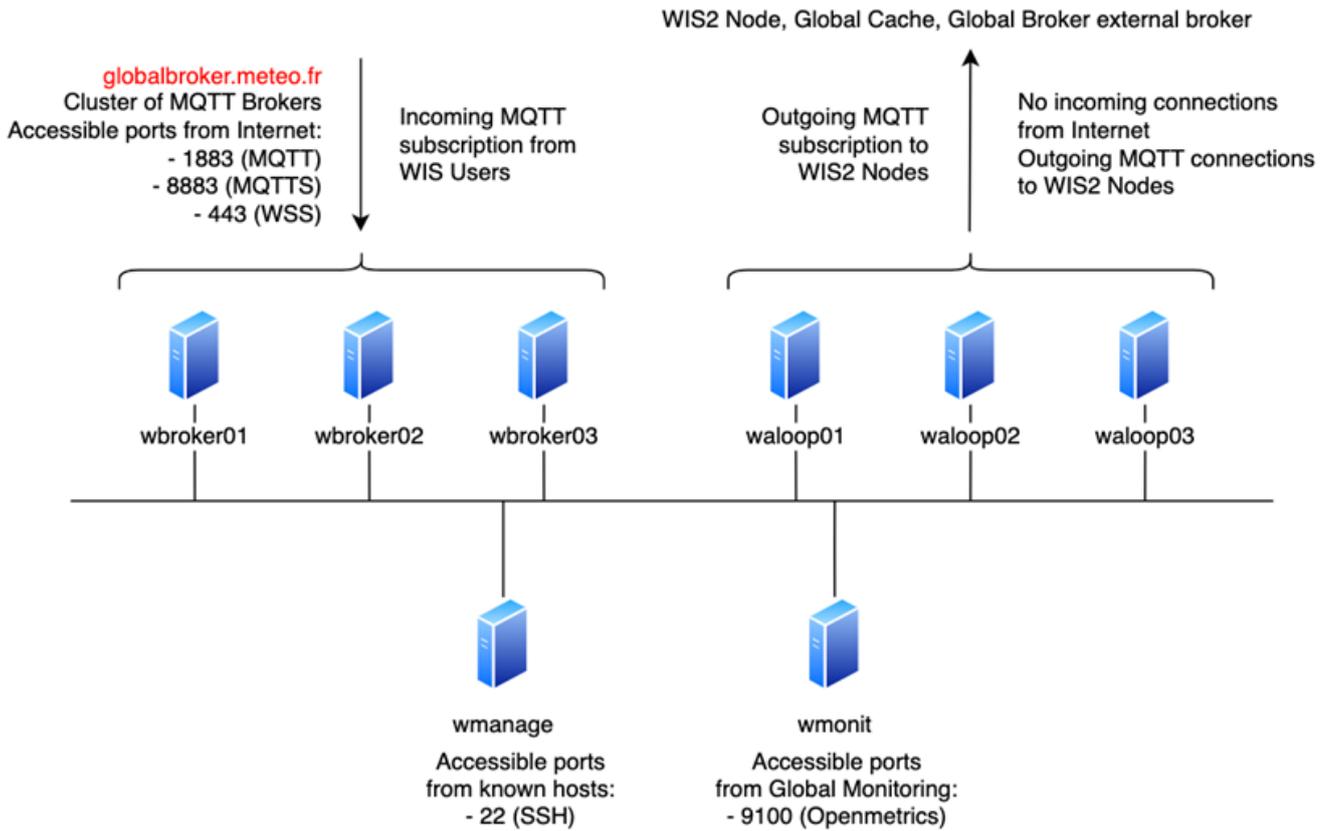


FIGURE 3 This is version 2 of the Node-RED-based anti-loop used by the Météo-France Global Broker. It has three improvements. First, the new version is using redis in a cluster. The *Get* and *Set* (with the red cube) are reading and writing on the redis database. Second, it is also possible to run multiple instances of the same container connected to the same WIS2 Node in an active/active manner. All instances are connected (*Subscriber* node) to the WIS2 Node. The *q-gate* node is either in 'open' mode, and the messages are sent to the downstream nodes, or in 'queue' mode, and the messages are kept in a holding queue. A separate flow, not shown here, is checking whether this instance is the primary one (and therefore process the messages) or a secondary one (and therefore queue the messages). If a secondary node is promoted to primary (the previous primary not being available any more), then the *q-gate* is 'flushed' (all queued messages are sent) and then put in 'open' mode. Thanks to this queuing mechanism, no message will be lost even in case of failure of the primary instance. The third improvement of the flow is the capability to *Validate* the messages. If configured, the messages can be syntactically checked (*Check msg*) and then *Discard* if the message is not compliant with the standardized format. In this case, statistics will also be produced for *Invalid* messages.

Running Météo-France Global Broker on the European Weather Cloud

Taking advantage of the OpenStack-based European Weather Cloud provided by ECMWF, Météo-France has decided to run the Global Broker first in Reading (December 2022 – September 2023) and then in Bologna. ECMWF is providing the security layer (firewall) and the load-balancing layer (Octavia service in Open Stack). The overall architecture is shown in Figure 4.

The three 'wbroker' hosts are hosting the VerneMQ software and are clustered to provide the reliable, scalable and redundant MQTT 3.1.1 and MQTT 5.0 pub/sub protocol. The three 'waloop' servers are hosting the anti-loop docker containers. There is one primary container per WIS2 Node, running on one of the three hosts. Traefik, another open-source tool, is used extensively to load-balance the traffic between the hosts. The two additional servers ('wmanage' and 'wmonit') are used to manage and monitor the entire environment.



- Traefik is installed on all hosts (except wmanage)
 - Redis (in docker) is installed on wbroker0x and waloop0x and is configured in cluster
 - On waloop0x, the antiloop software is running in primary/secondary mode
 - wmonit is gathering all metrics and stores in Prometheus, metrics can be collected by Global Monitoring
- MQTT (Message-Queueing Telemetry Transport)
 - MQTTS (The Secure Socket Layer based version of MQTT)
 - SSH (Secure SHell)
 - WSS (WebSocket Secure)

FIGURE 4 The architecture of the Météo-France Global Broker.

Conclusion

Having decided to use open standards and, as a consequence, being able to use off-the-shelf software is one of the key choices made by the architects of WIS2. The Météo-France Global Broker is an excellent example of the benefits of those choices. Built around VerneMQ, Node-RED, Redis, and Traefik, developing a reliable and scalable Global Broker has been easy and quick.

Thanks to the support offered by ECMWF (European Weather Cloud team and networking team), Météo-France has provided to the WIS2 community the first example of a Global Broker. During the pilot phase, the level of service has been extremely high. Running at ECMWF on the European Weather Cloud in Bologna, a state-of-the-art facility, the Global Broker will provide Météo-France with the environment needed to run one of the Global Services critical for the success of WIS2 and the upcoming migration from the GTS and WIS.

ECMWF publications

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

Technical Memoranda

911 **Haiden, T., M. Janousek, F. Vitart, Z. Ben-Bouallegue & F. Prates:** Evaluation of ECMWF forecasts, including the 2023 upgrade. *September 2023*

ESA Contract Reports

Fielding, M. & M. Janiskova: Assimilation system development for additional EarthCARE obs: Doppler velocity, cloud extinction, Rayleigh backscatter. *June 2023*

Fielding, M. & M. Janiskova: Assimilation system preliminary assessment for additional EarthCARE obs: Doppler velocity, cloud extinction, Rayleigh backscatter. *June 2023*

Fielding, M. & M. Janiskova: Optimising impact of radar reflectivity and lidar backscatter obs on analysis. *June 2023*

EUMETSAT Contract Reports

Lean, K., N. Bormann & S. Healy: Task 1.1 Evaluation of initial future EPS-Sterna constellations with 50 and 183 GHz. *October 2023*

ECMWF Calendar 2023/24

2023			
Oct 24–25	Finance Committee	Feb 5–8	Training course: Use and Interpretation of ECMWF products
Oct 25	Policy Advisory Committee	Apr 8–12	Data policy meetings of ECMWF, EUMETSAT & EUMETNET
Oct 30–Nov 3	Online training course: Introduction to ECMWF computing services (including MARS)	Apr 10–12	5th Workshop on Waves and Wave-Coupled Processes
Oct 31	Advisory Committee of Co-operating States (virtual)	Apr 23	Policy Advisory Committee (hybrid)
Nov 7	MAELSTROM Dissemination Workshop	Apr 24	Finance Committee (hybrid)
Nov 8–10	MAELSTROM Boot Camp	May 7–10	ECMWF–ESA Machine Learning Workshop (Italy)
Nov 13–17	Training course: A hands-on introduction to Numerical Weather Prediction Models: Understanding and Experimenting	Jun 3–6	8th CAMS General Assembly
Nov 20–24	Training course: Parametrization of subgrid physical processes	Jun 10–13	7th C3S General Assembly
Nov 27–Dec 1	Training course: Predictability and ensemble forecast systems	Jun 11–13	ROM SAF workshop
Nov 29	CAMS User Interaction Workshop	Jun 19–20	Council
Dec 7–8	Council	Sep 9–12	Workshop on Diagnostics in Forecasting
		Oct 7–9	Scientific Advisory Committee
		Oct 10–11	Technical Advisory Committee
		Oct 21–22	Finance Committee
		Oct 22	Policy Advisory Committee
		Dec 10–11	Council
2024			
Jan 17–19	1st Copernicus Climate Change Service Evolution (CERISE) General Assembly		

Contact information

ECMWF, Shinfield Park, Reading, RG2 9AX, UK
 Telephone National 0118 949 9000
 Telephone International +44 118 949 9000

ECMWF's public website www.ecmwf.int/

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

For any query, issue or feedback, please contact ECMWF's Service Desk at servicedesk@ecmwf.int. Please specify whether your query is related to forecast products, computing and archiving services, the installation of a software package, access to ECMWF data, or any other issue. The more precise you are, the more quickly we will be able to deal with your query.



Newsletter | **No. 177** | Autumn 2023

European Centre for Medium-Range Weather Forecasts

www.ecmwf.int