ECMWF

European Centre for Medium-Range Weather Forecasts

Newsletter

No. 158 | Winter 2018/19

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 The ECMWF Newsletter is not peer-reviewed.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 is available at *www.ecmwf.int/en/about/media-centre/media-resources*

Bologna and beyond

The year that lies ahead promises to be exciting on several fronts. Work to construct our new data centre in the Italian city of Bologna has started and is expected to be complete in the autumn. In parallel, an Invitation to Tender has been issued for ECMWF's next high-performance computing facility, to be installed at the Tecnopolo di Bologna by the first quarter of 2020. The new data centre will help us to achieve our strategic goal to produce forecasts with increasing fidelity on time ranges up to one year ahead.

Specific goals for 2025 include skilful ensemble predictions of high-impact weather up to two weeks ahead and predictions of large-scale patterns and regime transitions up to four weeks ahead. This Newsletter gives some clues as to where we stand now. ECMWF forecasts provided early indications of multiple weather hazards in Italy last autumn, but work on whether this kind of event can be predicted beyond ten days in advance is continuing. Research on the predictability of cold spells in Europe has resulted in a new product with useful skill up to two and a half weeks ahead. These examples show that progress is being made but that a lot of challenging work still lies ahead.

One area in which substantial progress will be made in 2019 is data assimilation. By feeding observations into the data assimilation system more continuously, significantly more observations can be used. This results in a better analysis and better forecasts. Separately, a 50-member Ensemble of Data Assimilations has been developed, an important step towards an integrated ensemble system. Both these developments will be implemented in the next upgrade of ECMWF's Integrated Forecasting System later this year. Work of a more long-term nature is under way on an alternative dynamical core option. First results show that the Finite-Volume Module is computationally efficient and performs well in benchmark tests.

This year will also see the retirement of a very venerable product: the ERA-Interim atmospheric reanalysis, to which no further data will be added after mid-2019. The tens of thousands of users of ERA-Interim can start using the new ERA5 atmospheric reanalysis, which now goes back to 1979 and will eventually go back all the way to 1950. ERA5 is an example of European cooperation at its best: successive EU-funded research projects enabled the development of ERA5, which is now produced operationally by the Copernicus Climate Change Service (C3S), implemented by ECMWF on behalf of the EU.

International collaboration has also made it possible to develop a new global reanalysis of atmospheric composition, CAMSRA: produced by the Copernicus Atmosphere Monitoring Service (CAMS), implemented by ECMWF on behalf of the EU, a range of European institutes helped to validate it.

Whether in forecast products, weather science or environmental services, there is thus much to look forward to in 2019 and beyond as we continue to work towards our strategic goals.

Florence Rabier
Director-General

Contents

Editor Georg Lentze • **Typesetting & Graphics** Anabel Bowen • **Cover image** alev gunebakanli/istock/Thinkstock

Predicting multiple weather hazards over Italy

Linus Magnusson, Luigi Cavaleri (CNR ISMAR, Italy)

Autumn 2018 was characterised by a continuation of dry weather in northern Europe while southern Europe experienced several episodes of very wet and windy weather. In this article, we will focus on a cyclone which developed from a large-scale trough over the western Mediterranean. On 29 October, it moved from Sardinia towards the Alps. In Sardinia, severe convection linked to the cyclone brought extreme accumulations of hail. In the Eastern Alps, recorded wind gusts reached 59 m/s. The storm uprooted or broke 11 million trees, and torrential rainfall led to flash floods. Later the Po river experienced moderate flooding. The storm was accompanied by significant wave heights up to 11 m in the Tyrrhenian and Ligurian Seas west of Italy and more than 9 m in the Adriatic. About six metres were recorded at the ISMAR oceanographic tower in the Gulf of Venice. There was a severe surge bringing flooding to Venice, thankfully mitigated by the fact that the meteorological peak occurred during low tide. This is an example of a single weather system bringing multiple hazards, and here we will discuss the predictability of some of them.

For the verification of 24-hour precipitation, we have access to high-resolution observations from many of our Member and Co-operating States thanks to a project to collect high-density observations. For the period 29 October 06 UTC to 30 October 06 UTC, three stations in Italy reported more than 300 mm, and the average precipitation from all observations in a 2°x2° box centred over north-eastern Italy was 100 mm. The maximum precipitation in ECMWF's high-resolution forecast (HRES) starting on 25 October 00 UTC reached 173 mm. Underestimation of orographic precipitation is a known model deficiency even though it has improved in recent years as a result of upgrades in model physics and model resolution. The area average in the 2°x2° box for this HRES forecast was 82 mm based on all observation points and 63 mm based on all grid points, on a similar level to short-range forecasts. The HRES and ensemble median consistently predicted values above the 99th percentile of the model climate from six days before the event and onwards.

The signal of extensive rainfall, violent wind gusts and high waves in northern Italy on 29 October began to appear in the ensemble forecast a week before the event. This resulted in a relatively strong Extreme Forecast Index (EFI) signal, as exemplified by forecasts for the three variables from 23 October 00 UTC. The early predictability was linked to the deep trough over the western Mediterranean.

The ensemble wave forecast gradually became more and more extreme as the event approached. Seven days before the event, the median of all ensemble forecasts was above the 75th percentile of the model climate and three days before above the 99th percentile in the Northern Adriatic. The last ensemble forecast before the event had a median significant wave height of 3 m while HRES gave 3.7 m. This is still significantly below the observed value of nearly 6 m. The discrepancy is not unexpected because of the global model's limited ability to resolve the Gulf of Venice and its bathymetry, and here limited-area wave models play a

Flooding in Venice. Flooding was one of the multiple weather-related hazards over parts of Italy in the last week of October 2018. (*Photo: Luciana Bertotti*)

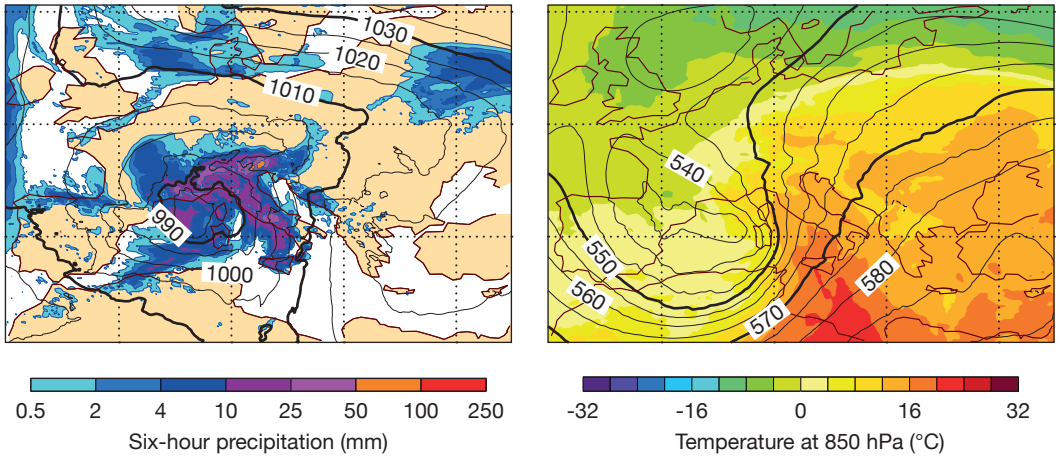

Meteorological situation on 29 October 2018. Analysis of mean sea-level pressure (contours, in hPa) and 6-hour precipitation (shading) (left) and geopotential height at 500 hPa (contours, in decametres) and temperature at 850 hPa (shading) (right) for 29 October 12 UTC.

role. The regional wave forecasts from the Institute of Marine Sciences (CNR ISMAR), Venice, use ECMWF's wind forecasts but adjust them (+16%) for a known underestimation in the Adriatic Sea. This simulation at higher resolution (about 0.08° for the wave model and down to about 100 m for the surge model) led to a very good forecast of wave and surge conditions, both accurately quantified several days in advance. For surge and flood forecasts, a key difficulty is to pinpoint with sufficient accuracy the time of the storm. Small phase differences between astronomical and meteorological components can lead to drastic differences in flooding. In the present case, the timing was sufficiently correct 24 hours in advance.

Overall, ECMWF forecasts provided an early indication of hazardous weather in this case, although they did not always capture just how extreme conditions were going to be. ECMWF is going to work on understanding if this type of event has predictability in the extended range (beyond 10 days) and also to explore how forecasts of extreme surface wind and precipitation can be improved.

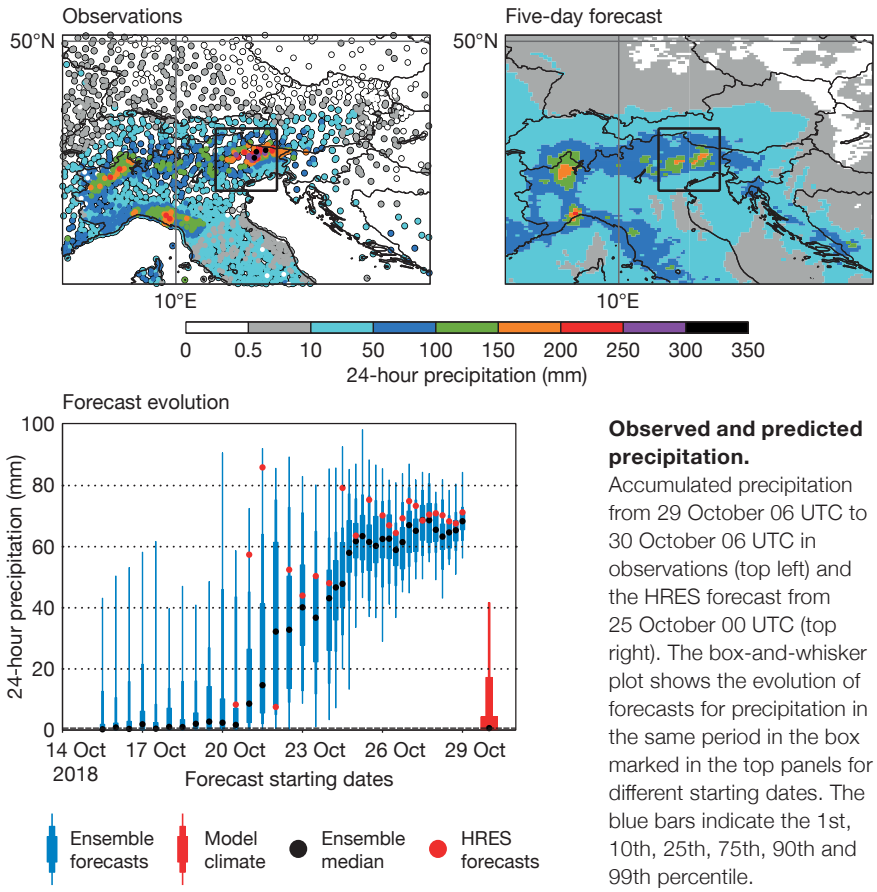

Observed and predicted precipitation. Accumulated precipitation from 29 October 06 UTC to 30 October 06 UTC in observations (top left) and the HRES forecast from 25 October 00 UTC (top right). The box-and-whisker plot shows the evolution of forecasts for precipitation in the same period in the box marked in the top panels for different starting dates. The blue bars indicate the 1st, 10th, 25th, 75th, 90th and 99th percentile.

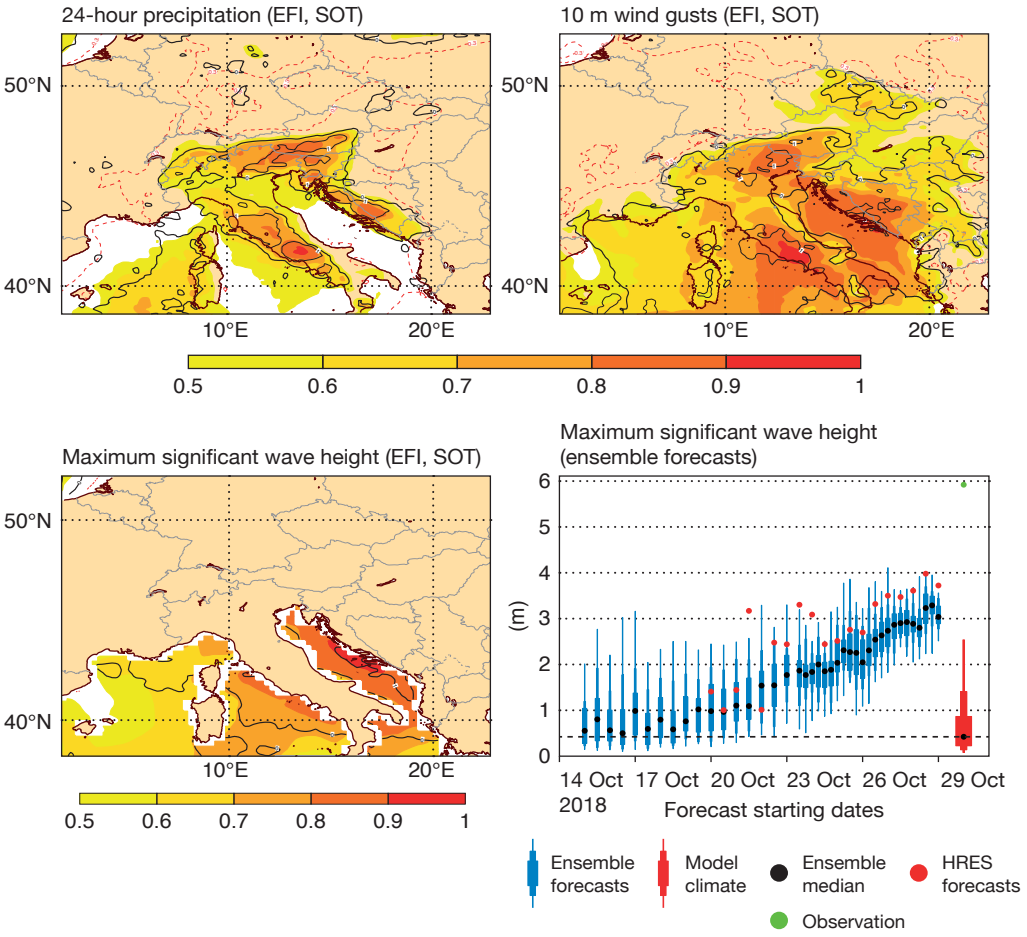

Extreme Forecast Index/ Shift of Tails and wave forecast. EFI and Shift of Tails (SOT) in the forecast from 23 October 00 UTC valid on 29 October for 24-hour precipitation (top left), 24-hour maximum wind gusts (top right) and maximum significant wave height (bottom left). The bottom right panel shows the evolution of maximum significant wave height forecasts for a point off the coast of Venice.

Assimilation of radiance data from GOES-16

Chris Burrows, Tony McNally

Since July 2018, ECMWF has been assimilating radiance data from the American satellite GOES-16. In November 2016, the satellite was launched into a geostationary orbit whose position allows continuous coverage of the Americas, including the western Atlantic and eastern Pacific. The official GOES-16 clear-sky radiance (CSR) product is not due to be available until May 2019, but close collaboration between ECMWF and the US National Oceanic and Atmospheric Administration (NOAA) has made it possible to start assimilating the data well before then.

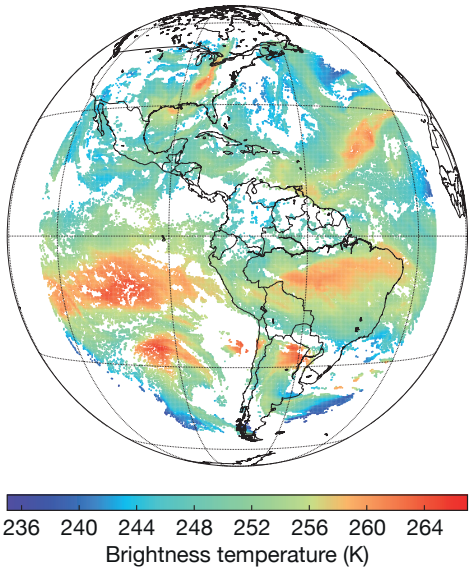

GOES-16 water vapour channel data. Clear-sky radiances as measured by the 7.0 μm water vapour channel of GOES-16 at 00:00 UTC on 26 July 2018. Observations considered cloud-affected have been identified and removed by NOAA.

GOES-16 capabilities

The satellite is equipped with an instrument which provides images of the Earth for a number of spectral bands (channels) covering wavelengths from the visible to the thermal infrared. For data assimilation, we are primarily interested in the channels which cover the wavelength range 5.5–7.5 μm as they are sensitive to atmospheric water vapour. Compared to GOES-13, its predecessor as the official 'GOES-EAST' satellite, GOES-16 has many additional capabilities thanks to the new ABI instrument (Advanced Baseline Imager). Notably, there is a significant improvement in spectral, spatial and temporal resolution, each of which has benefits for data assimilation.

As well as being sensitive to humidity, the water vapour channels also pick up strong signals from clouds, and any cloud-affected observations must be removed before the observations can be assimilated. For GOES-16, this is performed by NOAA. The resulting CSRs are box-averages of the clear-sky pixels.

Collaboration with NOAA

In January 2018, GOES-13 was retired and GOES-16 became the official GOES-EAST satellite. Since the official CSR product is not due to be available until May 2019, contact was made with NOAA to try to determine whether unofficial data could be made available to ECMWF before the 2018 hurricane season. NOAA/NESDIS/STAR were developing the CSR product, and ECMWF provided some assistance to ensure these data were in the correct format. Subsequently, the CSR data were made available to ECMWF and other centres via FTP, and the quality of the data was assessed and the analysis fed back to NOAA.

An example of the interaction between ECMWF and NOAA was related to horizontal stripes appearing in the images, which also affects the CSRs. Guidance from NOAA enabled us to implement a quality check whereby, if the standard deviation of the clear-sky pixels within an averaging box exceeds 4 K (in terms of brightness temperature), then the CSR is rejected from the assimilation to prevent degradation of the atmospheric analysis.

Assimilation

Several months after GOES-13 was retired, GOES-16 CSRs started to be assimilated operationally on 26 July 2018, following a period of passive monitoring during which bias correction coefficients were spun up and the quality was shown to be comparable to that of data from the Japanese satellite Himawari-8. As well as CSRs, wind information from GOES-16 is also assimilated directly. These winds are obtained by tracking cloud-features in consecutive visible and infrared images, hence this information is complementary to that provided by the CSR data. These retrievals are called atmospheric motion vectors (AMVs). In the period leading up to implementation, an assimilation experiment was performed to assess the impact on forecasts due to the assimilation of GOES-16 radiances. This demonstrated that the short-range forecast of humidity and wind was improved, as verified against independent satellite and in-situ observations.

Future plans

GOES-16 provides a full-disc scan every 15 minutes, unlike GOES-13, which produced one every 3 hours. Currently it is not possible to fully exploit the high temporal sampling in the ECMWF system, but we intend to investigate any benefits we may gain from it, particularly the potential improvements to the wind field which may be obtained by the 4D-Var system tracing water vapour features throughout the assimilation window. GOES-17, one of the three successors to GOES-16, has already been launched, and we plan to assess data from this satellite when it attains operational status, although some ongoing issues with the infrared channels may limit the volume of data we can use from this satellite.

ECMWF's ERA5 reanalysis extends back to 1979

Hans Hersbach

In January 2019, ECMWF released a further 21 years of its new global atmospheric reanalysis, ERA5, which now extends back to 1979. The data is available from the Copernicus Climate Change Service (C3S) implemented by ECMWF on behalf of the EU (https://climate.copernicus.eu/). This innovative dataset will eventually go back all the way to 1950.

Atmospheric reanalysis combines a weather model with observations to create a complete and consistent picture of meteorological conditions typically over several decades in the recent past. Reanalysis datasets are used for climate monitoring, climate studies and meteorological research, and they underpin an increasing variety of sectoral applications. They also support operational numerical weather prediction, for example by providing a reference against which the quality of forecasts from successive model versions can be compared.

ERA5 is by far the most popular dataset in the openly accessible C3S Climate Data Store. It contains estimates of atmospheric variables such as air temperature, pressure and wind at different altitudes, as well as surface variables such as rainfall, soil moisture content and ocean wave height. ERA5 replaces ECMWF's previous atmospheric reanalysis, ERA-Interim, which also covers the period from 1979 to the present. The main differences to ERA-Interim are that ERA5:

- benefits from 10 years of improvements in numerical weather prediction (it uses Cycle 41r2 of ECMWF's Integrated Forecasting System, implemented in March 2016)
- uses improved historical observations, boundary conditions and external forcing
- provides hourly information at a grid spacing of around 31 km globally and 137 levels in the vertical, up from six-hourly fields at a grid spacing of 79 km and 60 levels in ERA-Interim
- includes uncertainty information by providing an ensemble of 10 reanalyses.

The production of data back to 1950 will be complete near the end of this year. Continuous improvement of climate reanalysis (centennial, coupled, regional) features prominently in plans for the development of C3S.

ERA5

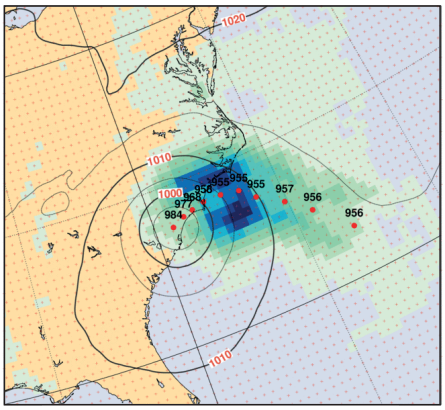

ERA-Interim

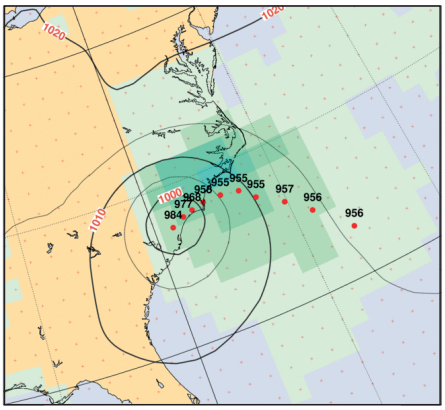

Hurricane Florence in ERA5 and ERA-Interim. The ERA5 reanalysis provides a detailed picture of the track, intensity (contours, in hPa) and accumulated rainfall (shading, in mm) associated with Hurricane Florence, which struck the east coast of the US in September 2018. The plots show the situation at 09 UTC on 15 September 2018 as reconstructed by ERA5 (left) and ERA-Interim (right).

New observations since July 2018

The following new observations have been activated in the operational ECMWF assimilation system since July 2018.

Observations	Main impact	Activation date
Clear-Sky Radiances from GOES-16 (replacing GOES-13, which has been excluded since 14 December 2017)	Tropospheric humidity and wind	26 July 2018
Radiances from CrIS on NOAA-20	Temperature, humidity, ozone, dynamics	11 September 2018

Atmospheric composition priorities for numerical weather prediction

Rossana Dragani

One of ECMWF's strategic goals for 2025 is to develop an integrated global model of the Earth system to produce forecasts with increasing fidelity on time ranges up to a year ahead. This will be achieved by incorporating an increased level of complexity of physical and chemical processes and of the interactions between the atmosphere, the ocean, sea ice and the land surface into the model. Atmospheric composition has the potential to be one of the sources of predictability at different timescales. Development activities at ECMWF have built a capacity to simulate and assimilate a variety of atmospheric composition species in the Integrated Forecasting System (IFS). Today this capacity is at the core of the Copernicus Atmosphere Monitoring Service (CAMS) implemented by ECMWF on behalf of the EU.

It is worth noting that the complexity of many atmospheric composition modules in the IFS means that they are often computationally unaffordable at the resolutions used in numerical weather prediction (NWP). Only stratospheric ozone is a prognostic variable in NWP applications, and it is included interactively in the radiative transfer model used for radiance assimilation. In contrast, the radiation scheme still relies on climatologies of aerosols, ozone and trace gases.

Members of the ECMWF ACWG

R. Dragani, A. Benedetti, J. Flemming, G. Balsamo, M. Diamantakis, A. Geer, R. Hogan, T. Stockdale, M. Ades, A. Agusti-Panareda, J. Barre, P. Bechtold, A. Bozzo, H. Hersbach, E. Hólm, A. Inness, S. Johnson, Z. Kipling, J. Letertre-Danczak, S. Massart, M. Matricardi, T. McNally, M. Parrington, I. Sandu, C. Soci, and F. Vitart.

Considerable collaborative effort has been made in the last few years to build these climatologies from CAMS products with the aim of improving the realism of the atmospheric composition fields in the radiative calculations and of the related weather feedbacks.

Working group recommendations

A cross-departmental working group on atmospheric composition (ACWG) was set up at ECMWF in spring 2018. Its main concern to date has been to identify a set of priority developments that have the potential to improve forecasts at all scales, from a few days to seasons ahead. The aim is to assess their potential in the time frame 2019–2022 and, where benefits exist, to proceed with operational implementation.

The ACWG suggests improving the representation of atmospheric composition in NWP to the level of complexity and coupling that is beneficial for NWP by leveraging the CAMS system and any other development related to atmospheric composition wherever possible. This can be achieved by building on the experience gained to date, hence the focus is on ozone, aerosols, and CO_2. The overarching goals for NWP applications in the 2019–2022 timeframe are:

1. Seamless fully interactive prognostic ozone; fully interactive refers to the coupling in both the radiation scheme and in data assimilation, targeting operational implementation by 2022.
2. Improved aerosol climatologies and, through systematic testing, seeking to demonstrate the benefits, at an affordable cost, of including a limited number of prognostic species (e.g. dust).
3. Enhanced representation of CO_2 in the radiative transfer model used for the radiance assimilation.

Each of these overarching goals requires substantial work in terms of development, and thorough assessment studies, and encompasses all areas from atmospheric composition modelling and coupling in physical processes to data assimilation (DA). The CAMS configuration will serve as an initial test bed for species not included in the current NWP configurations. Attention will also be paid to code structure and code efficiency, and, provided a higher level of coupling between atmospheric composition and meteorology is introduced, to how the model performance evaluation process might need to be adapted in the future. Finally, beyond improving forecast skill there are several products and diagnostics that could benefit from an improved atmospheric composition representation, such as visibility and high-spatial-resolution UV products. More details can be found in *ECMWF Technical Memorandum* No. 833, December 2018, by R. Dragani, et al.

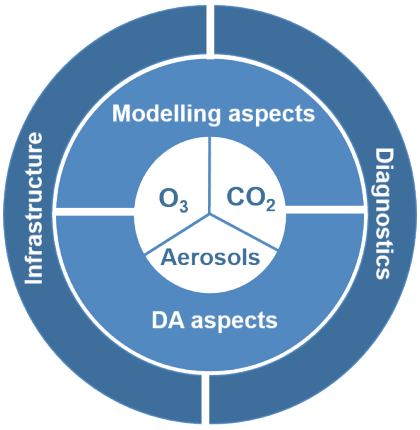

Summary diagram of the aspects considered by the ACWG. Recommendations are provided on modelling and data assimilation aspects for ozone, CO_2 and aerosols. The ACWG has also made recommendations on how to improve the atmospheric composition numerical infrastructure and code efficiency, and on how the model performance evaluation process will need to be modified in the future to account for atmospheric composition.

Water found on MARS

Fredrik Wetterhall, Blazej Krzeminski, Shahram Najm, Sébastien Villaume, Manuel Fuentes

In October 2018, the European Flood Awareness System (EFAS) started archiving its operational hydrological model output in the Meteorological Archival and Retrieval System (MARS) hosted by ECMWF. This marked the end of a long project to define new templates and data models for storing post-processed data in MARS in the GRIB2 format. The new capabilities in MARS have made the archiving of EFAS data more robust and more flexible for product generation purposes. They have also opened up MARS for storing data from any downstream application based on numerical weather prediction (NWP) output. ECMWF is the computational centre for EFAS, a flood forecasting system which is part of the EU-funded Copernicus Emergency Management Service (CEMS) (see diagram).

Requirements

Users often request output from EFAS. That is why a few years ago ECMWF decided to make EFAS output available through MARS. This would enable a more structured delivery of data and improve internal workflows for the validation of operational forecasts and experiments. However, it soon became evident that there were a number of obstacles to overcome. First, very few hydrological variables had GRIB code definitions. New parameters were proposed to the World Meteorological Organization (WMO) and were accepted after some revision. Second, there was no ready-made template for a hydro-meteorological forecast chain where the hydrological model is forced with NWP output. For our purposes, we had to define a new template where important metadata from the NWP fields used is stored in the GRIB header of the hydrological output.

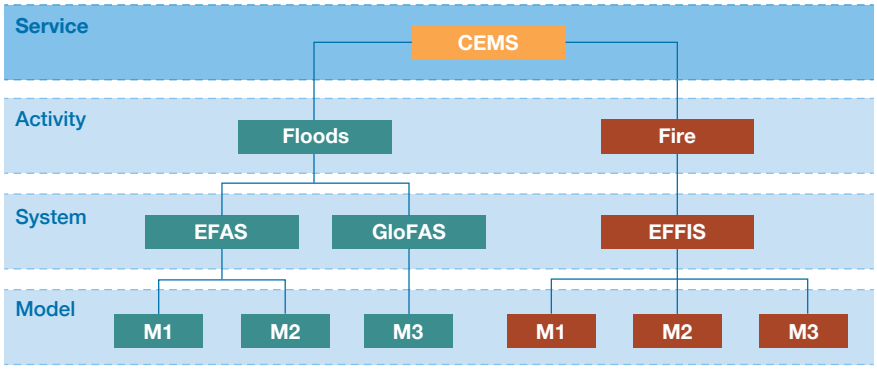

Overview of flood and fire activities in CEMS. ECMWF is the computational centre for EFAS, GloFAS and EFFIS. Those three operational systems are all part of CEMS and use a number of different models for different purposes.

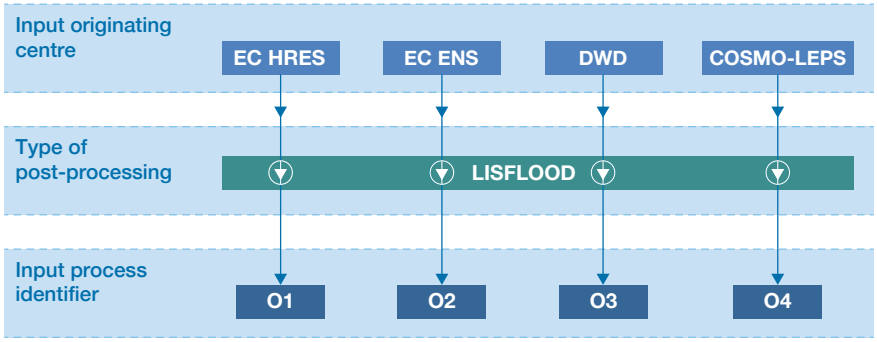

Operational chain in EFAS. Example of an operational chain where the hydrological model LISFLOOD is forced with input from different NWP models, which all need to be traceable in MARS.

Implementation

Using EFAS as an example, the implementation of the new GRIB template enables the storage of information on the forcing data through the key 'Input originating centre'. This should be the official code of the centre producing the NWP forcing data. For a multi-model system such as EFAS, this functionality is essential. The type of post-processing applied is stored in the key 'Type of post-processing'. This refers to a local table storing information on the models used in the post-processing. In our example, the model used is the hydrological model LISFLOOD. The third key, 'Input process identifier', is used to store information on the model version used, which can be a combined identifier based on the model versions used to produce the forcing and post-processed data, or just the version of the post-processing model. There are four versions of the template to account for deterministic and ensemble forecasts as well as instantaneous and statistically processed values.

Benefits and future work

The data model for post-processed data was deliberately made generic to account for different types of applications using NWP output, including calibrated forecasts, impact models and verification. The template has so far been implemented for EFAS operational forecasts and the EFAS climatology. Seasonal and sub-seasonal hydrological forecasts, including hindcasts, are expected to be added soon. Other downstream applications, such as the Global Flood Awareness System (GloFAS) and the fire forecasting output from the European Forest Fire Information System (EFFIS), will follow suit. EFAS data is currently archived under the new class 'Copernicus Emergency Management Service (CEMS)' in MARS. One of the immediate benefits of this implementation is that it can facilitate the delivery of non-restricted EFAS data through the Copernicus Climate Data Store for public access.

Major upgrade for global flood forecasts

Ervin Zsótér, Christel Prudhomme, Shaun Harrigan

On 14 November 2018, the Global Flood Awareness System (GloFAS) of the Copernicus Emergency Management Service (CEMS), for which ECMWF is the computational centre, was upgraded to version 2.0. The upgrade implemented a number of major system developments for improved global early warnings.

GloFAS aims to provide transboundary early flood guidance over all medium- to large-scale rivers of the world, helping national hydro-meteorological services, humanitarian agencies and commercial companies to improve their response to flood-related hazards. GloFAS has been developed in close cooperation between the European Commission's Joint Research Centre (JRC) and ECMWF, with support from national authorities and research institutions such as the University of Reading. GloFAS has a 30-day component (GloFAS 30-day), which has been producing probabilistic flood forecasts up to 30-days in a semi-operational fashion since 2011, and a seasonal hydrological outlook system (GloFAS Seasonal), which has been providing high- and low-flow anomalies up to 16 weeks ahead since November 2017. GloFAS became part of CEMS in April 2018, offering a fully operational, 24/7 service.

Main elements of the upgrade

The upgrade included a new calibration of the hydrological routing component to improve the river flood simulation worldwide. It also involved changes in the configuration to produce an improved hydrological reanalysis for more realistic climatological thresholds and improved real-time forecast initialisation. In response to requests from users, the upgrade also introduced a version numbering system for full traceability of GloFAS flood forecasts, significantly enhanced user documentation (now available on the web) and the publication of over 20 years of discharge forecast reruns.

Calibration: GloFAS version 2.0 implements a major calibration exercise for river routing parameters to better account for regional variability in hydrological processes. The calibration was conducted for 1,287 stations worldwide where daily discharge observations were available. As shown in the figure, the calibration has improved discharge simulations for the majority of stations. The GloFAS team is grateful to all data providers who have given us access to their data.

Reanalysis: The hydrological reanalysis provides atmosphere, land-surface and river state variables (e.g. runoff and river discharge) as close as possible to observed conditions. It is used in evaluation and to define the GloFAS discharge thresholds, which help to assess the severity of predicted flood events. In GloFAS version 2.0, the hydrological reanalysis uses high-resolution (31 km grid spacing) input data from ERA5, ECMWF's new atmospheric reanalysis.

Initialisation: To initialise the real-time ensemble discharge forecasts, the GloFAS state variables need to be updated every day. In GloFAS version 2.0, the simulations that produce the state variables also use high-resolution input data from ERA5. GloFAS version 2.0 ensures the highest possible consistency between the hydrological reanalysis and the real-time forecasts, addressing a known weakness of earlier GloFAS versions.

Versioning: A 2-digit version numbering system was introduced. The first number relates to major upgrades that affect the hydrological modelling chain, while the second number relates to other changes impacting the discharge. GloFAS

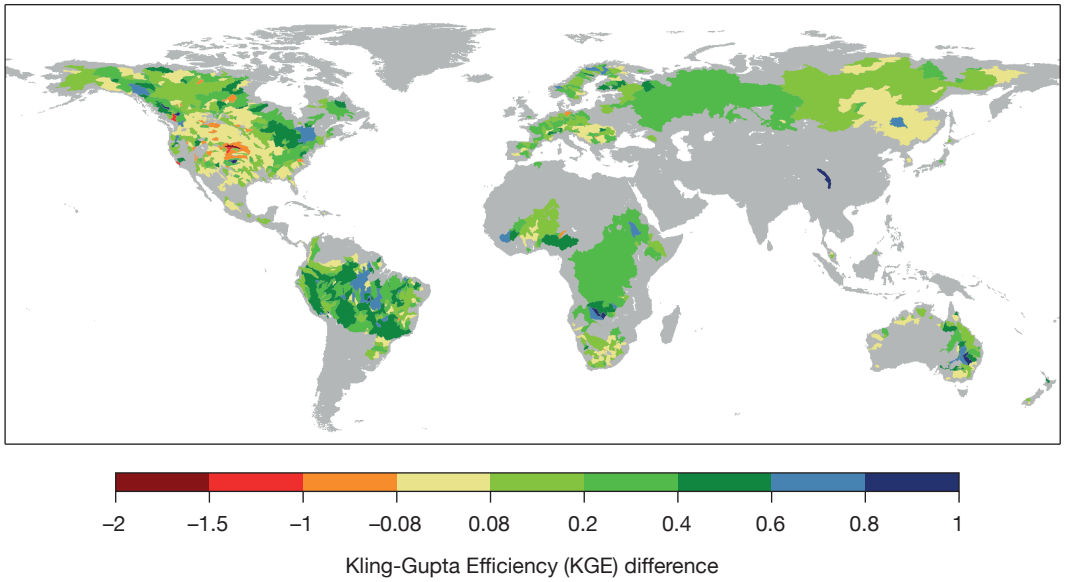

New calibration. GloFAS provides improved river discharge simulations after calibration over most of the catchments with available observations. Green/blue colours show higher skill, while orange/red colours indicate lower skill in simulations using the newly calibrated model. (Chart from an article by F.A. Hirpa et al. (2018) published in the Journal of Hydrology, doi:10.1016/j.jhydrol.2018.09.052, reproduced under the CC BY 4.0 licence.)

30-day and GloFAS Seasonal are considered as two independent systems, but after this upgrade they both share the version number 2.0.

Forecast reruns: The new discharge forecast reruns cover the period 1997 to 2018 for GloFAS 30-day and 1981 to 2017 for GloFAS Seasonal. They are simulations performed for a set of past dates with the same model version that is used for the real-time forecasts. These discharge forecast reruns provide a long, homogeneous forecast time series, ideal for analysing GloFAS performance.

Documentation: The GloFAS documentation has been expanded and updated with information on the modelling and different forecast configurations, as well as other user information. Details on the upgraded system are available in the General and Technical Information sections on the GloFAS website: www.globalfloods.eu.

ECMWF upgrades the Copernicus EFAS web portal

Maurizio Latini, Christel Prudhomme

ECMWF has redesigned the web portal of the European Flood Awareness System (EFAS) to create a more flexible, modern and personalised user experience. EFAS was developed by the European Commission's Joint Research Centre (JRC) and has been operational since 2012 as part of the Copernicus Emergency Management Service. The new portal (accessible via https://www.efas.eu/) became operational in January 2019.

ECMWF is the computational centre for EFAS. As such it is responsible for producing twice-daily flood forecasts and making them available to EFAS partners. The latter is done through EFAS-IS, an internet application that provides access to the forecasts through a map viewer. EFAS-IS also includes tools to issue flood notifications (formal, informal and flash flood) and to collect feedback on the quality of the forecasts and formal notifications.

The EFAS portal landing page. The landing page and most parts of the website are openly accessible, but full access to the map viewer and other tools depends on a user's credentials. The website has been designed following the Copernicus style guide recommendations.

Main features

The new portal uses predefined roles and associated permissions to control access to and management of the web portal, including the map viewer displaying the forecasts. User authentication is performed using a single-sign-on application. Once logged in, users can access various sections of the portal according to their credentials. A flexible, continuous integration development environment ensures a faster release cycle for new features. This was achieved using a container technology approach that decreases the time needed for development, testing and deployment of new versions of the EFAS-IS application.

The new web interface introduces modern tools, such as a dashboard, where users can create a personalised environment focusing on information they are most interested in, through widgets created from the map viewer. In addition, users with the right credentials can directly manage the portal by creating pages, editing text, uploading graphs etc. The portal also provides tools for EFAS forecasters based in dissemination centres to create flood notifications. This is managed through an intuitive system that behaves like a shopping cart which, once the notification process is complete, automatically sends warning emails to EFAS partners to inform them of possible flood events, and updates the notification layer so that maps of the ongoing flood warning can be visualised. The cart functionality also enables forecasters on duty to send a daily overview to the Emergency Response Coordination Centre (ERCC) of the European Commission containing information on ongoing and predicted floods in Europe.

Finally, the collection of in-situ qualitative evaluations of notifications and of missed events by EFAS partners is now facilitated and managed through the web interface. A simple webform, accessible directly from the viewer, collects user feedback in a way which ensures that all responses can be categorised and analysed, for consideration and integration in the EFAS development cycle.

The 18-month-long portal development work was supported by the JRC, other EFAS centres and EFAS partners, who helped in particular with the design and testing of the application.

How confident are predictability estimates of the winter North Atlantic Oscillation?

Antje Weisheimer, Damien Decremer

One of the main characteristics of seasonal forecasts in the extratropics is that forecast signals are relatively weak while noise levels are high. Seasonal predictions of winter atmospheric circulation over parts of the North Atlantic and Greenland during recent decades show that the correlation skill between observations and the ensemble mean is larger than one would expect from low signal-to-noise ratios. This is sometimes called the 'signal-to-noise paradox'. It implies that the real world appears more predictable than the forecast model suggests and that the forecast ensembles are underconfident (over-dispersive). A new study

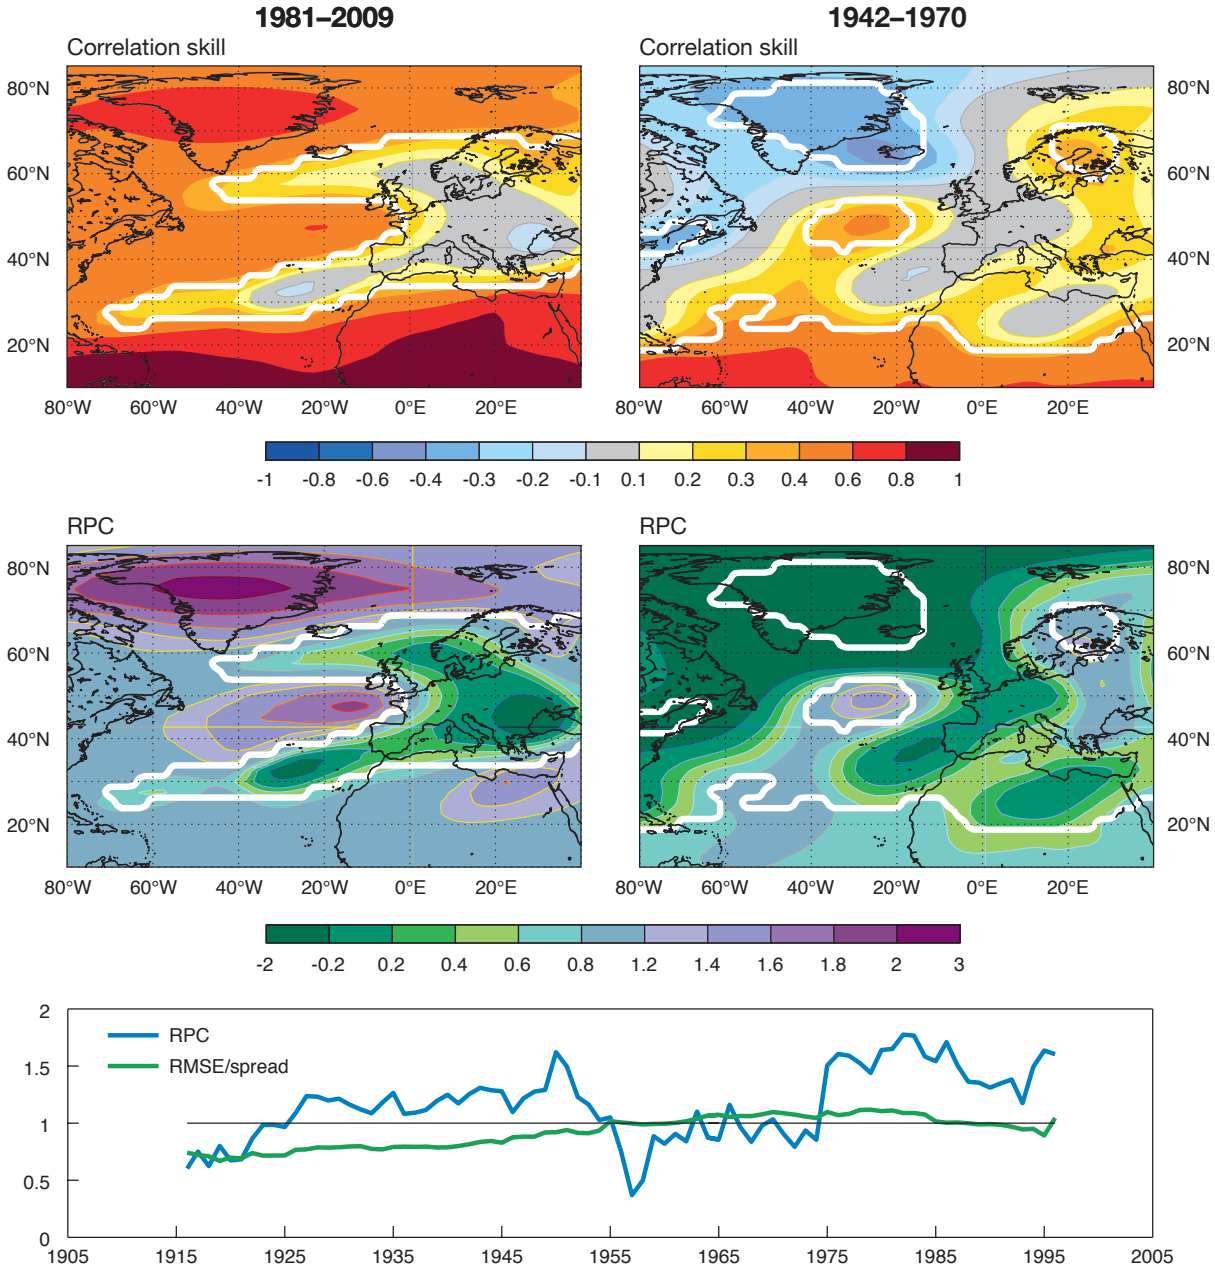

The signal-to-noise paradox over different timescales. The top four plots show anomaly correlation skill and Ratio of Predictable Components (RPC) for atmospheric seasonal hindcasts of geopotential height at 500 hPa during winter for two contrasting hindcast periods. The white contour lines show the 95% significance level of correlation skill. The time series below shows RPC and the ratio between RMSE and ensemble spread for the hindcast period from 1900 to 2010 computed over 30-year moving windows and plotted at the centre of each window.

analyses the behaviour of the ECMWF atmospheric model over a much longer hindcast period, from 1900 to 2009. It was found that the paradox disappears if the full 110 years are considered, but that there is also substantial decadal-scale variability. It follows that relatively short hindcasts are not representative of longer-term behaviour. In addition, correlation-based measures show larger sensitivities to small hindcast sample sizes than alternatives, such as error-based diagnostics.

Multi-decadal variability of predictability

The correlation skill of hindcasts of geopotential height at 500 hPa during winter (December to February) for the period 1981–2009 across Greenland and parts of the North Atlantic is statistically highly significant and larger than expected from purely model-based estimates. Here the hindcasts were generated with an atmosphere–land only configuration of the ECMWF model, but coupled hindcasts perform similarly.

A quantity called the Ratio of Predictable Components (RPC) can be used to measure the 'signal-to-noise paradox'. A perfect forecast system where the predictable components in the real world and in the model are similar would result in an RPC value of about one. If the model estimate of predictability is larger than the real-world predictable component, the RPC will be smaller than one. The forecasts are overconfident, or under-dispersive, as there is not enough variability in the ensemble. Seasonal forecast models generally tend to be overconfident, especially in the tropics.

In the situation of the 'paradox', the predictable components of the real world are larger than the model estimates of predictability, resulting in RPC values larger than one. As the figure shows, RPC in the period 1981–2009 is larger than one over Greenland and some eastern parts of the North Atlantic.

The availability of seasonal hindcasts over many decades makes it possible to trace the long-term evolution of RPC values. While the RPC of the North Atlantic Oscillation is very close to one if estimated over the full 110-year hindcast period 1900 to 2009, 30-year seasonal forecast skill and RPC vary on multi-decadal scales. For example, as shown in the figure, the mid-century period around the 1950s was characterised by a pronounced drop in skill. The corresponding RPC is either negative or close to 1, suggesting that the 'paradox' itself undergoes multi-decadal fluctuations.

The figure also shows the temporal evolution of the RPC for the NAO index during the 20th century. Within specific decades, RPC values can vary considerably. In contrast, the relation between the root-mean-square error (RMSE) and ensemble spread is much more stable, with no indication of underconfidence in any specific period. It is hypothesized that sampling uncertainty due to short hindcasts and non-stationarities in the climate system can contribute substantially to the occurrence of the 'signal-to-noise paradox'.

Further information can be found in an article by Antje Weisheimer et al. in the *Quarterly Journal of the Royal Meteorological Society*, doi:10.1002/qj.3446.

Improved user access to ECMWF publications

Helen Setchell, Paolo Battino

ECMWF publications are now easier to find and access following a combined effort in cataloguing, integrating third-party identifiers, development, and design.

A redesigned Publications landing page (https://www.ecmwf.int/en/publications) provides quick access to the most popular and important collections in ECMWF's eLibrary. The resulting search page makes it easier to filter publications that are part of the same collection, event or series. Search results for important ECMWF acronyms, such as IFS cycles (for example 'CY43R3'), have been improved. The eLibrary also now includes additional types of publication, previously not indexed, such as education material and software packages.

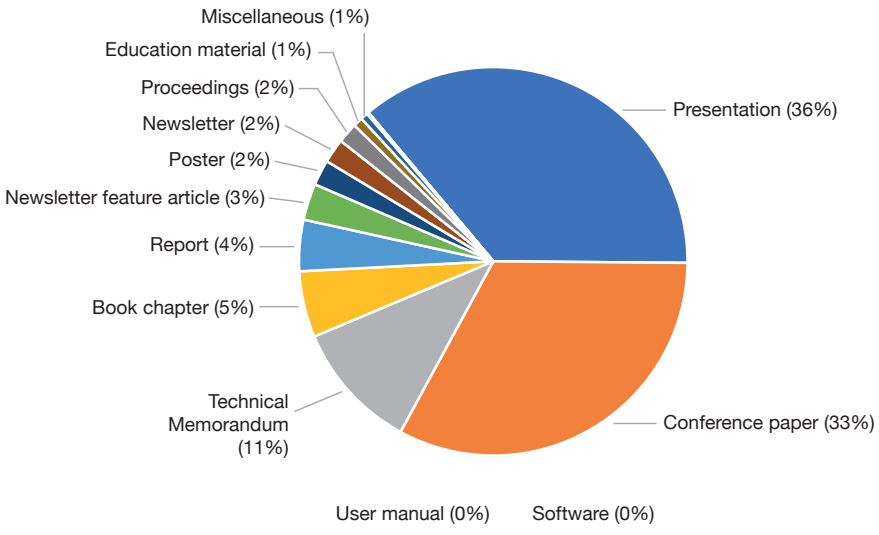

ECMWF eLibrary records by publication type. Presentations, conference papers and Technical Memoranda make up more than three quarters of ECMWF publication records.

Special pages elsewhere on the site have been created for Technical Memoranda (https://www.ecmwf.int/en/publications/technical-memoranda) and IFS Documentation (https://www.ecmwf.int/en/forecasts/documentation-and-support/changes-ecmwf-model/ifs-documentation) to address the specific ways users have told us they like to access these documents. Improvements in the publications metadata have enabled us to develop these user interfaces based on the same underlying catalogue, i.e. without the need to duplicate content.

There have been two other significant improvements which make the publications of ECMWF and its staff more accessible to users beyond our own website. The first is the use of persistent Digital Object Identifiers (DOIs) for ECMWF publications, which we have started to assign in batches to existing publications, beginning with ECMWF Newsletter feature articles. The second is integration with ORCID, a persistent digital identifier which uniquely distinguishes ECMWF staff and their publications.

This latter integration enables us to remove from our catalogue the records of staff publications which are published by others, and focus our efforts on ECMWF publications only. Staff are responsible for maintaining their own lists of publications. These will eventually be reused, among other things to populate our staff profile pages (https://www.ecmwf.int/en/about/who-we-are/staff-profiles).

We are always interested in hearing from our users on how we can improve the ECMWF website. If you have an idea, criticism, or suggestion, you can contact us via the website contact form (https://www.ecmwf.int/en/about/contact-us).

New 14-month forecast available for research

Ross N. Hoffman (NOAA/AOML, US), Sylvie Malardel (Météo-France, Réunion), Tanya Peevey (NOAA/CIRES, US)

A new uninterrupted 14-month-long forecast produced using a recent version of ECMWF's Integrated Forecasting System (IFS Cycle 43r1) is available to researchers. Such a forecast is known as a 'Nature Run' (NR) and is commonly used for Observing System Simulation Experiments (OSSEs). The new NR, called ECO1280, starts at 00:00 UTC on 30 September 2015 and includes the rapid intensification of Hurricane Joaquin. The ECO1280 uses the TCo1279L137 configuration of the IFS, which has a grid spacing of about 9 km and 137 vertical layers.

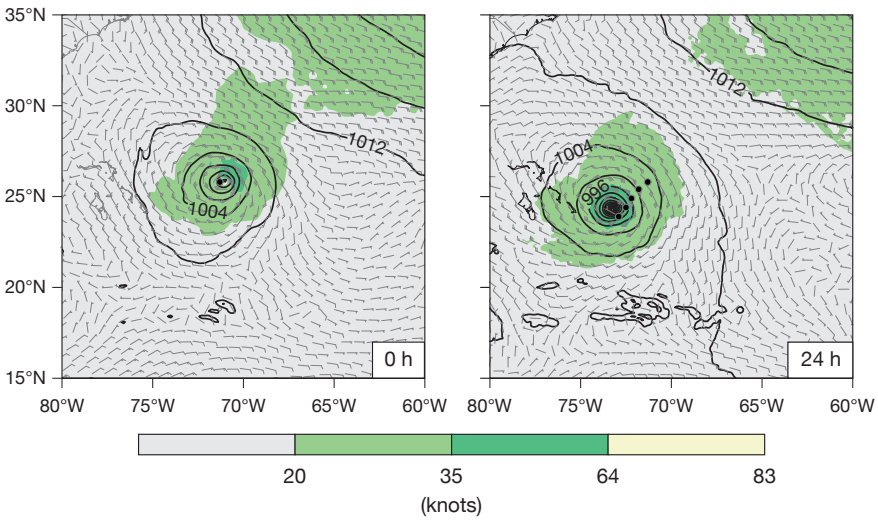

Hurricane Joaquin rapidly intensifies between 0 and 24 h in the ECO1280 as seen here in the surface pressure field (contours, in hPa) and 10 m wind field (shading). Joaquin is tracking towards the southwest. The black dots are best-track positions every six hours during the forecast period.

Motivation

An NR is used to represent the atmosphere in simulation experiments including OSSEs. Ideally the NR is generated by a state-of-the-art numerical model and realistically represents all phenomena that affect the new observing system. Since no forecast system is perfect, there should be realistic differences between the NR model and the model used for assimilation and forecasting. For researchers, the value of the new ECMWF NR is that the ECMWF model is very advanced, complete and accurate.

Data archive

All ECO1280 variables are archived every three hours. During the first month, data are archived every hour. Storage required for one month of data at three-hour archiving is about 4 TB. Instructions for downloading, converting, and interpolating the ECO1280 are given on the ECO1280 web page: https://www.cira.colostate.edu/imagery-data/ecmwf-nature-run/

Use and access

The ECO1280 may be used for any research purpose. Interested researchers should acquire the ECO1280 by contacting the Cooperative Institute for Research in the Atmosphere at Colorado State University (CIRA/CSU) to describe planned usage and to obtain data access instructions. Resulting publications or presentations must acknowledge ECMWF and CIRA/CSU. A useful aid to the design and interpretation of OSSEs is the OSSE checklist available at http://www.aoml.noaa.gov/qosap/osse-checklist/.

ECMWF joins Sea-Ice Drift Forecast Experiment

Helge F. Goessling (AWI, Germany), Steffen Tietsche (ECMWF), Simon Reifenberg (AWI), Axel Schweiger (University of Washington)

Since October 2018, ECMWF has contributed forecasts of the drift of Arctic sea-ice buoys derived from operational seasonal forecasts to the Sea-Ice Drift Forecast Experiment (SIDFEx). SIDFEx was launched in 2017 as part of the Year of Polar Prediction (YOPP). It is a community effort to collect and analyse Arctic sea-ice drift forecasts at lead times from days to a year, based on any method, for a number of sea-ice buoys and, ultimately, the research icebreaker Polarstern.

Aims of SIDFEx

SIDFEx is motivated in part by the need to determine an optimal start position for Polarstern's year-long MOSAiC drift across the Arctic commencing in autumn 2019. Specifically, it will test whether forecasting systems that account for initial conditions and simulate the evolving atmosphere, sea ice and ocean system can provide additional skill over drift forecasts that rely on historical sea-ice velocity fields. The MOSAiC drift provides an exciting opportunity to assess current capabilities to forecast sea-ice drift for a number of applications, ranging from logistics support for field experiments to potential search and rescue operations. In addition, the examination of sea-ice drift forecasts will help us to understand many aspects of the coupled atmosphere–sea-ice–ocean system and to identify and resolve model shortcomings. SIDFEx specifies and enforces a common data format for submission, but otherwise contributors are free to choose their method and time range. As of November 2018, ten groups from the US, Canada, Germany, Norway, and the UK have submitted almost 20,000 forecasts.

ECMWF's contribution

ECMWF contributes a fully coupled seasonal forecast every month in near real time. This has been made possible with the implementation of SEAS5 in November 2017, which predicts the movement of sea ice in response to surface winds, ocean currents, and internal forces within the ice. For the SIDFEx contribution, predicted trajectories of ice-tethered buoys are computed from SEAS5 daily-mean ice velocity fields archived on the native NEMO ocean grid. Although in principle these computations simply involve adding up daily displacements, the intricacies of the geometry of the tri-polar NEMO grid make this a non-trivial task. Therefore, we use the specialised software tool ARIANE, which has been developed by Bruno Blanke and Nicolas Grima at UBO/LOPS in Brest, France.

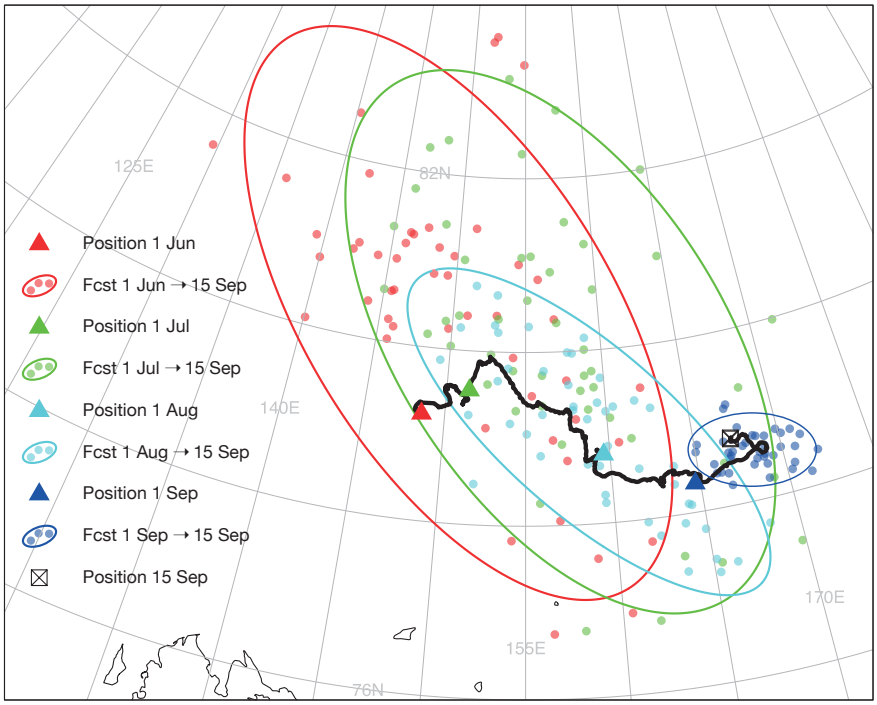

Drift forecasts for an Arctic sea-ice buoy. The forecasts are based on the ECMWF seasonal system SEAS5. Initial positions on the first day of successive months from June to September 2018 are marked by triangles on the observed drift trajectory (black curve). The dots denote ensemble forecast positions, all valid on 15 September 2018. Ellipses enclose 90% probability areas of a bivariate normal distribution fitted to the points comprising the forecasts. The verifying observed position is shown by the cross in the box.

First results

ECMWF trajectory forecasts have been produced and submitted to cover initialisation times from 1 June 2018 onward. The figure shows forecasts valid on 15 September for one of the SIDFEx buoys and four consecutive forecast start dates. It illustrates the considerable forecast uncertainty on the seasonal time scale, which decreases as the initial date approaches the target date. Although the buoy position observed on 15 September is close to the predicted buoy positions for individual ensemble members in each case, for two of the initial dates the observed position is outside the likely (90%) range of forecast outcomes. Results for other buoys are broadly similar. More cases are needed to establish whether this is chiefly the result of forecast model biases or of the atypical drift pattern that occurred this summer: the eastward drift in the Russian and Alaskan sectors was rather unusual. In any event, it is likely that there are forecast model biases which call for calibration to improve reliability, as is common practice for these timescales. As more trajectory

forecasts are added to the SIDFEx database, a clearer picture of the forecast quality will emerge.

Outlook

The SIDFEx community has only just started to exploit the constantly growing dataset, which is made publicly available in near real time (https://www.polarprediction.net/yopp-activities/sidfex/). ECMWF will benefit from the systematic forecast quality assessment of sea-ice drift carried out by the SIDFEx community, which will help to identify issues and ways to improve forecasts from days to seasons. The Centre is considering adding medium- and extended-range trajectory forecasts to its contribution. These could provide additional support for research and operations associated with the drifting icebreaker Polarstern.

Copernicus users remain highly satisfied

Xiaobo Yang

Users of the Copernicus Atmosphere Monitoring Service (CAMS) and the Copernicus Climate Change Service (C3S), which are both implemented by ECMWF on behalf of the EU, continue to be highly satisfied, according to a survey carried out by the Copernicus User Support Team at ECMWF. Between 26 June and 17 July 2018, 259 CAMS users and 1,323 C3S users took part in the survey. The overall satisfaction rating is 87.5% for CAMS (an average rating of 3.5 out of 4) and 86.25% for C3S (3.45 out of 4).

CAMS

Compared to the 2017 survey, the commercial sector's share of the use of CAMS increased from 18% to 22%, while that of the public sector decreased from 18% to 15%. Responses from EU-based users show a similar pattern: the EU commercial sector's share of the use of CAMS increased significantly from 21% to 28%, while that of the public sector decreased from 20% to 16%.

Most users are aware of data provision, documentation and user support services. Users consider these services useful, with data access ranked most useful. All non-data services receive high satisfaction ratings.

Sixty-one per cent of all responses came from the EU. Eastern EU countries are still underrepresented and efforts to increase user uptake in eastern Europe will continue. The overall user satisfaction rating is the same as last year, which indicates that satisfaction with CAMS continues to be high and that there is no need for any radical change to the CAMS strategy and portfolio.

C3S

Most respondents (74%) came from outside the EU. This reflects the fact that global reanalysis datasets – the most popular datasets in the past year – have been widely used outside the EU. Among EU respondents, most are based in western and central Europe. Future C3S user uptake activities should consider targeting specifically users in eastern Europe. The growing interest in C3S from the EU private sector seen in the survey (15% of all EU respondents, up from 13% in the previous survey) is encouraging and suggests that C3S is on track to further diversify its user base.

Climate reanalysis data, comprising the ERA5 and ERA-Interim datasets, are by far the most popular and receive the highest satisfaction ratings. Many other C3S products and services are not well known, partly because products are still under development (e.g. seasonal forecasts), are new (e.g. the European State of the Climate report), or cater for rather specialist interests (e.g. monthly sea-ice maps). It is expected that this will change as the C3S Climate Data Store (CDS) is developed further and used more widely.

C3S released the CDS to the public on 14 June 2018. An open question in the survey asked users about their expectations. The replies included many ideas and suggestions for the CDS which are already part of the CDS roadmap, indicating that C3S is well aligned with user requirements.

Over 90% of survey participants gave the service three or four stars out of four, indicating great satisfaction with the usefulness and quality of C3S products and services. The current service strategy appears to meet users' needs.

Full survey reports are available on the news pages of the CAMS (http://atmosphere.copernicus.eu) and C3S (http://climate.copernicus.eu) websites.

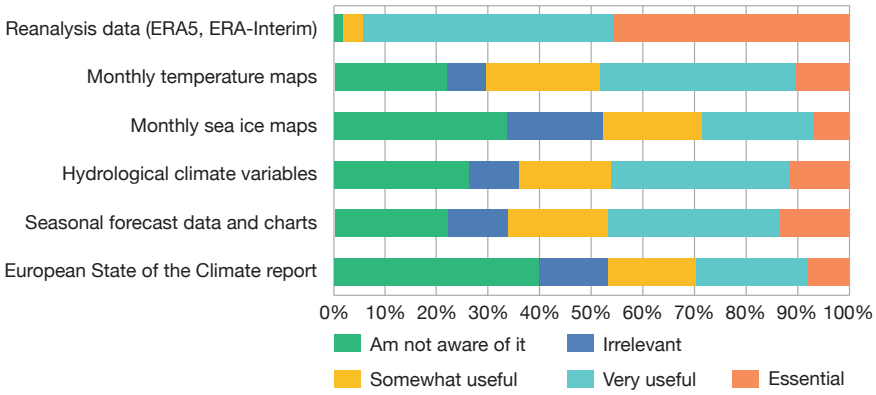

Usefulness of C3S products and services. C3S global reanalysis datasets are rated most highly, with more than 94% of respondents saying that they are 'very useful' or 'essential'.

A new product to flag up the risk of cold spells in Europe weeks ahead

Laura Ferranti, Linus Magnusson, Frédéric Vitart, David Richardson

Persistent severe cold temperatures pose serious threats to health and welfare. There is therefore a demand for the development of early warning systems that could allow more time for mitigating actions. In the medium range (up to ten days ahead), predictions of severe temperature conditions can be directly based on temperature forecast probabilities. As forecasts move beyond ten days, on the other hand, increasingly they cannot accurately represent day-to-day weather variability at individual grid points. They may, however, capture large-scale circulation patterns that last more than a week. Early indications of such patterns can help forecasters to identify the potential for temperature anomalies in the extended range. For this reason, ECMWF has developed a new product that visualises extended-range predictions for weather patterns in the Euro-Atlantic region associated with persistent cold spells in Europe. The predictions have been found to have useful skill up to about 2.5 weeks ahead.

Role of persistent highs

In Europe, anomalous surface weather may be the result of persistent high-pressure systems, also called blocking highs. Blocking highs tend to be nearly stationary for more than a week, disrupting the usual eastward progression of weather systems. Severe cold episodes in winter as well as dry spells and heat waves in summer are often associated with the occurrence of blocking. For example, during the August 2003 heat wave, the hot dry tropical continental air mass that characterised this event was pushed over western Europe by a persistent anti-cyclonic system over Central Europe. Similarly, the series of severe cold spells over northern and western Europe that occurred in winter 2009/2010 was associated partly with a blocking high over Scandinavia, and partly with a record persistence of anomalously high pressure over Greenland and low pressure over the Azores, a pattern known as the negative phase of the North Atlantic Oscillation (NAO-). NAO- brings a substantial reduction of the westerly flow across the Atlantic, and in Europe a strengthening of northerly winds from the Arctic. The large spatial scale and the low-frequency nature of such circulation patterns are crucial factors for successful predictions at the extended timescale. The link between flow patterns such as the NAO with tropical forcing and sudden warming events in the lower stratosphere (SSW) is another factor. In fact, during low-frequency phenomena

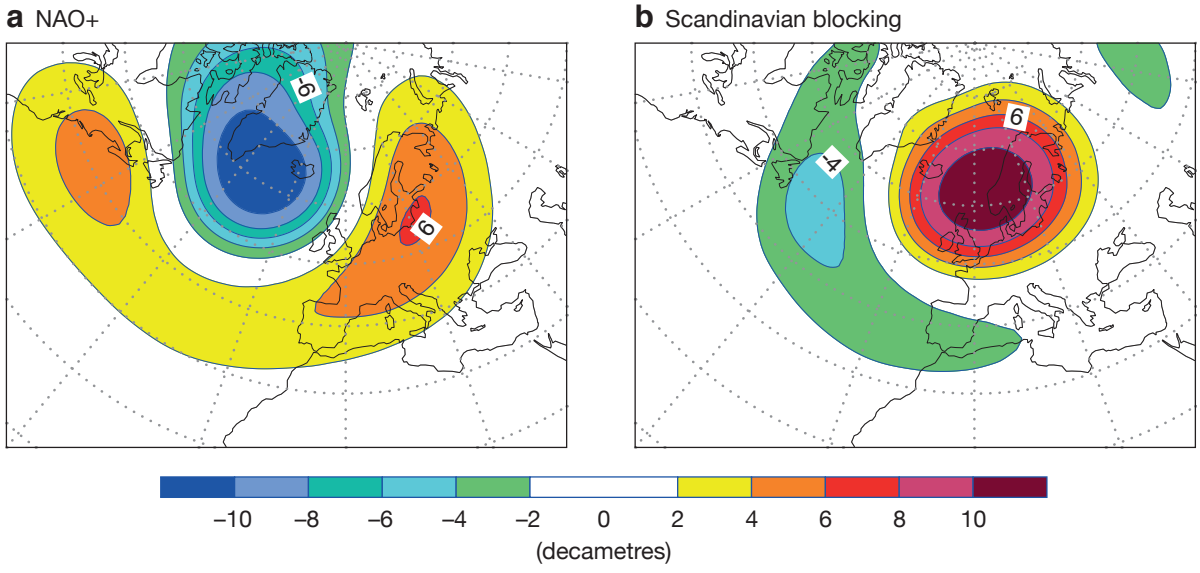

FIGURE 1 The two dominant wintertime patterns for 500 hPa geopotential height anomalies in the Euro-Atlantic region are (a) NAO+ and (b) Scandinavian blocking.

doi: 10.21957/k2rlf88oe1

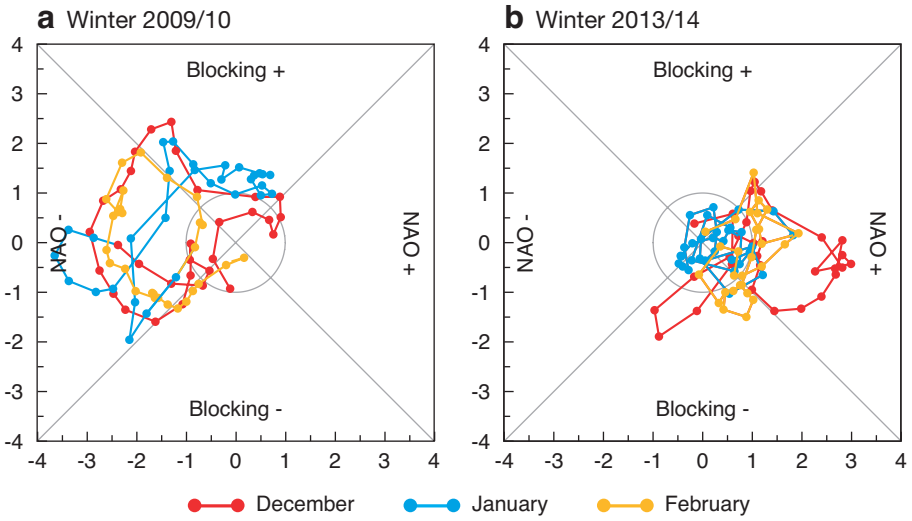

FIGURE 2 NAO–BLO diagrams showing the daily evolution of the analysis for the two dominant patterns for (a) winter 2009/10 and (b) winter 2013/14.

such as El Niño-Southern Oscillation (ENSO), SSW and Madden–Julian Oscillation (MJO) events, the skill of northern hemisphere forecasts is enhanced, creating a window of opportunity for extended-range predictions.

A new product

In order to visualise predicted changes between circulation patterns associated with high-impact temperature anomalies over Europe, we have developed a new product based on two dominant circulation patterns explaining the largest part of atmospheric variability over Europe (Figure 1). The first pattern (Figure 1a) represents the typical structure of the positive phase of the NAO circulation, and the second, with a high centred over Scandinavia and a low to the east over the Atlantic Ocean, describes the anomalous flow during northern European blocking events. The methodology used to identify the two dominant patterns is described by Ferranti et al. (2018).

These two circulation modes can be used to construct a diagram in which any geopotential height anomaly over the Euro-Atlantic region is characterised by the contribution made to it by the NAO± pattern ('NAO') and the blocking/anti-blocking (high/trough over Scandinavia) pattern ('BLO'). To give an example of the use of the NAO–BLO diagram, Figure 2 shows the daily evolution of the analysed 500 hPa geopotential anomalies during the winters of 2009/10 and 2013/14, respectively. Every dot in the diagram shows the amplitude of NAO and BLO anomalies on a given day. For example, the further to the left a dot in the diagram is, the closer the meteorological situation is to the NAO- pattern. The amplitude of the NAO and BLO circulation pattern anomalies is calculated relative to the standard deviation of the respective climatological distribution, which is represented by the circle. Large (small) amplitudes are outside (inside) the circle.

Figure 2a shows that, during most of December 2009 and February 2010, the flow circulation strongly resembled the negative phase of the NAO pattern, while for much of January 2010 the circulation was characterised by Scandinavian blocking. The persistent southward advection of cold Arctic air or westward advection of cold continental air associated with these patterns resulted in record-breaking cold temperatures over Europe. In contrast, the winter of 2013/2014 (Figure 2b) was dominated by NAO+ and westerly flow anomalies across the Atlantic. Consistent with the anomalous flow conditions, the winter of 2013/2014 saw a series of storms and severe rainfall but rather mild temperatures over Europe.

It might appear that the 2-dimensional NAO–BLO diagram provides a somewhat limited view of the complex structure of extratropical variability. In fact, most of the Euro–Atlantic low-frequency winter variability is typically explained by variations of four climatological regimes represented by the positive phase of the NAO; the negative phase of the NAO; blocking; and Atlantic ridge patterns. Nevertheless, since NAO- and blocking are the patterns typically associated with severe and persistent temperature anomalies, the NAO–BLO diagram provides a more appropriate and simplified framework to help predict temperature extremes.

Relationship with cold spells

The systematic relationship between the NAO–BLO diagram and severe European cold spells is highlighted in Figure 3. The figure shows the distribution of severe cold events in the NAO–BLO diagram for November to February winter periods from 1980 to 2015. The cold events were identified using 2 m temperature reanalysis data. They are shown in the NAO–BLO diagram for nine predefined regions covering the European domain.

A severe event was defined as an event in which daily mean temperatures are cooler than the 10th percentile

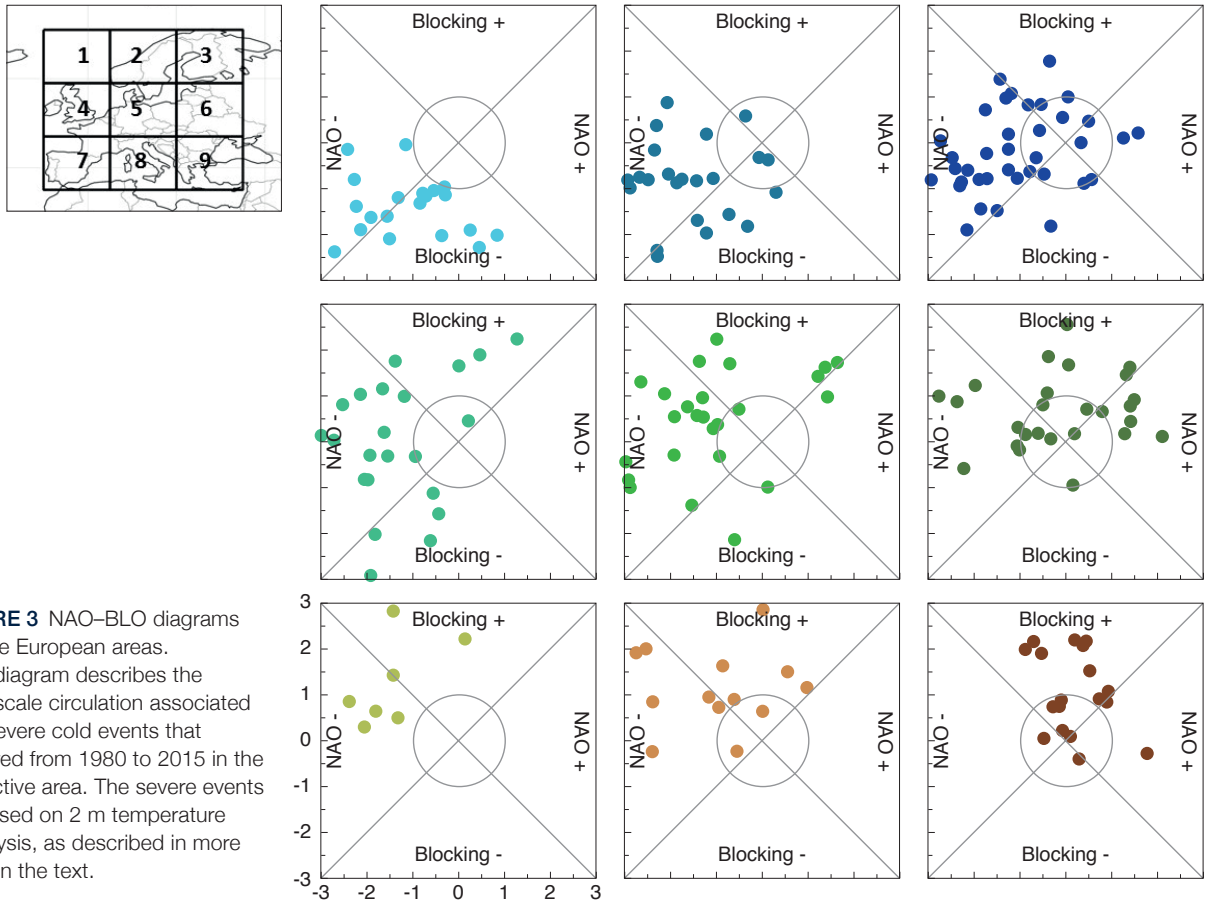

FIGURE 3 NAO–BLO diagrams for nine European areas. Each diagram describes the large-scale circulation associated with severe cold events that occurred from 1980 to 2015 in the respective area. The severe events are based on 2 m temperature reanalysis, as described in more detail in the text.

of the daily climate for at least 60% of the grid points located in the respective region and in which this criterion is satisfied for four consecutive days. The location of each event in the NAO–BLO diagram was determined from the reanalysis 500 hPa geopotential field from the first day of the event. For each of the nine regions, the NAO–BLO diagram shows all events detected in that region. For the northern regions 2 and 3, most severe cold episodes correspond to NAO- events with rather large amplitudes. For region 1, a roughly equal number of cold episodes are characterized by the NAO- and the anti-blocking type of circulation. For the central and southern regions, severe cold events are associated with NAO- and with blocking. The amplitudes in the NAO–BLO diagram are generally very large, indicating that this diagram can be used to describe most severe cold episodes. For the eastern European domain (regions 3, 6 and 9), there is a larger number of cases with relatively small amplitudes. This indicates that for these regions the NAO–BLO diagram is less suitable for representing severe cold spells. However, overall Figure 3 shows that there is a strong link between large amplitudes in the NAO- and BLO sectors and the occurrence of severe cold spells over Europe.

An illustration of the predicted evolution of atmospheric flow in the Euro-Atlantic region, as represented by ensemble forecast (ENS) trajectories in the NAO–BLO diagram, is given in Figure 4. The forecast trajectories from day 0 to day 10 of the ENS initiated on 23 February 2018 indicate a clear transition from Scandinavian blocking to NAO- by day 5. The evolution of forecast uncertainties (rather low in this case) is well depicted by the NAO–BLO diagram. The verifying analysis and the ensemble mean trajectory match quite closely. Indeed, in late February to early March, a spell of severe low temperatures with significant snowfall was associated with the establishment of a large high-pressure system over Scandinavia, which later developed into a NAO- circulation.

In the extended range, the day-to-day trajectories are affected by substantially larger uncertainty. Therefore, the evolution of forecast anomalies is represented by probability density functions (PDF) based on daily values of individual members in a given week.

Figure 5 shows the PDFs for forecasts starting on 15 February and 19 February 2018 for two subsequent target weeks. During the week starting 26 February 2018, the analysed NAO- daily values have very large amplitudes in the tail of the NAO climatological distribution. In that week, mean temperature anomalies ranged between –6°C and –3°C over most of Europe. The PDF for that week for the forecast starting on

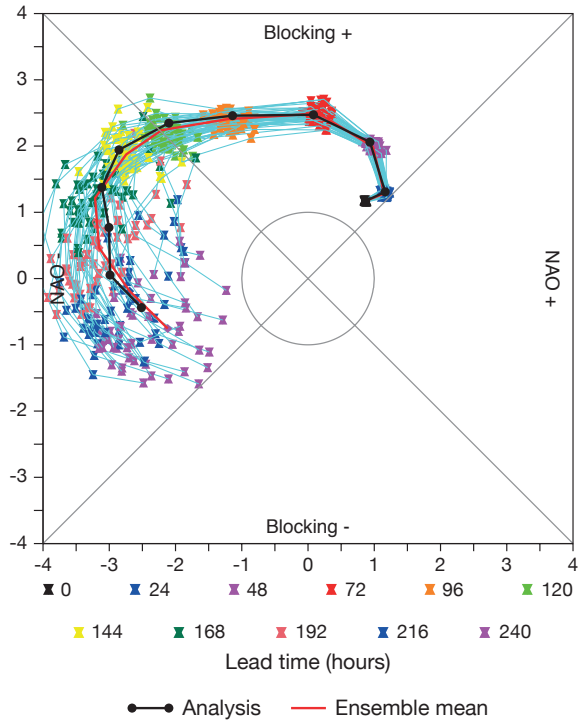

FIGURE 4 Evolution of ENS forecast up to day 10 in the NAO–BLO diagram. The ENS forecast starts at 00 UTC on 22 February 2018.

19 February indicates high probabilities for strong blocking and extreme NAO- values (Figure 5a). The PDF for the same week from the forecast starting on 15 February has its absolute maximum over the NAO- sector (Figure 5b). It thus provides a strong indication of the likelihood of an NAO- event. However, at this range the amplitude of the event is underestimated, and the uncertainties are rather large. During the subsequent week, starting 5 March, the NAO- persisted while its amplitude gradually reduced. The cold anomalies persisted mainly over Scandinavia and eastern Europe. The PDFs from forecasts for that week starting on 19 and 15 February provide a good indication of the likelihood of NAO- persistence as well as of its reduced amplitude (Figure 5c,d).

The change in atmospheric circulation that led to the cold spell was predicted about 2.5 weeks in advance. The high level of accuracy in this forecast is likely to be

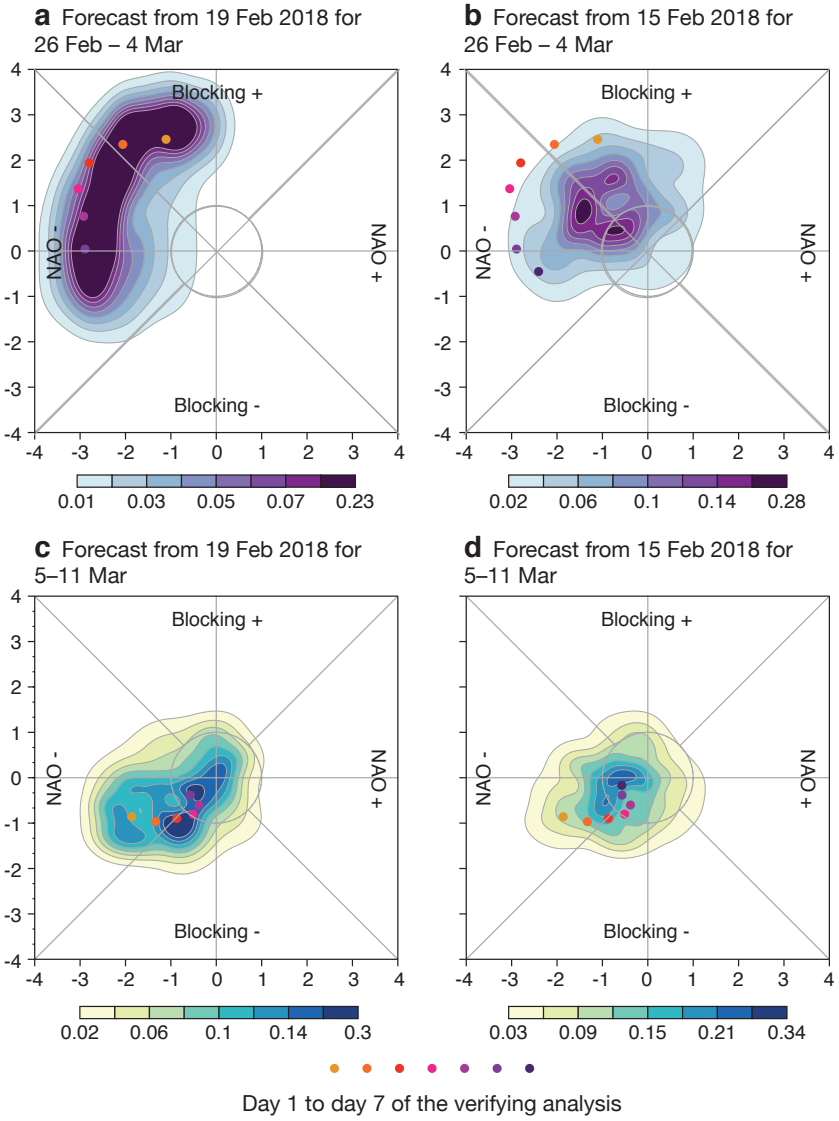

FIGURE 5 Probability density functions for (a) an ensemble forecast starting on 19 February 2018 for the week starting on 26 February, (b) an ensemble forecast starting on 15 February 2018 for the same week, (c) an ensemble forecast starting on 19 February 2018 for the week starting on 5 March and (d) an ensemble forecast starting on 15 February 2018 for the same week. Daily values of the verifying analysis are represented by dots.

linked to two elements: a very intense MJO event propagating across the West Pacific, and a stratospheric warming event (SSW) that took place around 11 February. The temperature anomaly vertical cross section (not shown) indicates that the SSW was probably not a crucial factor in bringing about the change in circulation although it increased the amplitude of the anomalies.

Skill assessment

First, we assessed the skill in predicting NAO and BLO values separately, then we considered the skill in predicting the evolution in the NAO–BLO diagram to better understand the ability to capture transitions between different flow patterns. Figures 6a and 6b show the skill in predicting the amplitudes of the NAO and BLO patterns, respectively. In this assessment, we considered ECMWF's ensemble forecasts and five additional ensemble systems available from the sub-seasonal to seasonal (S2S) prediction research project archive. The S2S project, established by the World Weather Research Programme/World Climate Research Programme, has constructed an extensive database containing sub-seasonal forecasts (up to 60 days ahead), three weeks behind real time, and reforecasts from 11 operational centres. The skill metric is the anomaly correlation coefficient (ACC) between the observed NAO/BLO amplitudes and the NAO/BLO amplitudes predicted by the ensemble-mean forecast, computed for the common period of reforecasts available in the S2S archive covering 12 years (December/January/February 1999–2010). Since beyond two weeks ahead the forecast is not expected to have day-to-day accuracy, a 5-day running mean (centred on the day) was applied to the verifying analysis and the ensemble mean prior to the ACC computation. We use the S2S models to estimate the current range of skill of sub-seasonal predictions.

Looking at the forecast range at which the ACC drops below 0.5, the limit of skill ranges from 11 to 17 days for the NAO predictions and from 9 to 13 days for the BLO predictions. The skill of ECMWF forecasts in predicting BLO drops below 0.5 at about 13 days – a few days earlier than for NAO. Probabilistic skill scores for S2S predictions of NAO and BLO are consistent with the ACC values obtained for the ensemble mean (not shown).

The accuracy of the forecast trajectories in the NAO–BLO diagram is evaluated by assessing the temporal correlation between the predicted and analysed NAO and BLO amplitudes (Figure 6c). This metric, known as bivariate correlation, has been documented by Gottschalck et al. (2010). The bivariate correlation can generally be used to measure the temporal correlation between two vectors defined in a two- or higher-dimensional space. It is thus appropriate for assessing trajectories in the NAO–BLO diagram. Although this score measures only the accuracy of the forecast trajectories without highlighting the prediction skill in terms of strong versus weak events, it provides an objective skill measure for forecasting transitions between the circulation patterns

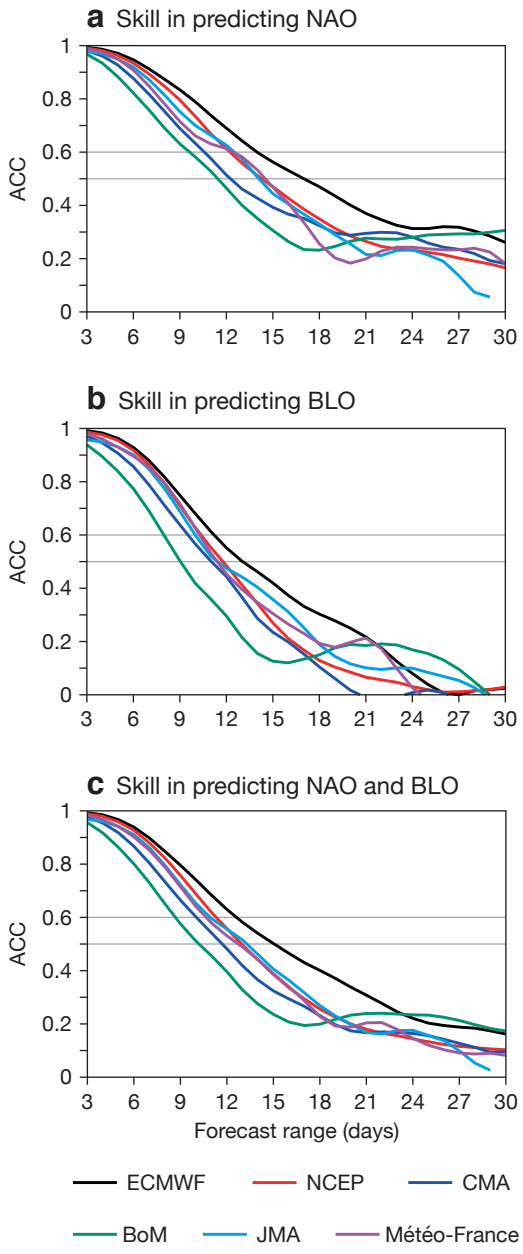

FIGURE 6 Ensemble mean anomaly correlation coefficient (ACC) as a function of forecast lead time for the prediction of (a) the NAO+/NAO- pattern (westerly/easterly flow across the Atlantic) and (b) the blocking/anti-blocking pattern; and (c) bivariate correlation as a function of forecast lead time for the prediction of the combined pattern. ACC values are based on a 5-day running mean applied to the forecasts and verifying analysis data. Skill scores have been calculated for forecasts produced by ECMWF, NCEP (US National Centers for Environmental Prediction), CMA (China Meteorological Administration), BoM (Australian Bureau of Meteorology), JMA (Japanese Meteorological Agency), and Météo-France.

associated with high-impact weather over Europe. The lead times at which the bivariate correlation coefficient drops below 0.5 range from 10 to 15 days for S2S forecasts, indicating some potential to predict the onset of cold spells over Europe beyond the medium range.

Conclusion and outlook

The ability to predict the onset of a period with severe temperature anomalies weeks ahead is closely linked with the ability to accurately forecast the evolution of anomalies in the large-scale atmospheric circulation. Therefore, reliable extended-range forecasts of flow patterns such the NAO and blocking can support the prediction of severe cold events over Europe. The NAO–BLO diagram is an effective tool to assess the likelihood of regime transitions associated with the occurrence of severe cold episodes in the exended range. The NAO–BLO forecast trajectory and PDF diagrams are currently available to registered users as test products at: *https://confluence.ecmwf.int/display/FCST/Test+products*. The success of forecasting, weeks ahead, changes in large-scale flow that lead to cold conditions depends on the type of transition. The ECMWF ensemble is able to provide reliable probabilities of cold conditions associated with the establishment of the NAO- pattern beyond the medium range. The predictability of such events is enhanced by tropical–extratropical teleconnections resulting from MJO activity. On the other hand, providing probabilities in the extended range for the occurrence of cold events associated with a transition to blocking presents a bigger challenge. The skill of blocking predictions is not highly sensitive to the existence of an MJO in the initial conditions and, consistent with that, blocking exhibits lower predictability than NAO-. Understanding these flow-dependent variations in forecast skill, and using the new NAO–BLO diagrams, will help users to exploit periods of enhanced extended-range predictability. An assessment of whether it is possible to develop a similar product for European heat waves is planned for 2019.

Further reading

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

Gottschalck, **J.**, **M. Wheeler**, **K. Weickmann**, **F. Vitart**, **N. Savage**, **H. Lin**, **H. Hendon**, **D. Waliser**, **K. Sperber**, **M. Nakagawa**, **C. Prestrelo**, **M. Flatau** & **W. Higgins**, 2010: A framework for assessing operational Madden–Julian Oscillation forecasts. *BAMS*, 1247–1258. doi:10.1175/2010BAMS2816.1

Continuous data assimilation for the IFS

Peter Lean, Massimo Bonavita, Elías Hólm, Niels Bormann, Tony McNally

Accurate weather predictions are not possible without an accurate specification of the initial state of the Earth system. To this end, billions of observations of the Earth system are made every day. Data assimilation combines many of those observations with model information to arrive at a full set of initial conditions called the analysis. In the current ECMWF operational configuration, by the time the analysis is complete, the most recent observations that have gone into producing it are almost two hours old. This article presents a revised configuration of ECMWF's 4D-Var data assimilation system designed to allow the analysis to benefit from more recent observations. In the new, more continuous framework, we are able to assimilate observations taken around one and a half hours later than in the current system. In addition, more accurate albeit more time-consuming data assimilation configurations can be accommodated inside the operational schedule, providing scope for further improvements in analysis and forecast skill in upcoming cycles. The new data assimilation configuration has been found to improve forecast quality significantly. It is due to be implemented in the next upgrade of ECMWF's Integrated Forecasting System (IFS Cycle 46r1), scheduled for 2019.

Operational schedule

For an observation to be useful for numerical weather prediction (NWP), it must be delivered in a timely manner. To allow for the inevitable delay between the time that an observation is made and the time at which it is available for use, most operational centres (including ECMWF) structure their operational data assimilation (DA) and forecasting schedule in two distinct phases. First, there is a data collection phase, during which we wait for the observations to arrive. After a predefined data cut-off time, the computation of the analysis and the forecast begins (see Figure 1).

Since the time at which forecasts are issued to users is usually fixed, the choice of the data cut-off time involves a trade-off between two competing requirements: allowing more time for observations to arrive (which should lead to a more accurate analysis and hence better forecasts); and allowing more time for the DA and forecast computations to be completed (which should also result in a more accurate analysis and better forecasts through, e.g., more sophisticated algorithms, increased resolution, etc.).

In ECMWF's current operational schedule, the DA computations do not begin until an hour after the end of the assimilation window (the time during which observations to be used in a data assimilation cycle are made). For example, computing the 00 UTC analysis with a 21 to 03 UTC assimilation window starts at around 04 UTC. The observation quality control and DA computations take about an hour to complete. This means that by the time the analysis has been produced, the most recent observations that went into producing it are about two hours old.

Current DA configuration

Since its operational implementation in the late nineties (Rabier et al., 2000), ECMWF has used incremental 4D-Var as its atmospheric DA algorithm. Incremental 4D-Var (Courtier et al., 1994) works by iteratively minimising a cost function to achieve the best possible fit between a short-range forecast (the background) on the one hand and observations on the other. The model and the observation operators are repeatedly linearised

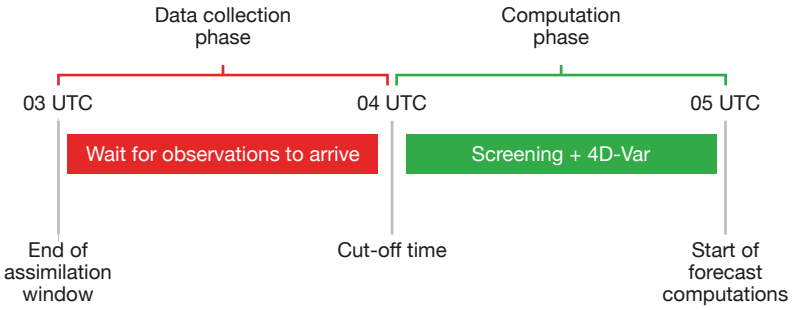

FIGURE 1 Schematic representation of the data collection and computation phases of the current operational DA schedule.

doi: 10.21957/9pl5fc37it

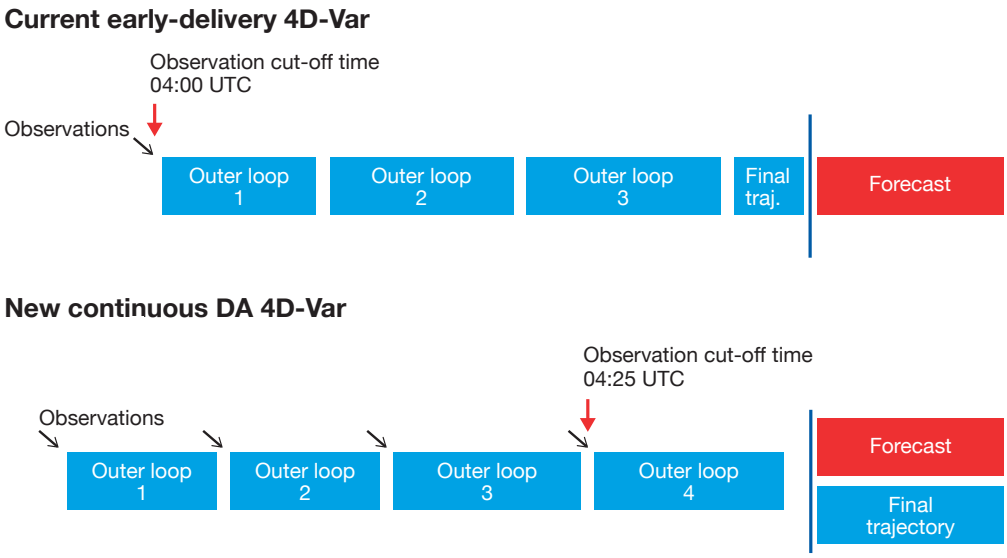

FIGURE 2 Schematic representation of the current early-delivery 4D-Var configuration and the proposed continuous DA 4D-Var configuration. The black arrows indicate at which stage of the 4D-Var computations new observations are ingested. The blue vertical line denotes the time when the forecast computations start, which is unchanged.

around a progressively more accurate model trajectory solution. These relinearisations are called outer loops. The outer loop mechanism has recently been shown to be one of the key drivers of analysis and forecast accuracy in the IFS (Bonavita et al., 2018). It is, however, a strictly sequential algorithm in the sense that successive linearisations and minimisations need to be performed one after the other for the procedure to converge. This imposes rather stringent limits on the complexity of the incremental 4D-Var configurations that can be run within ECMWF's operational time constraints.

A partial answer to the time-constraint problems described above was provided with the introduction of the early-delivery suite at ECMWF (Haseler, 2004). As the data cut-off is only one hour after the end of the assimilation window, many observations have not arrived by this time. A second, 'delayed cut-off' assimilation cycle is run which assimilates all observations that have arrived within 4 hours of the end of the assimilation window. The delayed cut-off cycle is thus able to make use of a much larger number of observations and thus to provide an analysis (called the long-window data assimilation analysis, or LWDA analysis) of higher accuracy than the one produced by the early-delivery suite. The LWDA analysis is then used to produce a short-range forecast (the background) which is used in the next early-delivery cycle. In this way, late-arriving observations can benefit the quality of the following analysis through a more accurate background. However, the limitations on the complexity of the 4D-Var algorithm that can be used in operational practice remain.

Continuous DA configuration

The general idea of a continuous DA system based on 4D-Var is not new. Variations on it were proposed in the mid-nineties (Järvinen et al., 1996; Pires et al., 1996) and even used operationally for a while at Environment Canada (Gauthier et al., 2007). In a continuous DA framework, we do not wait for all the observations to arrive before starting the analysis computations. In effect, the computation phase overlaps with the data collection phase. We call these schemes 'continuous DA' as, in principle, they can be run continuously in the operational schedule with new observations being fed to the assimilation system as they arrive. Motivated by the aim of allowing more recent observations into the analysis, a variant of this concept will be introduced in IFS Cycle 46r1.

The incremental 4D-Var outer loop provides a convenient mechanism by which newly arrived observations can be introduced into the system after the assimilation has started. The idea is that, instead of stopping observations entering the 4D-Var analysis after a fixed cut-off time, new observations are allowed in between successive outer loops (Figure 2). As each outer loop takes about 15 minutes to complete, in the new continuous DA we have extended the effective cut-off time by around 25 minutes.

From an algorithmic point of view, instead of solving a fixed minimisation problem, we solve a series of slightly different minimisation problems as the number of observations increases from one outer loop to the next.

To fully benefit from the later cut-off time, we have extended the assimilation window of the early-delivery analysis from 6 hours to 8 hours. This ensures that it extends right up to the time at which the DA runs. In our current system, the 00 UTC analysis has an assimilation window from 21 UTC to 03 UTC. The DA for this cycle begins after the 04 UTC cut-off time. Many observations taken in the 03 to 04 UTC period have already arrived by this time, but are not used. By

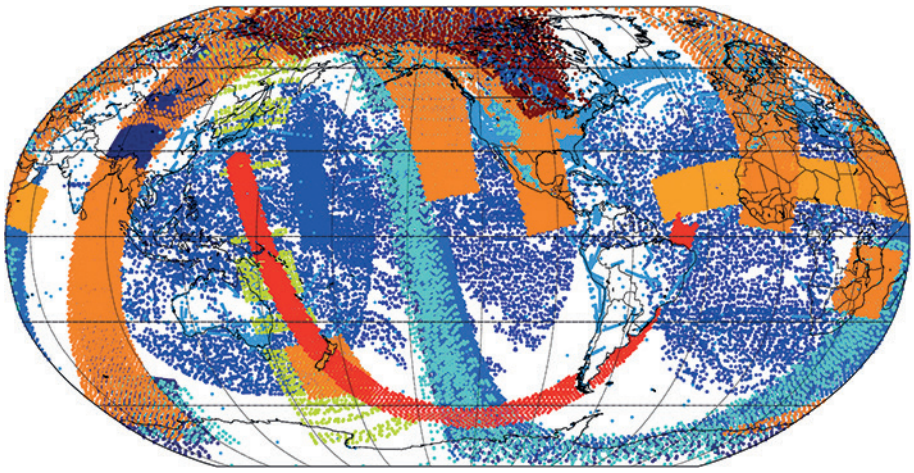

FIGURE 3 Example of extra observations assimilated in a single continuous DA cycle compared to the current operational setup. They include satellite observations from a large number of instruments as well as in situ measurements.

extending the assimilation window by two hours to 05 UTC, we can assimilate these very valuable observations and later ones, enabling the state to be constrained by observations a further one and a half hours into the forecast.

All the newly arrived observations need to also go through quality control. This is achieved by performing observation screening in each outer loop on all available observations (not just those that have arrived since the previous outer loop). An advantage of this approach is that the quality control is performed against a more accurate model trajectory than the previous model background, which should lead to improved screening decisions.

Our previous operational schedule was based upon a trade-off between the time allocated to the arrival of observations on the one hand and the time allocated to the DA computations on the other. In the continuous DA framework, we no longer need to wait until 04 UTC to begin the 00 UTC analysis, as later-arriving observations will be captured in successive 4D-Var outer loops. As a result, we are now able to start the 4D-Var computations earlier. In Cycle 46r1, we will start the assimilation around 10 minutes earlier and use this time to increase the number of outer loops from three to four. Even this simple change has been shown to have a statistically significant positive impact on analysis and forecast skill (Bonavita et al., 2018).

The changes described so far are applied to the time-critical early-delivery assimilation. As the LWDA cycles (which provide the background for the early-delivery assimilation) already have a very late observation cut-off time, there is no need to use the continuous DA approach here, and the length of the assimilation window is kept at 12 hours. The only change to the LWDA cycles is the addition of an extra outer loop to further improve the accuracy of the background.

Outer Loop	Number of observations in EXP relative to CTRL
1	101%
2	107%
3	110%
4	114%

TABLE 1 Number of observations used in each outer loop of 4D-Var in a six-month continuous DA experiment (EXP), relative to the number of observations used in the current operational setup (CTRL).

Outer Loop	Average number of iterations in CTRL	Average number of iterations in EXP
1	31.5	31.1
2	29.9	31.2
3	30.4	28.8
4	-	29.6

TABLE 2 Average number of iterations required to achieve convergence in the 4D-Var DA system in a six-month experiment using the operational setup (CTRL) on the one hand and continuous DA on the other (EXP).

Results

In the new early-delivery assimilation, the number of observations increases in each outer loop. In an experiment covering six months (EXP), the first outer loop had slightly more observations than the first outer loop in an experiment using the current system (CTRL), even though it has an earlier cut-off (see Table 1). The reason is that the assimilation window in EXP is longer than in CTRL. Overall, the continuous DA configuration

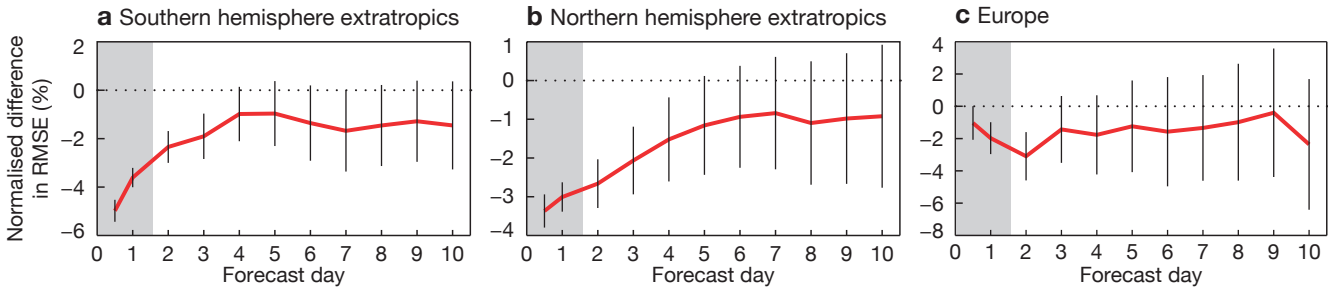

FIGURE 4 Normalised difference in root-mean-square error (RMSE) between forecasts of geopotential height at 500 hPa produced using the current early-delivery DA configuration on the one hand and continuous DA on the other, for (a) the southern hemisphere extratropics (20°S to 90°S), (b) the northern hemisphere extratropics (20°N to 90°N), and (c) Europe (35°N to 70°N and 10°W to 40°E), verified against the operational analysis. Negative values indicate smaller RMSE when using continuous DA. Vertical bars indicate 95% confidence intervals. The plots are based on six months of experimentation between December 2016 and August 2017. Verification in the early stages of the forecast (shaded grey) is sensitive to the choice of reference used and may be unreliable.

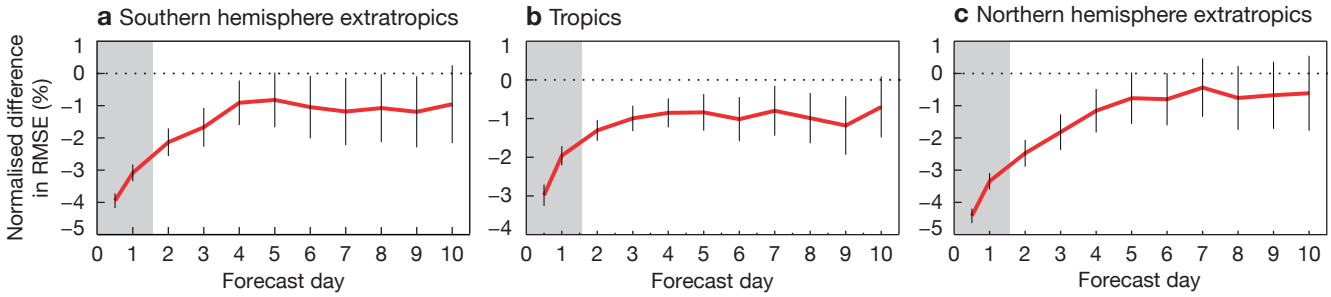

FIGURE 5 Normalised difference in root-mean-square error (RMSE) between forecasts of vector wind at 500 hPa produced using the current early-delivery DA configuration on the one hand and continuous DA on the other, for (a) the southern hemisphere extratropics (20°S to 90°S), (b) the tropics (20°S to 20°N) and (c) the northern hemisphere extratropics (20°N to 90°N), verified against the operational analysis. Negative values indicate smaller RMSE when using continuous DA. Vertical bars indicate 95% confidence intervals. The plots are based on six months of experimentation between December 2016 and August 2017. Verification in the early stages of the forecast (shaded grey) is sensitive to the choice of reference used and may be unreliable.

uses around 14% more observations than the current early-delivery DA. An example of the geographic distribution of the extra observations assimilated in a single cycle using continuous DA is shown in Figure 3.

Before running these experiments, an initial concern was that introducing new observations at each outer loop changes the minimisation problem being solved. It was thought that this might make the preconditioning techniques used to accelerate the convergence of 4D-Var less effective and might even cause numerical problems. Table 2 shows that there were no issues in this respect. A possible explanation is that the relatively small change in observation counts from one outer loop to the next (up to about 6% of total observation numbers) does not change the topology of the cost function significantly. Another reason is that the preconditioners represent only broad, large-scale features of the underlying cost function and thus tend to be insensitive to small, localised changes in observation coverage.

In general, medium-range forecast scores are improved by around 2% at day 3 in the extratropics for most variables. Examples of the improvements can be seen in Figures 4 and 5. In the tropics, 500 hPa wind shows a statistically significant improvement out to day 9. The improvements are widespread and geographically fairly uniform (Figure 6). These results are consistent with an improvement in predictability of around 2 to 3 hours. This comes from the extra one and a half hours of observational data, combined with the improved assimilation enabled by the extra outer loop.

Observation timeliness

Up until now, any observations that arrive within 1 hour of the time that they were made will be processed in the early-delivery cycle. With a continuous DA configuration, the analysis quality becomes more responsive to changes in the timeliness of observation delivery.

Figure 7 shows an example of the geographical distribution of observations made between 04:00 and 04:25 UTC (towards the end of the extended 4D-Var assimilation window) and assimilated in the continuous DA configuration. The coverage is far sparser than for observations made during earlier parts of the assimilation window. Only observations which are delivered within at most 25 minutes of the observation

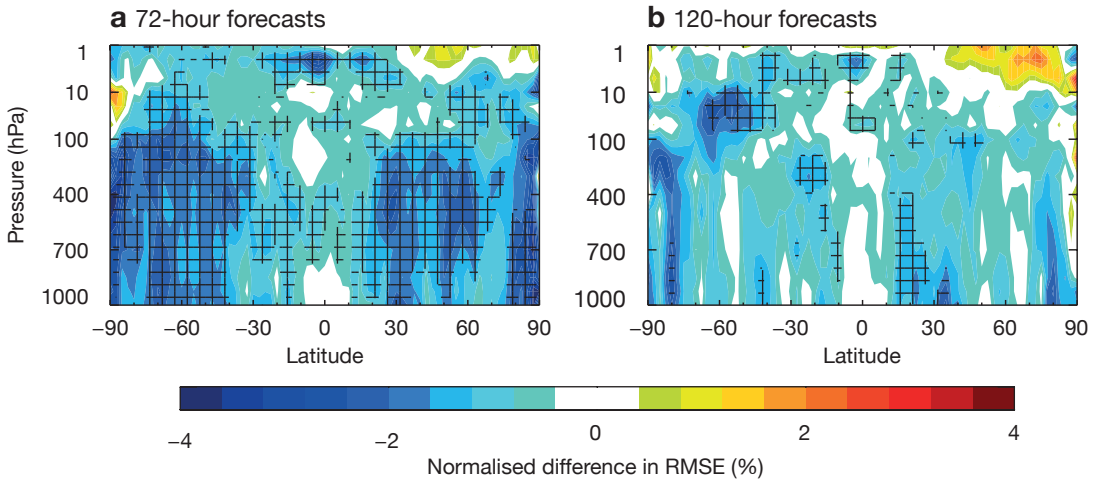

FIGURE 6 Zonal (east–west) mean difference in root-mean-square error (RMSE) between forecasts of vector wind produced using the current early-delivery DA configuration on the one hand and continuous DA on the other, for lead times of (a) 72 hours and (b) 120 hours, verified against the operational analysis. Negative values indicate smaller RMSE when using continuous DA. Hatched areas indicate statistically significant changes at the 95% confidence level. The plots are based on six months of experimentation covering December 2016 to August 2017.

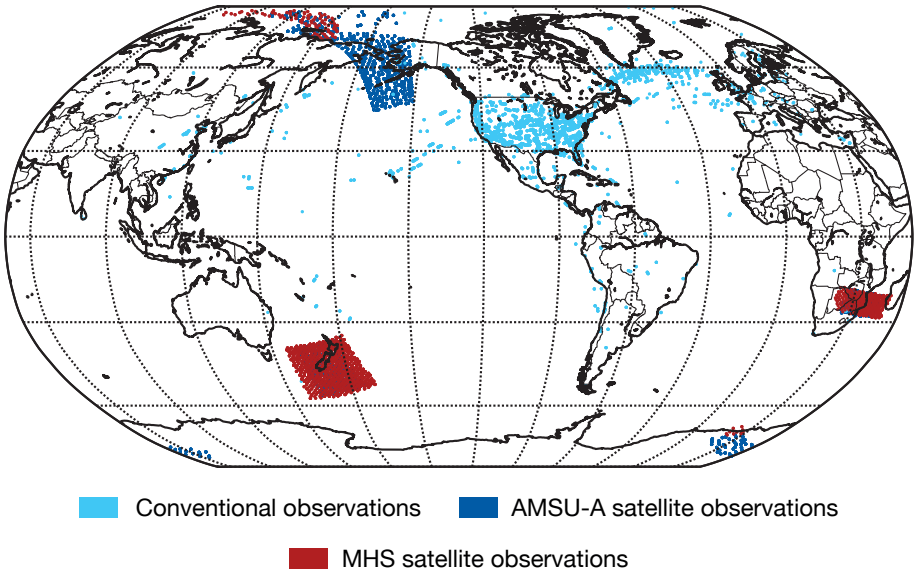

FIGURE 7 Example of the geographic distribution of observations made between 04:00 and 04:25 UTC and assimilated in the continuous DA configuration (with an observation data cut-off at 04:25 UTC).

time will be used in this period. In this example, we see many conventional observations, including some AMDAR reports from aircraft over the North Atlantic, Europe and North America. In addition, data from polar-orbiting satellites as they downlink data over the ground stations in the polar latitudes and at local reception stations of the Direct Broadcast Network (DBNet) can be seen.

It is easy to see how improving the timeliness of all observation types would increase the coverage at the end of the window, further constraining the analysis and leading to increased predictability in the forecasts. The DBNet initiative, coordinated by the World Meteorological Organization, is an example of how this can be achieved in a very cost-effective way (*http://www.wmo.int/pages/prog/sat/rars_en.php*). It uses a network of ground stations that process local satellite overpasses in a globally consistent way and can thus achieve very good timeliness. Expanding such initiatives will bring even greater benefits in the new continuous DA framework. In the continuous DA configuration, we will still be able to use less timely data, but improved timeliness will substantially increase the number of observations that can be used.

Future possibilities

Continuous DA opens the door to exciting new avenues of research and new configurations of our system. It makes it possible for significantly more expensive, but correspondingly more accurate 4D-Var configurations to fit into the operational schedule.

In Cycle 46r1, we take advantage of this by starting the assimilation 10 minutes earlier to accommodate an extra outer loop. However, this could be taken further in upcoming cycles. For example, the possibility of running higher-resolution minimisations or using stricter convergence criteria is currently being explored. Taking this development to its logical extreme, we can envisage the possibility that, one day, our assimilation system could run quasi-continuously with new observations being drip-fed into the system as they arrive, progressively refining our analyses.

Conclusions

The current Global Observing System (GOS) produces a continuous flow of observations that can be used to improve the initial conditions and forecast skill of NWP systems. The pre-processing and initial quality control of these observations is also transitioning at ECMWF towards a near-continuous framework in the context of the Continuous Observation Processing Environment (COPE) project. The continuous data assimilation system described in this article represents a first step at ECMWF to better exploit the steady stream of available observations and, more generally, to adapt the ECMWF data assimilation system to the changing characteristics of the GOS. An important by-product of this development is the fact that it makes more time available for DA computations. This will enable us to introduce more effective 4D-Var configurations in operations while maintaining the same product dissemination schedule and reducing pressure on the time-critical part of the operational suite. Further improvements are possible in this area (e.g. increasing the resolution of the minimisations, more accurate linearisation states, etc.) and are currently being explored.

In the continuous DA framework, we can assimilate observations taken around one and a half hours later than in the current system, as well as delayed observations that would previously have arrived after the cut-off time. In addition, the extension of the assimilation window up to the cut-off time makes the quality of our forecasts more responsive to changes in the timeliness of observation delivery.

Experiments with this first version of continuous DA reflect the benefits of being able to assimilate more observations and producing a more accurate 4D-Var analysis, with predictability gains of two to three hours for typical skill scores.

As important as these forecast improvements are, a possibly more consequential change is the new possibilities that the continuous DA framework offers for the evolution of the ECMWF DA system. Conceptually it is easy, for example, to extend the continuous DA setup towards an assimilation system running continuously 'in the background' and continuously providing updated estimates of the initial conditions of the Earth system based on the steady stream of incoming observations.

Further reading

Bonavita, **M.**, **P. Lean** & **E. Holm**, 2018: Nonlinear effects in 4D-Var, *Nonlinear Processes in Geophysics*, **25**, 713–729.

Courtier, **P.**, **J.-N. Thépaut** & **A. Hollingsworth**, 1994: A strategy for operational implementation of 4D-Var, using an incremental approach, *Q. J. Roy. Meteor. Soc.*, **120**, 1367–1387, https://doi.org/10.1002/qj.49712051912.

Gauthier, **P.**, **M. Tanguay**, **S. Laroche**, **S. Pellerin** & **J. Morneau**, 2007: Extension of 3DVAR to 4DVAR: Implementation of 4DVAR at the Meteorological Service of Canada, *Mon. Wea. Rev.*, **135**, 2339–2354.

Haseler, **J.**, 2004: Early-delivery suite. *ECMWF Technical Memorandum* **No. 454**. https://www.ecmwf.int/en/elibrary/9793-early-delivery-suite

Järvinen, **H.**, **J.-N. Thépaut** & **P. Courtier**, 1996: Quasi-continuous variational data assimilation, *Q. J. Roy. Meteor. Soc.*, **122**, 515–534.

Pires, **C.**, **R. Vautard** & **O. Talagrand**, 1996: On extending the limits of variational assimilation in nonlinear chaotic systems, *Tellus A*, **48**, 96–121.

Rabier, **F.**, **H. Järvinen**, **E. Klinker**, **J.-F. Mahfouf** & **A. Simmons**, 2000: The ECMWF operational implementation of four-dimensional variational assimilation. Part I: Experimental results with simplified physics. *Q. J. Roy. Meteor. Soc.*, **126**, 1143–1170.

A 50-member Ensemble of Data Assimilations

Simon Lang, Elías Hólm, Massimo Bonavita, Yannick Trémolet (JCSDA, US)

Since 2010, ECMWF has run an Ensemble of Data Assimilations (EDA) to help determine the initial conditions for its ensemble forecasts and its higher-resolution deterministic forecast. The EDA is an ensemble of 4D-Var data assimilations that reflects uncertainties in observations; atmospheric boundary conditions, such as sea-surface temperature; and the model physics. The role of the EDA is twofold: it contributes to the high-resolution initial conditions (the high-resolution analysis) by providing flow-dependent estimates of the errors in the short-range forecasts (the background) used in the data assimilation system; and it helps to determine the initial conditions for ensemble forecasts by providing EDA-based perturbations to the high-resolution analysis. In 2013, the number of EDA members was increased from 10 to 25. We have now developed a new, optimised 50-member EDA configuration that has a comparable computational footprint to the current operational 25-member configuration. The increase in ensemble size improves both the high-resolution analysis and the EDA-based perturbations to the initial conditions for the 50-member ensemble forecast. The change is due to be implemented in the next upgrade of ECMWF's Integrated Forecasting System (IFS Cycle 46r1), scheduled for 2019.

Benefits of more EDA members

The EDA uses a Monte Carlo approach to simulate the impact of observation, boundary condition and model uncertainties on the 4D-Var data assimilation system (Isaksen et al., 2010). This makes it possible to quantify uncertainty in the background and in the analysis. Since each EDA member is an independent 4D-Var data assimilation with multiple outer loops, the EDA is capable of sampling non-Gaussian posterior probability density functions. The maintenance cost of the EDA is low, as there is no need to support a separate data assimilation system for the ensemble assimilation component, and it is straightforward to propagate new developments in the high-resolution 4D-Var to the EDA. The EDA scales perfectly with ensemble size and contributes significantly to the skill of ECMWF's forecasts, but its computational cost is high. As a result, the EDA members are run at lower resolution than the high-resolution 4D-Var and the number of EDA members is currently limited to 25. However, it is advantageous to increase the number of EDA members, for two reasons. First, a bigger ensemble size improves the sampling of flow-dependent errors of the background used in ECMWF's 4D-Var system. Improved estimates of background errors in the data assimilation system in turn enable an improved high-resolution analysis. Second, increasing the EDA ensemble size to 50 makes it possible to assign independent EDA-based initial perturbations to each of the 50 ensemble forecast members. The resulting exchangeability of ensemble members is particularly important for research experiments involving ensemble forecasts.

Figure 1 illustrates the impact of the new 50-member EDA on the uncertainty estimates used in ECMWF's 4D-Var system. Due to better sampling, the ensemble standard deviation of the 50-member EDA background forecasts is spatially smoother than the ensemble standard deviation of the 25-member EDA background

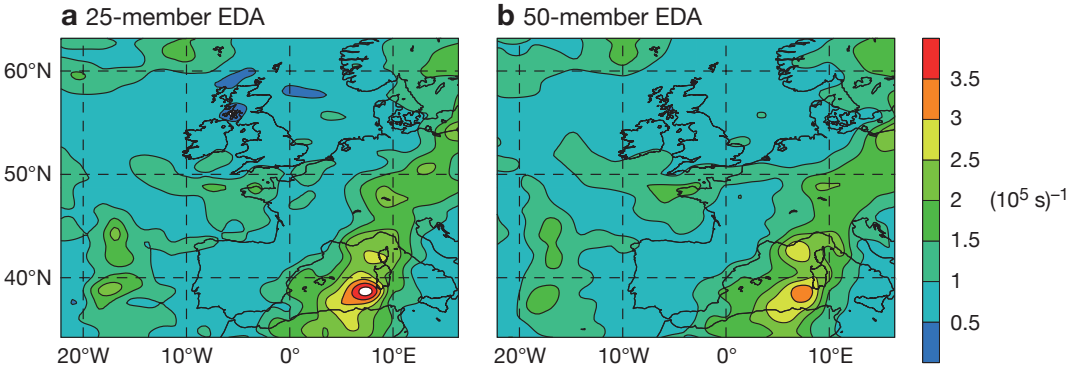

FIGURE 1 EDA 3-hour background forecast ensemble standard deviation of relative vorticity at model level 100 starting at 18 UTC on 31 August 2017 for (a) the 25-member EDA and (b) the 50-member EDA.

doi: 10.21957/nb251xc4sl

forecasts. Going to 50 EDA members reduces the sampling noise. As a result, the number of points with very small and very large standard deviations is reduced. This in turn improves the use of observations, as observations in regions with unrealistically small (unrealistically large) standard deviations are given too little (too much) weight in the 4D-Var system.

Computational efficiency

The computational cost per individual member in the 50-member configuration has been reduced so that the overall cost is the same as for the current 25-member configuration. This has been achieved by extensive optimisation, for example in the preconditioning of the first 4D-Var minimisation of the perturbed members with information from the minimisation of the unperturbed control member; the blacklisting of passive observations in the perturbed members; and improved observation thinning and optimisation of the observation operator setup for lower-resolution inner loops. In addition, non-essential tasks have been deactivated. For example, by default we now run a reduced number of final trajectories, which are merely used for diagnostic purposes, and the surface analysis is run with an optimal interpolation scheme instead of the operational simplified extended Kalman filter of the high-resolution analysis. Extensive tests have shown that the optimisation does not negatively impact the skill of individual EDA members or the EDA's uncertainty estimates.

A welcome side effect of the optimisation work is that the observation-related changes will also reduce the computational cost of all standard deterministic, low-resolution research experimentation and hence make it possible to achieve a higher throughput for research experiments.

Exchangeability

The EDA members are used together with singular vectors to generate perturbations which are added to the high-resolution analysis to construct the perturbed initial conditions for ensemble forecasts. Currently the initial perturbations of the ensemble forecasts have a plus–minus symmetry to distribute the 25 EDA members between the 50 ensemble forecast members: the EDA perturbation generated from the first EDA member is assigned to the first ensemble forecast member, the first EDA perturbation with the sign reversed is assigned to the second ensemble forecast member, and so on. A disadvantage of this scheme is that the ensemble forecast members are not exchangeable. Exchangeability is desirable for research ensemble forecast testing with small ensemble sizes (Leutbecher, 2018) to estimate the skill of the full 50-member ensemble. By increasing the number of EDA members to match the number of ensemble forecast members, it is possible to assign one individual EDA perturbation to each ensemble forecast member. As a result, the plus–minus symmetry of the initial perturbations is no longer required and the ensemble forecast members become exchangeable.

Figure 2 shows the mean absolute difference between pairs of members in an ensemble forecast experiment. In the ensemble forecast experiment that has the plus–minus symmetry of the initial perturbations from 25 EDA members, there is a distinct difference between pairs of members that share the same initial perturbation with only the sign reversed on the one hand, and pairs of members with independent initial perturbations on the other. In contrast, the members of the ensemble forecast experiment with perturbations from 50 EDA members and without the plus–minus symmetry are indistinguishable.

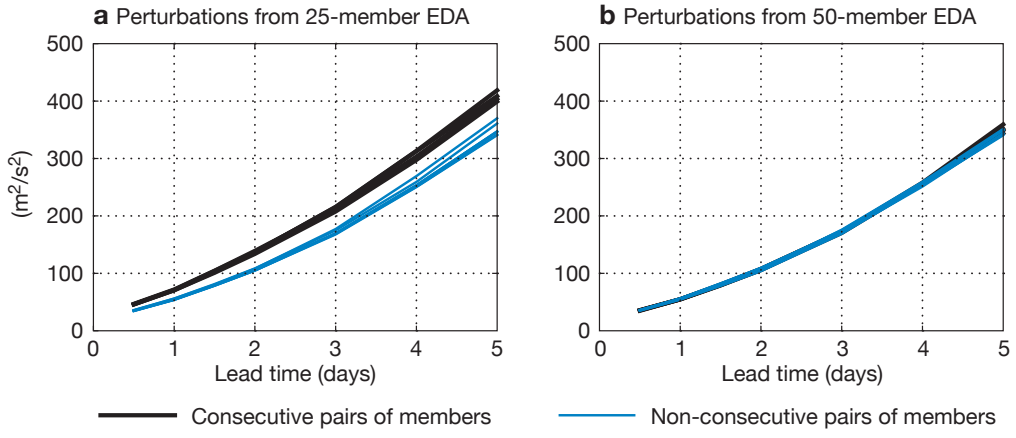

FIGURE 2 Mean absolute difference of 500 hPa geopotential in the northern extratropics between consecutive pairs of ensemble forecast members and non-consecutive pairs for (a) an ensemble forecast experiment with plus–minus symmetry of the initial perturbations (based on a 25-member EDA) and (b) an ensemble forecast experiment without such symmetry (based on a 50-member EDA). The differences are averaged over a period of 41 days.

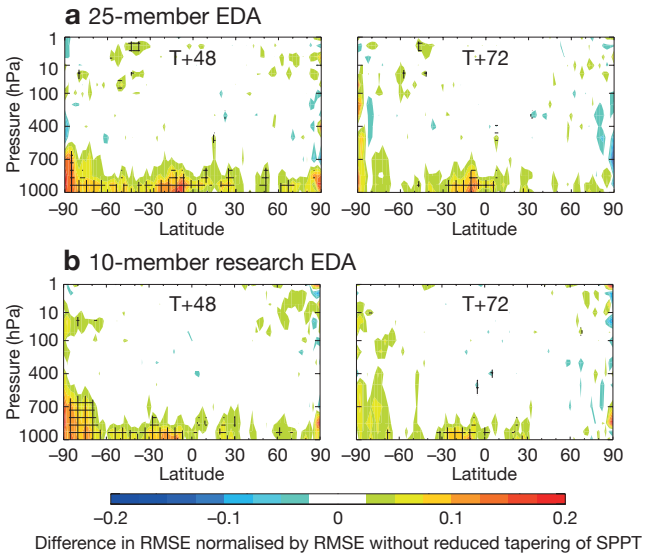

FIGURE 3 Impact of reduced tapering of SPPT (stochastic physics) in the boundary layer on the root-mean-square error (RMSE) in temperature for forecasts at different lead times based on the EDA control members in (a) the current EDA configuration with 25 members and (b) the proposed research test configuration with 10 members and lower-resolution outer loops and lower-resolution inner loops for the perturbed EDA members. Green to red shading indicates larger RMSE with reduced tapering of SPPT.

Research configuration

In addition to the new 50-member configuration, we have introduced computationally inexpensive research and testing EDA configurations that run with only 10 members. They will reduce the cost of EDA experimentation and streamline the testing and validation process for new model cycles. Specifically, from the next IFS upgrade, the operational EDA will have 50 members with outer loops at TCo639 resolution and two inner loops at TL191 resolution. The research configuration, on the other hand, has only 10 members, TCo399 outer loops and TL95/TL159 inner loops. Only the control member is produced using the same 4D-Var configuration as the standard low-resolution (TCo399) deterministic research experiments used for model development. This makes it considerably cheaper to assess the impact of model changes on the EDA and the feedback this has on the analysis. The cost of the EDA research configuration is only approximately three times that of a standard deterministic research experiment. The cost of equivalent experimentation with the currently operational 25-member EDA is over 20 times higher than a deterministic research experiment.

For research and development, it is desirable to run experiments at lower resolution to reduce computational cost. However, it is important for conclusions derived from research experiments to be relevant to the full-resolution operational system. For example, a new development that improves forecast quality in a research experiment compared to a reference experiment should also improve the high-resolution operational system. In order to test if this is indeed the case for the research EDA configuration, we look at a change to the Stochastically Perturbed Parametrization Tendencies (SPPT) scheme, which is used in the EDA to simulate model uncertainties. The focus here is not a decrease or increase in forecast skill but whether the full operational EDA configuration and the research EDA configuration react to the change in the same way. Figure 3 shows the impact on the skill of forecasts based on the unperturbed EDA control member resulting from a change to the SPPT. The change influences the EDA ensemble standard deviation and thus the uncertainty estimates used in the 4D-Var system. Consequently, it impacts the forecast skill of each EDA member, including the unperturbed EDA control member. The plot shows the results obtained in two sets of forecast skill experiments, one run with the currently operational 25-member EDA configuration – which is equivalent in cost to the new optimised 50-member configuration – and one run with the new 10-member research configuration. For each configuration, the skill of the unperturbed control forecasts is compared to a reference experiment which is run without the change to the stochastic physics scheme in the EDA. The difference between the experiment with the stochastic physics change and the reference experiment is very similar regardless of whether the 25-member configuration or the 10-member research configuration is used. This shows that it is possible to derive robust results from the research EDA configuration for a fraction of the cost of running the full operational configuration. However, it is essential to compare like with like and not to mix the different configurations.

Outlook

The current method to initialise ensemble forecasts is to apply EDA-based perturbations and singular vectors to the high-resolution deterministic analysis. That analysis is more accurate than the EDA analyses thanks to its higher resolution and the fact that it contains six extra hours of observations compared with the EDA. The long-term goal, however, is to develop a high-resolution EDA which runs in parallel with the high-resolution 4D-Var and includes the same observations, and which can provide the initial conditions for ensemble forecasts directly. In addition to the benefits described in this article, the development of a computationally efficient 50-member EDA is a step towards that goal.

Further reading

Isaksen, **L.**, **M. Bonavita**, **R. Buizza**, **M. Fisher**, **J. Haseler**, **M. Leutbecher** & **L. Raynaud**, 2010: Ensemble of Data Assimilations at ECMWF, *ECMWF Technical Memorandum* **No. 636**.

Leutbecher, **M.**, 2018: How many ensemble members are desirable? *ECMWF Newsletter* **No. 157**, 5.

A nonhydrostatic finite-volume option for the IFS

Christian Kühnlein, Piotr K. Smolarkiewicz

Over the next decade, many aspects of ECMWF's Integrated Forecasting System (IFS) may need to change in light of the efficiency of higher-resolution global forecasts called for by ECMWF's long-term strategy. The dynamical core lies at the heart of the model infrastructure. It numerically solves the fundamental governing equations, which take the form of physical conservation laws describing the resolved atmospheric dynamics. The dynamical core is coupled to parametrizations of subgrid-scale physical atmospheric processes and to models of other Earth system components. The current IFS dynamical core depends on the spectral-transform method to solve the governing equations. ECMWF is continuing to develop this dynamical core, which also includes a nonhydrostatic option, to make it as computationally efficient as possible. For added flexibility, it is also developing a new, nonhydrostatic dynamical core which uses the finite-volume method. This 'Finite-Volume Module' (FVM) has been shown to perform well compared to the current dynamical core in benchmark tests, and it holds the promise of greater computational efficiency for global nonhydrostatic forecasts at very high resolution run on future exascale high-performance computing (HPC) facilities.

The current dynamical core

The dynamical core of the operational IFS solves the hydrostatic primitive equations (HPEs) using the spectral-transform (ST) technique. It will be referred to as the IFS-ST in this article. The ST technique was introduced in the ECMWF forecasting model in the 1980s. It is still successfully employed today in combination with a semi-implicit semi-Lagrangian (SISL) integration scheme (see Wedi et al., 2015, for details and a comprehensive list of references). Recent advances, such as the cubic truncation and the Fast Legendre Transform, will ensure the computational efficiency of the ST method in the next decade. Nevertheless, there are uncertainties regarding the long-term viability of that method. These uncertainties are rooted in the scalability of the high-volume non-local data communications associated with it. At some point, this could lead to an excessive amount of time spent on parallel communication, preventing the timely delivery of the forecast (Wedi et al., 2015). In response to these issues, ECMWF is widening its options with respect to the dynamical core.

On the one hand, ECMWF is improving the computational efficiency of the ST method and of the nonhydrostatic extension of the IFS-ST that has been kindly made available by Météo-France and the ALADIN consortium. On the other hand, ECMWF is pursuing the development of methods with fundamentally different parallel communication patterns and a complementary nonhydrostatic dynamical core design, such as in the FVM.

The Finite-Volume Module

The Finite-Volume Module of the IFS (henceforth IFS-FVM) is currently under development at ECMWF (Smolarkiewicz et al., 2016, Kühnlein et al., 2018). An important property of the finite-volume (FV) method applied in the IFS-FVM is that solutions to the governing equations are calculated at discrete places on a meshed geometry (Box A). This means that there is a distributed-memory communication footprint that is predominantly local and performed via thin 'halo' cells shared with the nearest neighbours. This contrasts with the high-volume non-local communications required in the IFS-ST. In addition to the different communication patterns, the IFS-FVM provides numerous capabilities that are currently unavailable in the IFS. Among the advantages of the FV method are inherently local conservation and the ability to operate in complex mesh geometries (Kühnlein & Smolarkiewicz, 2017). The lack of conservation is a common issue with standard semi-Lagrangian schemes and a shortcoming in the operational IFS, which is suspected to contribute to moisture and temperature biases and to affect the forecast quality, particularly at sub-seasonal to seasonal forecast ranges. The ability of FV methods to operate in complex mesh geometries makes it possible to implement efficient quasi-uniform-resolution meshes that circumvent the polar anisotropy of the classical regular longitude-latitude grids. Furthermore, it enables variable and/or adaptive resolution with locally finer mesh spacings in sensitive regions (e.g. storm-tracks) and hence a more accurate and efficient representation of multi-scale interactions.

doi: 10.21957/sd92ack6p3

a Finite-volume method

The finite-volume method is an approach to the approximate integration of partial differential equations (PDEs) describing natural conservation laws. The IFS-FVM uses the 'median-dual' finite-volume method. 'Finite volume' refers to the small volume (grey lines in the figure) surrounding each node point (black dots) on a mesh where solutions are computed. When considering the PDEs, integrals of the divergence and gradient terms over these finite volumes are converted into surface integrals using the Gauss divergence theorem. On a discrete mesh, these surface integrals are then evaluated as a sum of all fluxes through individual surfaces bounding each finite volume. The fluxes are evaluated at the centre of the edges (black lines) connecting the nodes.

Since the flux entering a given volume is identical to that leaving the adjacent volume, these methods are conservative. Another important advantage of the finite-volume method is that it can be formulated for complex meshes.

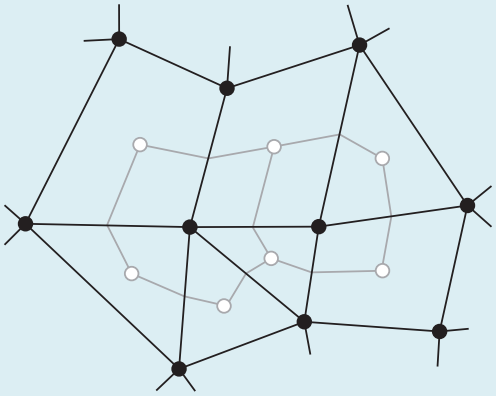

A novel atmospheric dynamical core formulation for NWP represents a long-term development. In addition to the dynamical core itself, it involves various other aspects of the model infrastructure. The design of the IFS-FVM facilitates its incorporation into the existing IFS. Table 1, adapted from Kühnlein et al. (2018), lists the main features of the IFS-FVM on the one hand and the hydrostatic and nonhydrostatic IFS-ST dynamical cores on the other. Despite fundamental differences in the governing equations and the discretisation, the different dynamical cores can share certain features. The most prominent examples are the octahedral reduced Gaussian grid (hereafter referred to as the 'octahedral grid') of the operational IFS and the co-located arrangement of all prognostic variables at the same nodes of the grid. This means the FV mesh is developed about the nodes of the same octahedral grid that supports the IFS-ST (Figure 1). The FV mesh is defined in terms of edges connecting the nodes and dual volumes around the nodes (Box A). The use of the same octahedral grid for the different dynamical cores has numerous benefits for the general model infrastructure, data assimilation and model initialisation, as well as comparison studies. In the IFS-FVM, mesh generation, mesh data structures and nearest-neighbour distributed-memory communication are provided by ECMWF's Atlas framework (Deconinck et al., 2017), which will also provide support for heterogeneous HPC architectures.

To make the IFS-FVM a useful option for global medium-range weather forecasting, it needs to be shown that it can provide forecast quality that is at least comparable to that of the current IFS. Another important question concerns the computational efficiency of the IFS-FVM, which must be sufficient to meet the tight constraints of operational scheduling. In the following, we study the questions of accuracy and computational

Model aspect	IFS-FVM	IFS-ST (hydrostatic)	IFS-ST (nonhydrostatic option)
Equation system	fully compressible	hydrostatic primitive	fully compressible
Prognostic variables	$\rho_d, u, v, w, \theta', \varphi', r_v, r_l, r_r, r_i, r_s$	$\ln p_s, u, v, T_v, q_v, q_l, q_r, q_i, q_s$	$\ln \pi_s, u, v, d_4, T_v, \hat{q}, q_v, q_l, q_r, q_i, q_s$
Horizontal coordinates	λ, ϕ (lon-lat)	λ, ϕ (lon-lat)	λ, ϕ (lon-lat)
Vertical coordinate	generalised height	hybrid sigma-pressure	hybrid sigma-pressure
Horizontal discretisation	unstructured finite-volume (FV)	spectral-transform (ST)	spectral-transform (ST)
Vertical discretisation	structured FD/FV	structured FE	structured FD/FE
Horizontal staggering	co-located	co-located	co-located
Vertical staggering	co-located	co-located	co-located/Lorenz
Horizontal grid	octahedral Gaussian/arbitrary	octahedral Gaussian	octahedral Gaussian
Time-stepping scheme	2-TL SI	2-TL constant-coefficient SI	2-TL constant-coefficient SI with ICI
Advection	conservative FV Eulerian	non-conservative SL	non-conservative SL

TABLE 1 Summary of the main features of the IFS-FVM and the IFS-ST (hydrostatic and nonhydrystatic). The abbreviations stand for finite-difference (FD), finite-element (FE), spectral-transform (ST), finite-volume (FV), two-time-level (2-TL), semi-implicit (SI) and iterative-centred-implicit (ICI). See Kühnlein et al. (2018) for further details.

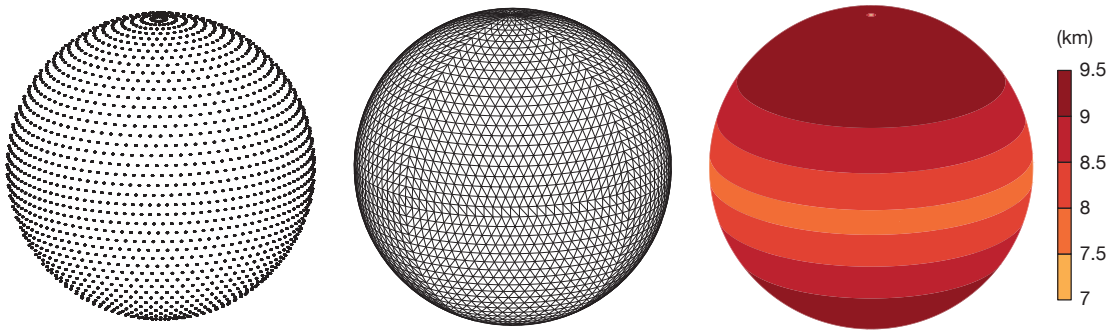

FIGURE 1 The octahedral grid used in the IFS-FVM and the IFS-ST, showing the locations of grid nodes using the O24 example, with 24 latitude lines between the pole and the equator (left), edges connecting the nodes for the O24 grid (middle) and the local spacing of the dual mesh cells for the O1280 grid used in ECMWF's current highest-resolution deterministic forecasts (right).

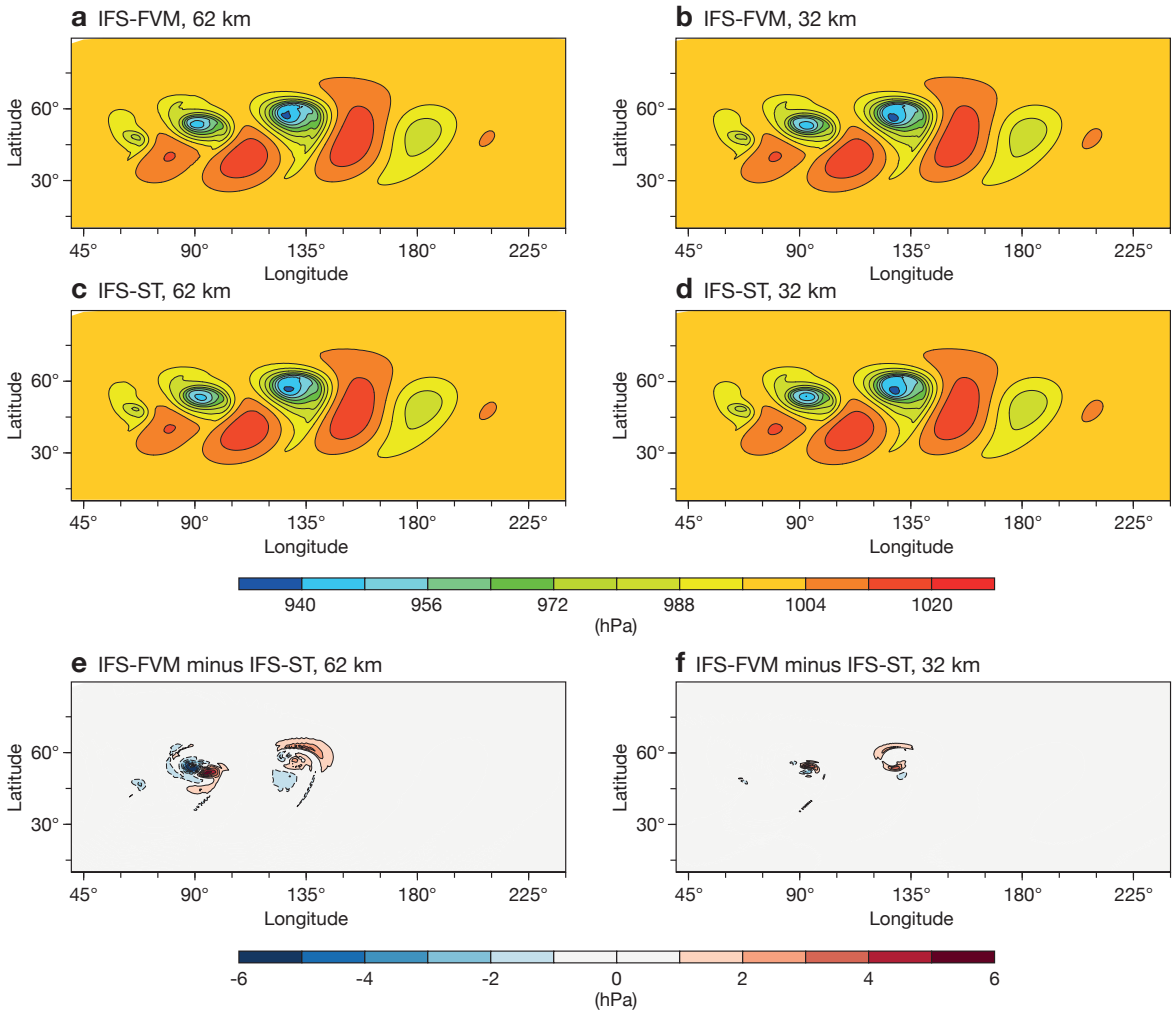

FIGURE 2 Baroclinic instability at forecast day 10, showing pressure at the lowest full level obtained with (a) the IFS-FVM with the O160 horizontal grid, (b) the IFS-FVM with the O320 grid, (c) the IFS-ST with the TCo159 grid, (d) the IFS-ST with the TCo319 grid, (e) the difference between the IFS-FVM and the IFS-ST with the O160/TCo159 grid and (f) the difference between the IFS-FVM and the IFS-ST with the O320/TCo319 grid.

efficiency, considering test cases of intermediate complexity at the current stage of development. More details about the IFS-FVM model formulation used and the comparison to the IFS-ST can be found in Kühnlein et al. (2018).

Comparison to spectral-transform IFS

Basic evaluation of the IFS-FVM dynamical core has been performed using the baroclinic instability benchmark (see Kühnlein et al., 2018, for a description of the setup). This is a commonly used test to evaluate

the performance of NWP models in the large-scale hydrostatic regime. Here, the proven hydrostatic IFS-ST dynamical core that is used for operational forecasting at ECMWF will serve as the reference.

Figures 2 and 3 show solutions for the two dynamical cores using the same octahedral grid and the difference between them. Results are shown for an example forecast at day 10, when a large-amplitude baroclinic wave has developed and has formed sharp fronts in the lower troposphere. The difference plots in Figures 2e,f and 3e,f show that there is very close agreement between the IFS-FVM and the IFS-ST. The solutions show essentially similar phase propagation and amplitude of the baroclinic wave throughout the whole depth of the simulation domain. There are only small differences at a grid spacing of about 62 km (O160 for the IFS-FVM and TCo159 for the IFS-ST) and they become even smaller when reducing the grid spacing to about 32 km (O320 and TCo319, respectively). Overall, the results attest to the high quality of the IFS-FVM dynamical core.

Computational efficiency

Computational efficiency of NWP models is crucial. For current HPC architectures at operational resolutions, the IFS-ST at ECMWF represents one of the most efficient dynamical core formulations for global NWP. The IFS-FVM is envisaged for future applications in the nonhydrostatic regime running on future HPC architectures, but its computational performance on the current HPC facility at ECMWF sheds light on its potential. We have assessed the computational efficiency of the dynamical cores using the dry baroclinic instability benchmark presented in Figures 2 and 3.

The IFS-FVM is a recent development, and considerable progress has been made in terms of its computational efficiency. Over the last two years, the IFS-FVM has been speeded up by at least a factor of seven. This has been achieved by optimising various aspects of the time stepping and by extensive recoding of the algorithms. For details, see Kühnlein & Smolarkiewicz (2017) and Kühnlein et al. (2018).

Figure 4 shows runtimes of the nonhydrostatic IFS-FVM against the hydrostatic IFS-ST, both configured similarly to ECMWF's current high-resolution forecast (HRES). The time-to-solution that can currently be achieved with the IFS-FVM is only about twice as long as for the

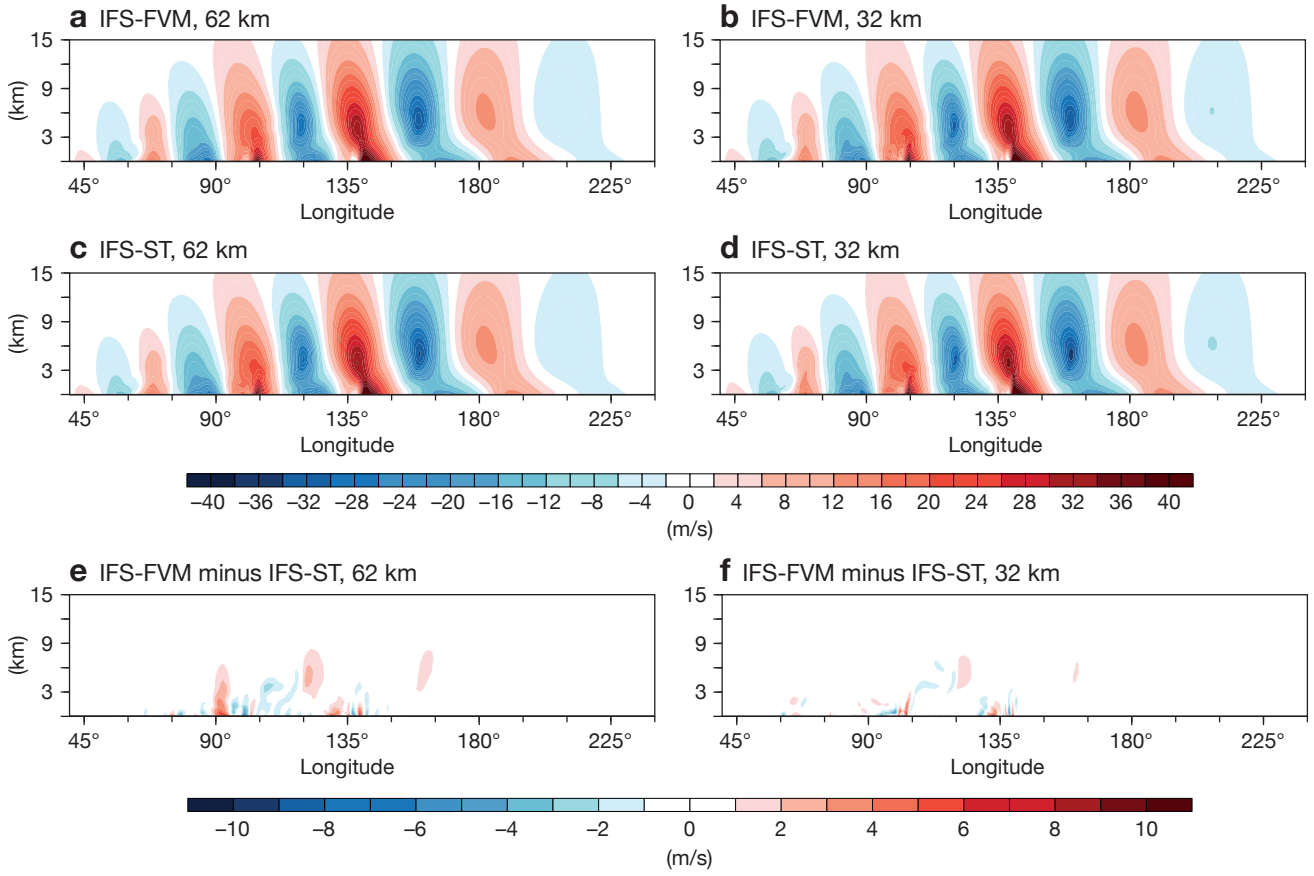

FIGURE 3 Baroclinic instability at forecast day 10, showing meridional wind along an east–west–height cross section at 50°N obtained with (a) the IFS-FVM with the O160 horizontal grid, (b) the IFS-FVM with the O320 grid, (c) the IFS-ST with the TCo159 grid, (d) the IFS-ST with the TCo319 grid, (e) the difference between the IFS-FVM and the IFS-ST with the O160/TCo159 grid and (f) the difference between the IFS-FVM and the IFS-ST with the O320/TCo319 grid.

operational hydrostatic IFS-ST. Indeed, the IFS-FVM already compares favourably with other nonhydrostatic dynamical cores. Moreover, the local FV method offers the prospect of better scalability and efficiency with respect to future HPC. Another important aspect is that the IFS-FVM employs smaller time steps (typically smaller by a factor of 5 to 6 compared to the IFS-ST), which can be beneficial for accuracy. Overall, the results highlight the potential of the IFS-FVM to become competitive for operational global weather forecasting.

Moist dynamics and coupling to parametrizations

The results presented above were obtained using the basic dry dynamical cores. At the current stage of development, the IFS-FVM has been extended to moist/precipitating dynamics and coupling to selected IFS physical parametrizations (Kühnlein et al., 2018). Work on coupling to the full IFS physical parametrization package is ongoing. Figures 5 and 6 compare the IFS-FVM to the reference IFS-ST using the baroclinic instability benchmark with moist/precipitating processes. For this particular test, coupling to the operational prognostic single-moment bulk microphysics parametrization of the IFS for large-scale condensation and precipitation is applied, implemented by means of a generic interface suitable for all parametrizations. The current generic interface calls and evaluates the physical parametrizations sequentially at each of the IFS-FVM semi-implicit time steps, which is the standard procedure in the operational IFS-ST. Because of the significant cost of the IFS physical parametrizations (about 30% of the forecast model), the physics–dynamics coupling with the IFS-FVM needs to be revised in order to maintain computational efficiency, given the smaller time steps of the IFS-FVM.

Figure 5 shows an example surface precipitation field at forecast day 10. Five rain bands which are virtually identical for the two dynamical cores appear at the two chosen horizontal grid spacings. The rain bands, which are the result of air being lifted along sharp frontal zones, are essentially identical in phase. This is highlighted in Figure 5 by overlaying the IFS-FVM precipitation rate shading with IFS-ST 0.5 mm/day contours and vice versa. In addition, precipitation rates are overall similar. Figure 6 shows a time series of the minimum near-surface pressure and the area-integrated precipitation rate over the course of a 15-day simulation. It indicates close agreement in terms of the magnitude and temporal evolution of these quantities, and no systematic biases. Again, the presented results attest to the high quality of the IFS-FVM dynamical core with moist/precipitating dynamics and the coupling interface to physical parametrizations.

The ability of next-generation dynamical cores to produce global forecasts in the nonhydrostatic regime is

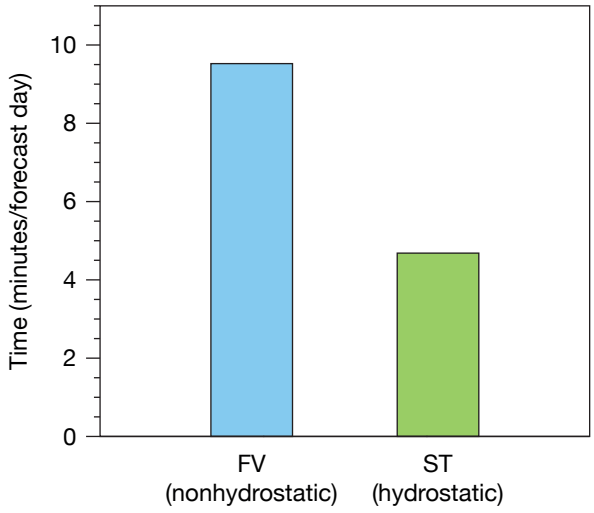

FIGURE 4 Elapsed time to run one day of the dry baroclinic instability benchmark test similar to the current high-resolution forecast (HRES) configuration, set up for the O1280/TCo1279 horizontal grid with 137 vertical levels, and using 350 nodes of ECMWF's Cray XC40 supercomputer.

essential. The IFS-FVM participates in the current round of the Dynamical Core Model Intercomparison Project (DCMIP), in which next-generation nonhydrostatic dynamical cores from various weather forecasting centres and research institutions are compared using common test cases. Figure 7 provides an illustration from a 'splitting supercell' test case, run in a small planet configuration at about 1 km horizontal grid spacing. The grid spacing approaches the horizontal extent of convective updrafts so that the experiment is a test case for resolved nonhydrostatic dynamics. The dynamical core solutions feature different responses in small scales depending on the numerical discretisation and filtering mechanism. The properties of different numerical formulations for global atmospheric models in the regime where convection is (partially) resolved will be studied more systematically over the next few years.

Outlook

Supporting substantially higher resolution in global NWP may ultimately demand local numerical discretisations to solve the governing nonhydrostatic equations in a computationally efficient manner. The IFS-FVM successfully implements such a local discretisation and thus complements the operational hydrostatic IFS-ST at ECMWF. At the same time, the IFS-FVM introduces several useful new features into the IFS, such as conservative and monotone advective transport, deep-atmosphere nonhydrostatic governing equations, and fully flexible unstructured FV meshes with optional variable resolution or meshes defined about the nodes of the operational octahedral grid. Furthermore, the recently extended perturbation form of the IFS-FVM equations offers

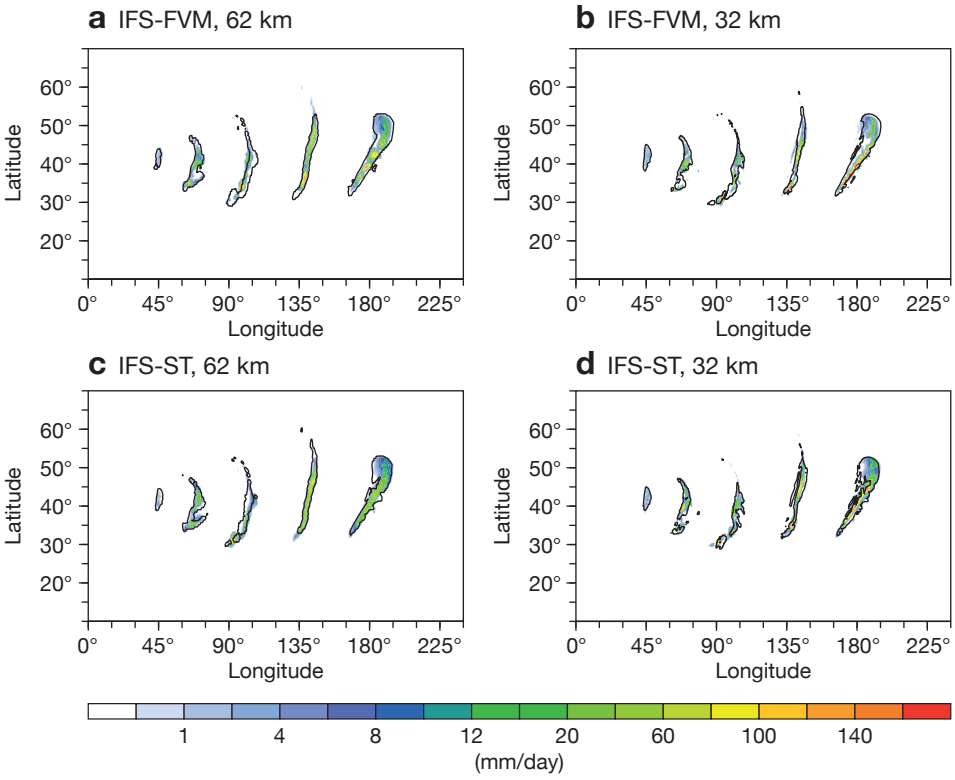

FIGURE 5 Baroclinic instability with moist/precipitating processes at forecast day 10, showing the surface precipitation rate (shading) with the IFS-FVM and the IFS-ST coupled to the same IFS cloud scheme for (a) the IFS-FVM with the O160 horizontal grid, (b) the IFS-FVM with the O320 grid, (c) the IFS-ST with the TCo159 grid and (d) the IFS-ST with the TCo319 grid. For comparison, the IFS-FVM precipitation rate shading has been overlaid with IFS-ST 0.5 mm/day contours and vice versa.

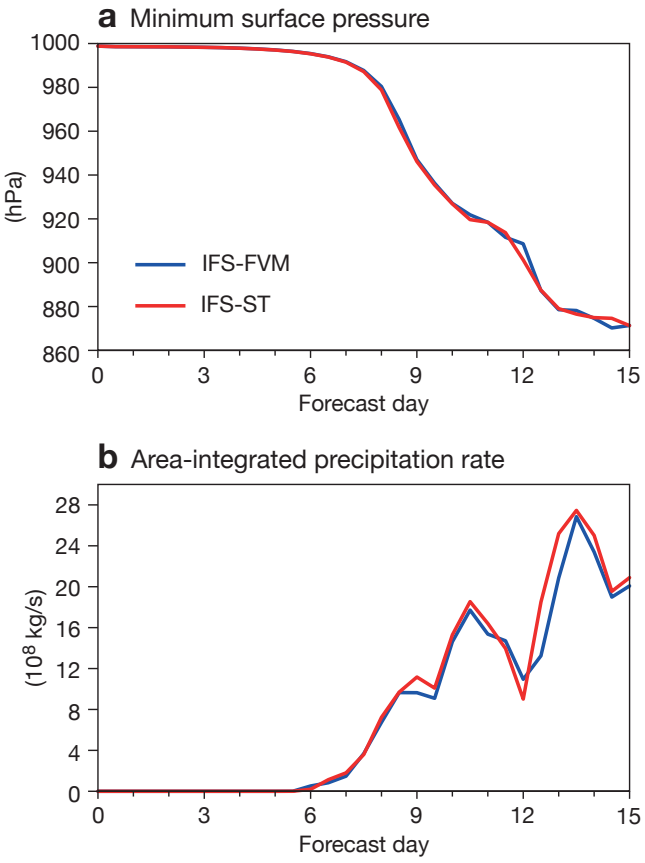

FIGURE 6 Baroclinic instability with moist/precipitating processes, showing (a) the evolution of minimum pressure at the lowest full level by forecast day and (b) the evolution of the domain-averaged precipitation rate by forecast day for the IFS-FVM and the IFS-ST with a 32 km grid spacing (O320 and TCo319, respectively).

significant potential for long-range predictions (Smolarkiewicz et al., 2019).

As shown in this article, the IFS-FVM has advanced to a stage where it can provide forecast quality comparable to that of the operational ST dynamical core for benchmark tests of intermediate complexity. Substantial progress has been made in terms of computational efficiency of the IFS-FVM dynamical core, and further improvements are in preparation.

Ongoing and future work on the IFS-FVM involves all the technical and scientific aspects of coupling it to the full IFS physical parametrization package and configuring the overall model to run global medium-range re-forecast experiments for comparison with the operational IFS-ST. Further research and development will extend the IFS-FVM numerical schemes; explore multigrid techniques; explore alternative physics–dynamics coupling for the smaller time steps; introduce variable precision; and adopt I/O procedures of the IFS. Tangent-linear and adjoint versions of the nonlinear IFS-FVM model required in the context of 4D-Var data assimilation will be developed.

Adaptation of the IFS-FVM to novel low-energy heterogeneous HPC architectures will be facilitated by the data structures of Atlas, which makes use of specific programming paradigms. This is considered to be a basis from which the IFS-FVM can reach out to extreme scalability and into the realm of global nonhydrostatic convection-resolving weather forecasting.

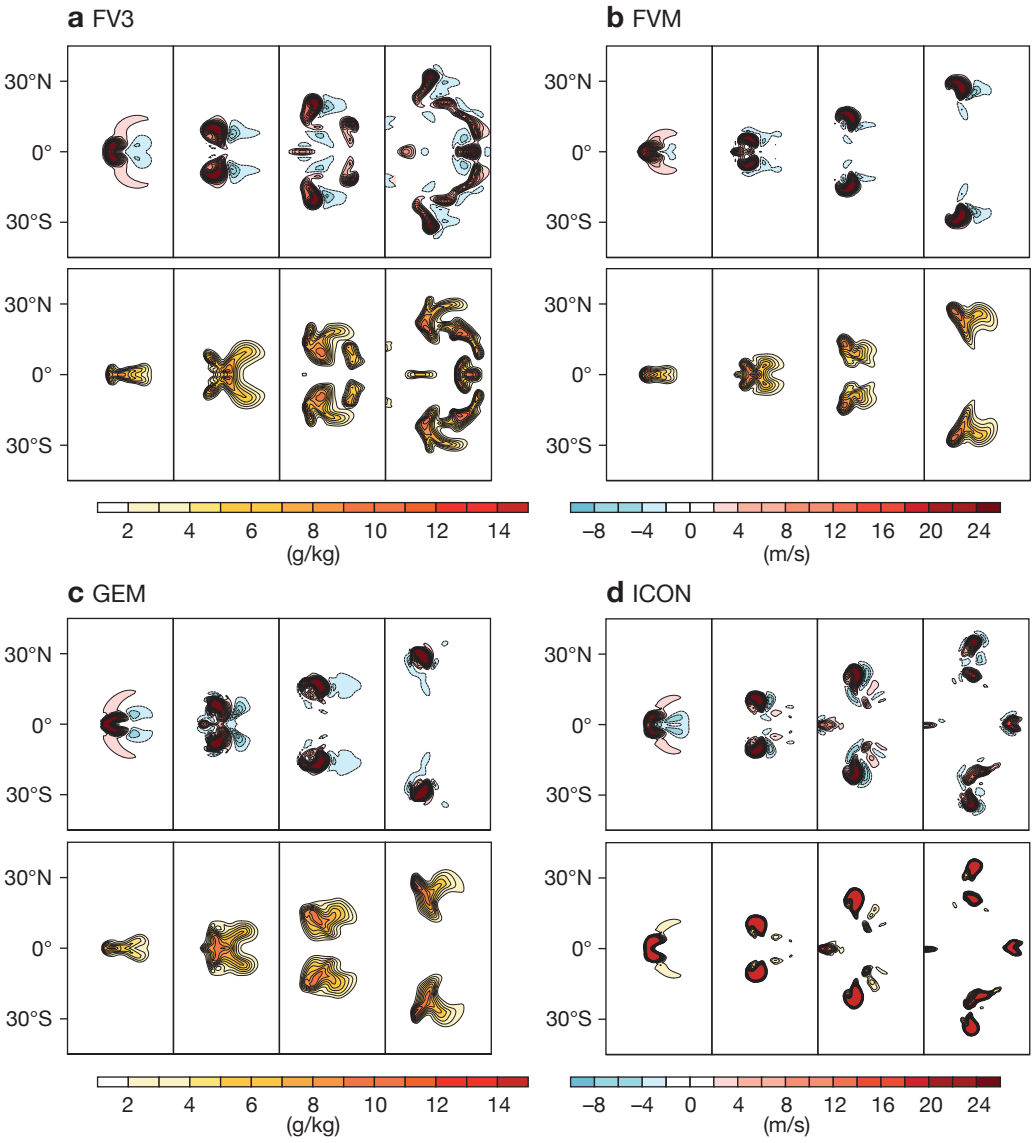

FIGURE 7 Evolution of a splitting supercell over two hours of simulation (0.5, 1, 1.5, 2 h) showing vertical velocity (top half of each panel) and rain water (bottom half of each panel) in a horizontal cross section at 5 km above the surface for four models: (a) FV3 (US Geophysical Fluid Dynamics Laboratory), (b) FVM (ECMWF), (c) GEM (Environment Canada) and (d) ICON (German national meteorological service, DWD/Max-Planck Institute for Meteorology). These are four of the ten participating dynamical cores in the DCMIP study of this test case (plots from an article under review by C. M. Zarzycki et al., 2018, doi:10.5194/gmd-2018-156, used under the CC BY 4.0 licence).

Figures 1, 2, 3, 5 and 6 reproduced from Kühnlein et al. (2018) under the CC BY 4.0 licence.

Further reading

Deconinck, **W.**, **P. Bauer**, **M. Diamantakis**, **M. Hamrud**, **C. Kühnlein**, **P. Maciel**, **G. Mengaldo**, **T. Quintino**, **B. Raoult**, **P.K. Smolarkiewicz** & **N.P. Wedi**, 2017: Atlas: A library for numerical weather prediction and climate modelling. *Comput. Phys. Commun.*, **220**, 188–204.

Kühnlein, **C.** & **P.K. Smolarkiewicz**, 2017: An unstructured-mesh finite-volume MPDATA for compressible atmospheric dynamics. *J. Comput. Phys.*, **334**, 16–30.

Kühnlein, **C.**, **W. Deconinck**, **R. Klein**, **S. Malardel**, **Z.P. Piotrowski**, **P.K. Smolarkiewicz**, **J. Szmelter** & **N.P. Wedi**, 2018: FVM 1.0: A nonhydrostatic finite-volume dynamical core formulation for IFS. *Geosci. Model. Dev. Discuss.*, doi:10.5194/gmd-2018-237, accepted.

Smolarkiewicz, **P.K.**, **W. Deconinck**, **M. Hamrud**, **C. Kühnlein**, **G. Mozdzinski**, **J. Szmelter** & **N.P. Wedi**, 2016: A finite-volume module for simulating all-scale global atmospheric flows. *J. Comput. Phys.*, **314**, 287–304.

Smolarkiewicz, **P.K.**, **C. Kühnlein** & **N.P. Wedi**, 2019: Semi-implicit integrations of perturbation equations for all-scale atmospheric dynamics. *J. Comput. Phys.*, **376**, 145–159.

Wedi, **N.P.**, **P. Bauer**, **W. Deconinck**, **M. Diamantakis**, **M. Hamrud**, **C. Kühnlein**, **S. Malardel**, **K. Mogensen**, **G. Mozdzinski** & **P.K. Smolarkiewicz**, 2015: The modelling infrastructure of the Integrated Forecasting System: Recent advances and future challenges. *ECMWF Technical Memorandum*, **No. 760**.

The new CAMS global reanalysis of atmospheric composition

Antje Inness, Richard Engelen, Johannes Flemming

ECMWF has a long history of providing global meteorological reanalyses, the latest of which is ERA5. Meteorological reanalyses are datasets providing a complete and consistent record of meteorological conditions for recent decades. They are produced by combining model information with observations through data assimilation. This ensures that the resulting gridded datasets are comprehensive and consistent over time. Over the last decade, reanalysis activities have been extended to include other components of the Earth system, such as the land surface, the ocean and the chemical composition of the atmosphere. This is in line with the emphasis in ECMWF's current ten-year Strategy on the need to account for all relevant interactions between different components of the Earth system in ECMWF's Integrated Forecasting System (IFS). In the framework of the Copernicus Atmosphere Monitoring Service (CAMS), implemented by ECMWF on behalf of the EU, ECMWF released a new reanalysis of atmospheric composition in September 2018. This 'CAMS reanalysis' (CAMSRA) covers the period from 2003 to 2016 and will be extended to subsequent years by adding one year each year. It provides consistent information on aerosols and reactive gases, such as ozone (O_3), carbon monoxide (CO), nitrogen dioxide (NO_2) and many more chemical species, using a fully integrated atmospheric composition modelling and data assimilation system based on the IFS.

Comparison with previous reanalyses

CAMSRA follows in the footsteps of the earlier GEMS, MACC (MACCRA) and CAMS interim (CIRA) reanalyses (see Table 1). It has a greater horizontal resolution (T255 or about 80 km) than CIRA (T159 or about 125 km). It also provides data on more chemical species at a better temporal resolution than the previous reanalyses. Great care was taken to ensure that the emission datasets used in CAMSRA were consistent in time and that consistent anthropogenic, biogenic and biomass burning emissions were used in the aerosol and chemistry schemes. Furthermore, a more recent, improved IFS model cycle was used to produce the CAMS reanalysis, and more and newly reprocessed satellite datasets were assimilated. These include satellite retrievals of CO, NO_2, O_3 and aerosol optical depth (AOD). A step towards the coupling between composition and weather is that in CAMSRA prognostic ozone and aerosol fields from CAMSRA are used in the IFS radiation scheme, while the previous atmospheric composition reanalyses used ozone and aerosol climatologies in the radiation scheme.

As shown in more detail in the section on evaluation results, the CAMS reanalysis performs better than the

Period	Name	EXP	Class	IFS Cycle	Resolution	Model	Production Period
01/01/2003 – 24/05/2009	GEMS reanalysis	eac1	MC	32r3	T159/L60	IFS/MOZART 3.5 coupled system	Mar 2008 – Sep 2009
01/01/2003 – 31/12/2012	MACC reanalysis (MACCRA)	rean	MC	36r1	T255/L60	IFS/MOZART 3.5 coupled system	Mar 2010 – Feb 2012
01/01/2003 – near real time	CAMS interim reanalysis (CIRA)	eac3	MC	40r2/41r1	T159/L60	IFS(CB05)	Dec 2014 – Dec 2016, then continued in near real time
01/01/2003 – near real time	CAMS reanalysis (CAMSRA)	eac4	MC	42r1	T255/L60	IFS(CB05)	Jan 2017 onwards

TABLE 1 Reanalyses of atmospheric composition produced with the GEMS, MACC and CAMS systems. EXP and CLASS information is needed to retrieve the data from ECMWF's meteorological archive (MARS).

doi: 10.21957/64gx3ai8cv

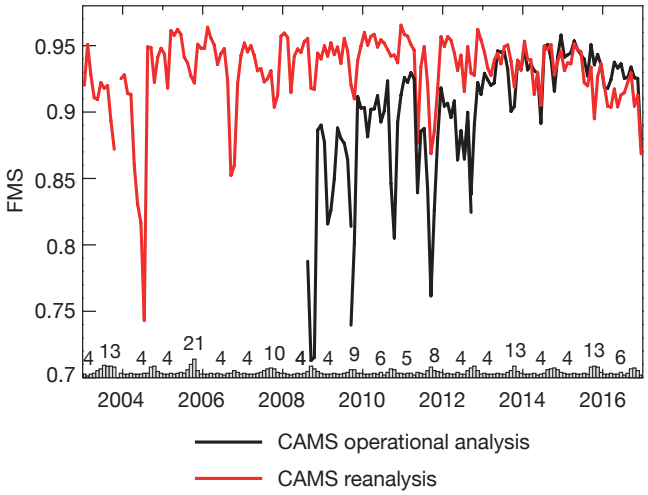

FIGURE 1 Time series from 2003 to 2016 of the FMS score for ozone profiles (1,000–3 hPa) at the Antarctic Neumayer station from the CAMS operational analysis and the CAMS reanalysis. The bar chart at the bottom shows the number of ozonesonde profiles per month used for validation.

previous atmospheric composition reanalyses: it has smaller biases compared to independent O_3, CO, NO_2 and AOD observations and is more consistent in time, especially compared to the MACC reanalysis.

Ensuring consistency

ECMWF has been producing atmospheric composition forecasts and analyses for over a decade (Flemming et al., 2017a). The model and data assimilation system used for this was developed as a European effort by a consortium of partners in the EU-funded GEMS and MACC projects. Operated by CAMS, the forecasting system has been fully operational at ECMWF since January 2015. Since the model, the data assimilation system, and the observations and emissions used have changed considerably over time, it is difficult if not impossible to compare operational forecast data from a recent period with earlier data in a meaningful way (e.g. to determine trends or seasonal anomalies). This is why reanalyses are produced: they analyse atmospheric composition over a long period of time using a single version of the modelling and data assimilation system, while taking care to minimise changes in the versions of the emissions used or satellite retrievals assimilated. Such a system provides the temporal consistency needed to identify trends or to provide maps of annual or seasonal anomalies.

To illustrate the advantage a reanalysis has over analysis data produced by a continuously evolving operational model, Figure 1 shows the 'Figure of Merit in Space' (FMS) ozone score at the Antarctic Neumayer station for the CAMS operational analysis and the CAMS reanalysis. The FMS score is a measure of the fit between model ozone profiles and ozonesonde profiles (here calculated from the surface to 3 hPa) and has a score between 1 (perfect fit) and 0. Figure 1 illustrates the improvements in the CAMS operational analysis over the years. For instance, in the earlier years the CAMS system did not adequately reproduce the low values and vertical distribution of the Antarctic ozone hole, which is reflected by the low scores in the austral spring from 2008 to 2012. Model improvements and the assimilation of improved O_3 retrievals led to better O_3 scores in more recent years. In contrast, CAMSRA, which uses a model version introduced as recently as 2015, has much better scores than the operational analysis during the earlier years and a better consistency in performance throughout the period.

Key features

The chemistry scheme of the IFS used in CAMSRA is an extended version of the Carbon Bond Mechanism 5 (CB05) as implemented in the Chemical Transport Model (CTM) Transport Model 5 (TM5) and is documented in Flemming et al. (2017b). This is a tropospheric chemistry scheme. For stratospheric ozone, the chemical tendencies above the tropopause are computed by a parametrization based on Cariolle & Teyssèdre (2007). The CAMS aerosol model component of the IFS is a hybrid bulk/bin scheme with 12 prognostic tracers, consisting of three bins for sea salt depending on size (0.03–0.5, 0.5–5 and 5–20 μm), three bins for dust (0.030–0.55, 0.55–0.9 and 0.9–20 μm), hydrophilic and hydrophobic organic matter (OM) and black carbon (BC), plus sulphate aerosol and a gas-phase sulphur dioxide (SO_2) precursor (Morcrette et al., 2009). Updates to the chemistry and aerosol schemes that are specific to CAMSRA are given in Inness et al. (2018).

The CAMS system uses MACCity anthropogenic emissions, biomass burning emissions from the CAMS Global Fire Assimilation System and biogenic emissions from the MEGAN model. Great care was taken to ensure that the emission datasets for CAMSRA were consistent in time and that consistent anthropogenic and biomass burning emissions were used for the aerosol and chemistry fields.

The data assimilation system for the atmospheric composition fields remains unchanged from the one described by Inness et al. (2015). CAMSRA uses 4-dimensional variational data assimilation (4D-Var) with 12-hour assimilation windows. The atmospheric composition fields for O_3, CO, NO_2 and AOD are included in the 4D-Var minimisation, which is carried out together with the meteorological variables.

Figure 2 shows a time series of the observations of atmospheric composition data that were assimilated in the CAMS reanalysis. These include O_3 retrievals from a range of satellite sensors (SCIAMACHY, OMI, GOME-2, MIPAS, MLS, SBUV/2); total column CO (TCCO)

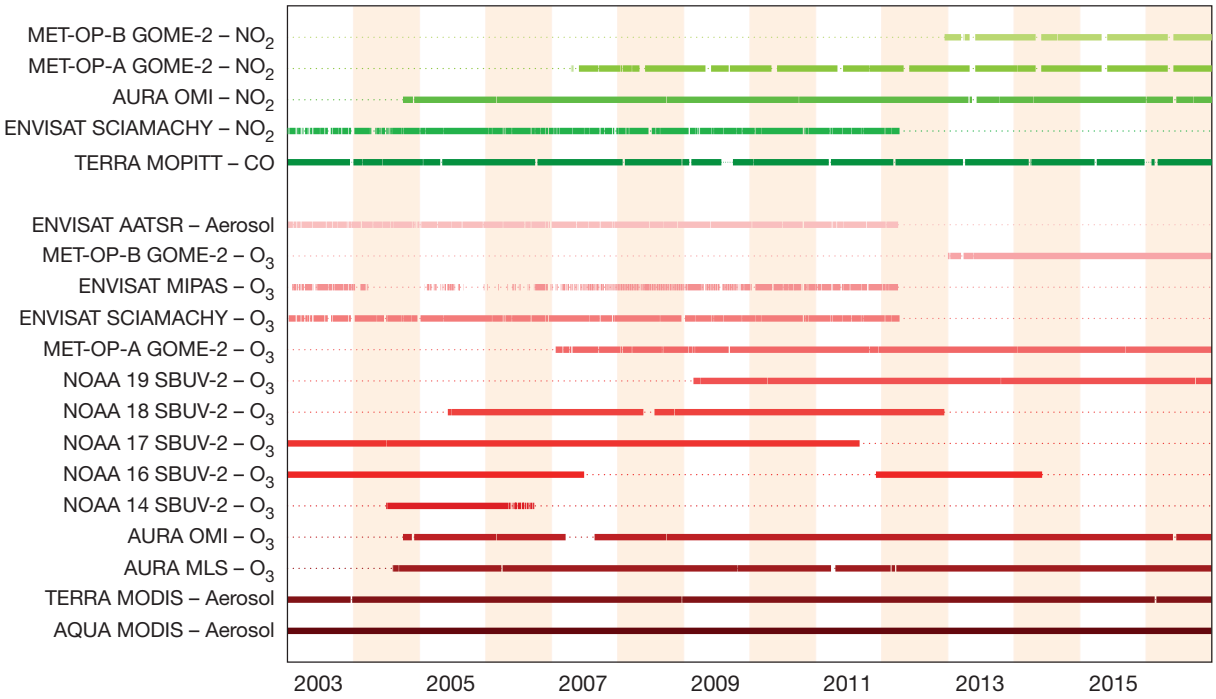

FIGURE 2 Atmospheric composition data assimilated in the CAMS reanalysis between 2003 and 2016.

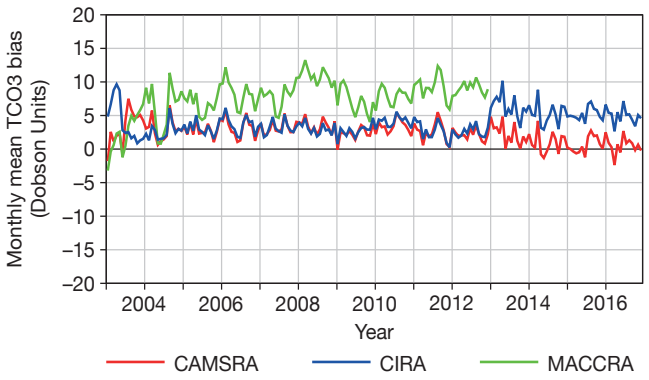

FIGURE 3 Time series of global monthly mean TCO3 bias for CAMSRA, CIRA and MACCRA data validated against data from the WOUDC ozone database. About 50–60 WOUDC stations were available from 2003 to 2014 and about 40 stations after 2014.

retrievals from MOPITT; tropospheric column NO_2 retrievals from SCIAMACHY, OMI and GOME-2; and retrievals of total AOD at 550 nm from MODIS and AATSR. Variational bias correction is applied to the total column O_3, NO_2 and AOD data to ensure good time consistency when blending the various datasets. More details about the model, the emission datasets, the data assimilation system and the assimilated satellite data used in CAMSRA, as well as references for the datasets and time series that show the quality of the datasets, can be found in Inness et al. (2018).

First evaluation results

First results on the quality of CAMSRA data for O_3, CO and AOD are shown here by comparing the fields from CAMSRA with independent observations and fields from MACCRA and CIRA. The quality assessment of the reanalysis is a combined effort of ECMWF staff and a large consortium of European partner institutes contracted to validate CAMS global services. This external quality assessment makes use of Europe-wide expertise to provide an independent assessment of the data products. A comprehensive quality assessment of CAMSRA is given in a validation report available on the CAMS website (Eskes et al., 2018) and a summary is also given in Inness et al. (2018).

The seasonal mean total column O_3 (TCO3) fields from CAMSRA and CIRA agree to within 1% when averaged over the period 2003–2016. The differences between CAMSRA and MACCRA are larger but still within 5%. All reanalyses have small positive biases with respect to independent TCO3 observations from the WOUDC ozone database, with MACCRA having the largest biases and CAMSRA the smallest (Figure 3). O_3 from CAMSRA is more consistent in time than from MACCRA. Agreement with ozonesondes is within 10% in the long-term global mean (Figure 4). The best agreement between the reanalyses and the sondes is found in the stratosphere, where the assimilated O_3 observations constrain the analyses well. Differences between the reanalyses are larger in the troposphere, where the impact of the assimilation is smaller (Inness et al., 2015) and differences in the chemistry schemes, emissions and transport become more important. CAMSRA and CIRA agree better with ozonesondes in the tropical mid- to upper troposphere than MACCRA,

which shows a large underestimation here (–30%). CAMSRA agrees better with ozonesondes above 15 hPa than MACCRA, which overestimates O_3 there. This makes CAMSRA a better dataset to be used as the climatology in e.g. radiation schemes or radiance observation operators.

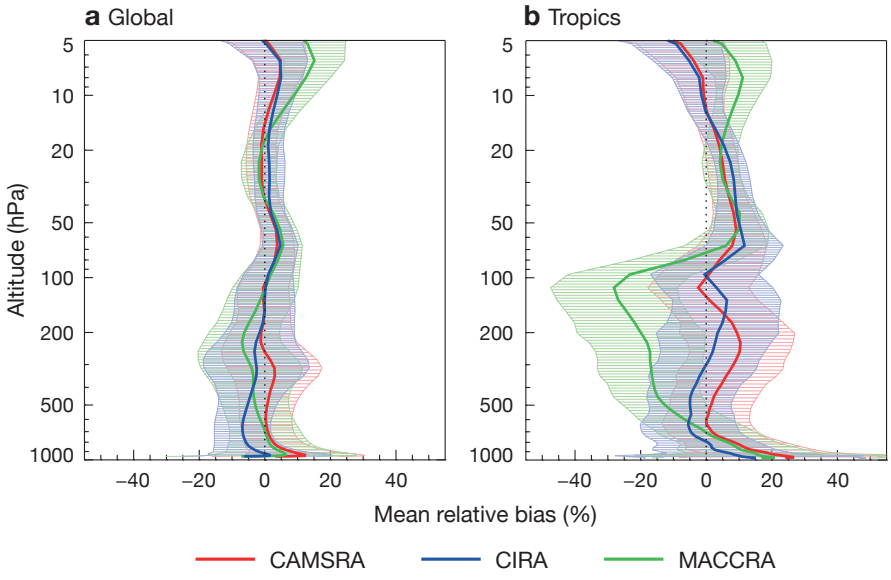

FIGURE 4 Mean relative O_3 bias for CAMSRA, CIRA and MACCRA data validated against ozonesonde data, averaged (a) over the globe and (b) over the tropics. The hatched areas show one respective standard deviation. For CAMSRA and CIRA, the average is calculated over the period 2003–2016, for MACCRA over the period 2003–2012.

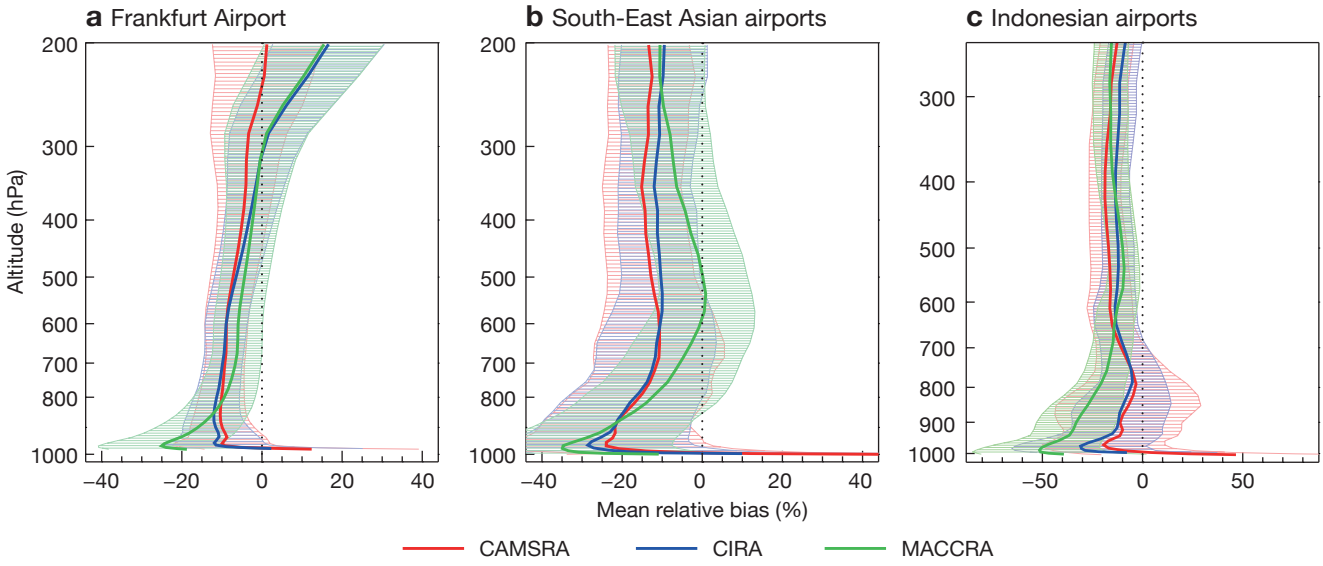

FIGURE 5 Mean relative CO bias for CAMSRA, CIRA and MACCRA data validated against IAGOS aircraft data at (a) Frankfurt Airport, (b) South-East Asian airports and (c) Indonesian airports (note the different x-axis scale). The hatched areas show one respective standard deviation. For CAMSRA and CIRA, the average is calculated over the period 2003–2016, for MACCRA over the period 2003–2012.

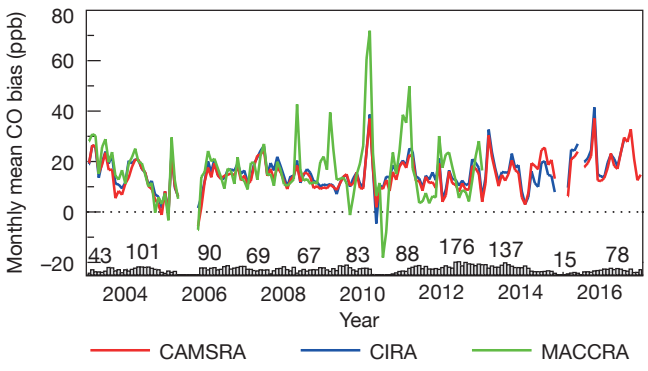

FIGURE 6 Time series of monthly mean CO biases for CAMSRA, CIRA and MACCRA data validated against IAGOS CO data in the lower troposphere (950–700 hPa) at Frankfurt Airport. The bar chart at the bottom shows the number of IAGOS observations used for validation.

Differences in TCCO between CAMSRA and CIRA are smaller than 5% when averaged over the period 2003–2016. The differences between CAMSRA and MACCRA are larger, firstly because MACCRA employed a different chemistry scheme, and secondly because different fire emissions were used. Comparisons with IAGOS aircraft observations show an underestimation of CO in the free troposphere at Frankfurt Airport and South-East Asian airports for all three reanalyses, with larger underestimation in the lower troposphere (Figure 5). This underestimation is similar in CAMSRA and CIRA, while MACCRA has larger negative biases in the lower troposphere. Over Indonesia, CAMSRA and CIRA are very different from MACCRA below 700 hPa and have smaller biases. This is likely due to differences in the fire emission data used.

A time series of the CO bias in the lower troposphere against IAGOS data at Frankfurt Airport shows that CAMSRA is more consistent in time than MACCRA (Figure 6). In MACCRA, the assimilation of IASI TCCO satellite retrievals was included from 2008, which led to a change in the CO burden, while in CAMSRA and CIRA only TCCO data from the MOPITT satellite instrument were assimilated to achieve better temporal consistency.

Total AOD values in CAMSRA are lower than in CIRA or MACCRA in many areas, but larger over India and South-East Asia, and they agree better with observations of total AOD from the AERONET database (Figure 7). The validation with AERONET data also shows some biases at measuring stations near outgassing volcanoes (in particular Mauna Loa in Hawaii and Popocatepetl near Mexico City). The very large AOD biases at those locations in CAMSRA degrade the global average bias. This is a side effect of possibly erroneous model treatment of diffuse volcanic emissions. The model-resolution orography does not resolve the height of the volcanoes. As a result, the

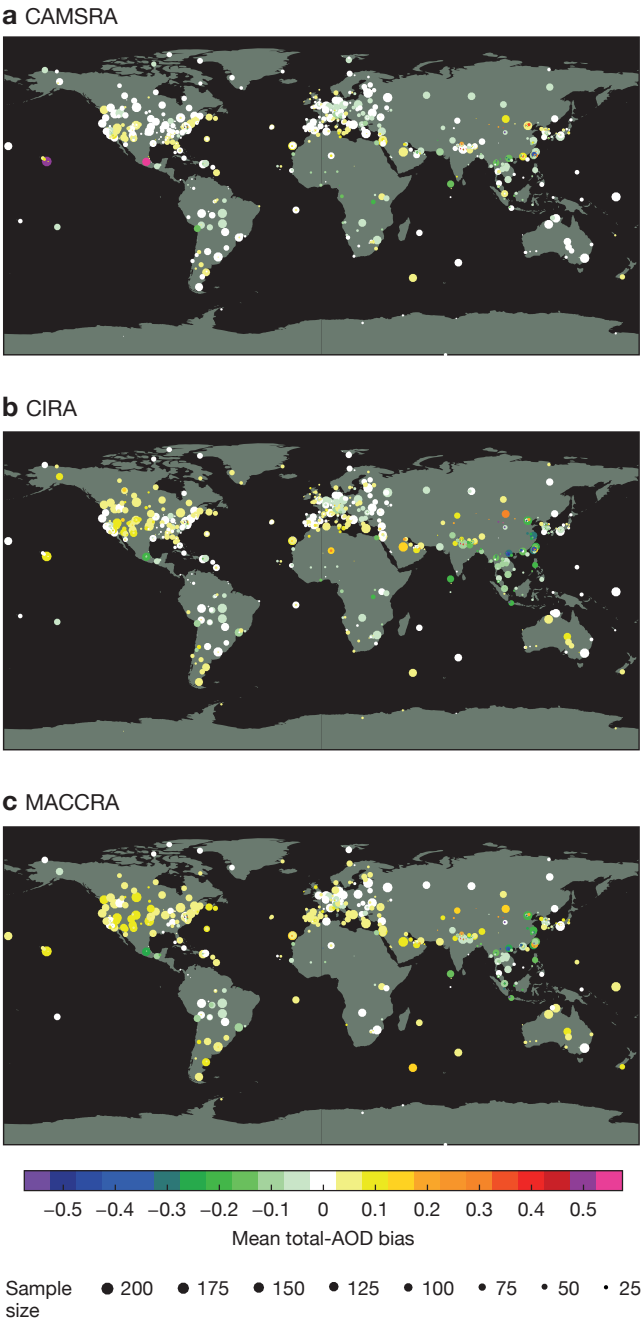

FIGURE 7 Mean total-AOD bias for (a) CAMSRA, (b) CIRA and (c) MACCRA data, validated against AERONET observations. For CAMSRA and CIRA, the average is calculated over the period 2003–2016, for MACCRA over the period 2003–2012.

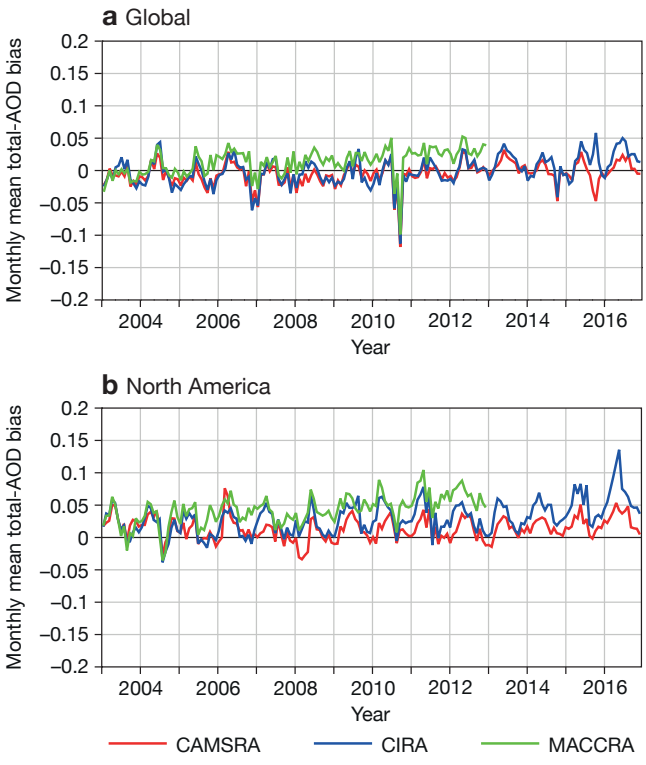

FIGURE 8 Time series of monthly mean total-AOD bias for CAMSRA, CIRA and MACCRA data validated against AERONET observations for (a) the globe and (b) North America. The Mauna Loa and Mexico City stations were excluded from these time series as they are unrepresentative and skew the statistics.

model misrepresents the altitude of the volcanic plumes. In combination with recent enhancements to the SO_2 oxidation scheme, which improve aerosol on the global scale, this gives rise to the local biases. When calculating global mean AOD statistics, it is advisable to exclude the Mauna Loa and Mexico City stations as unrepresentative.

AOD in CAMSRA is more consistent in time than in CIRA and MACCRA, especially over Europe and North America, where CIRA and MACCRA show an increasingly positive bias with time (Figure 8).

There are large differences in the different aerosol species between the three reanalyses (Figure 9). Part of

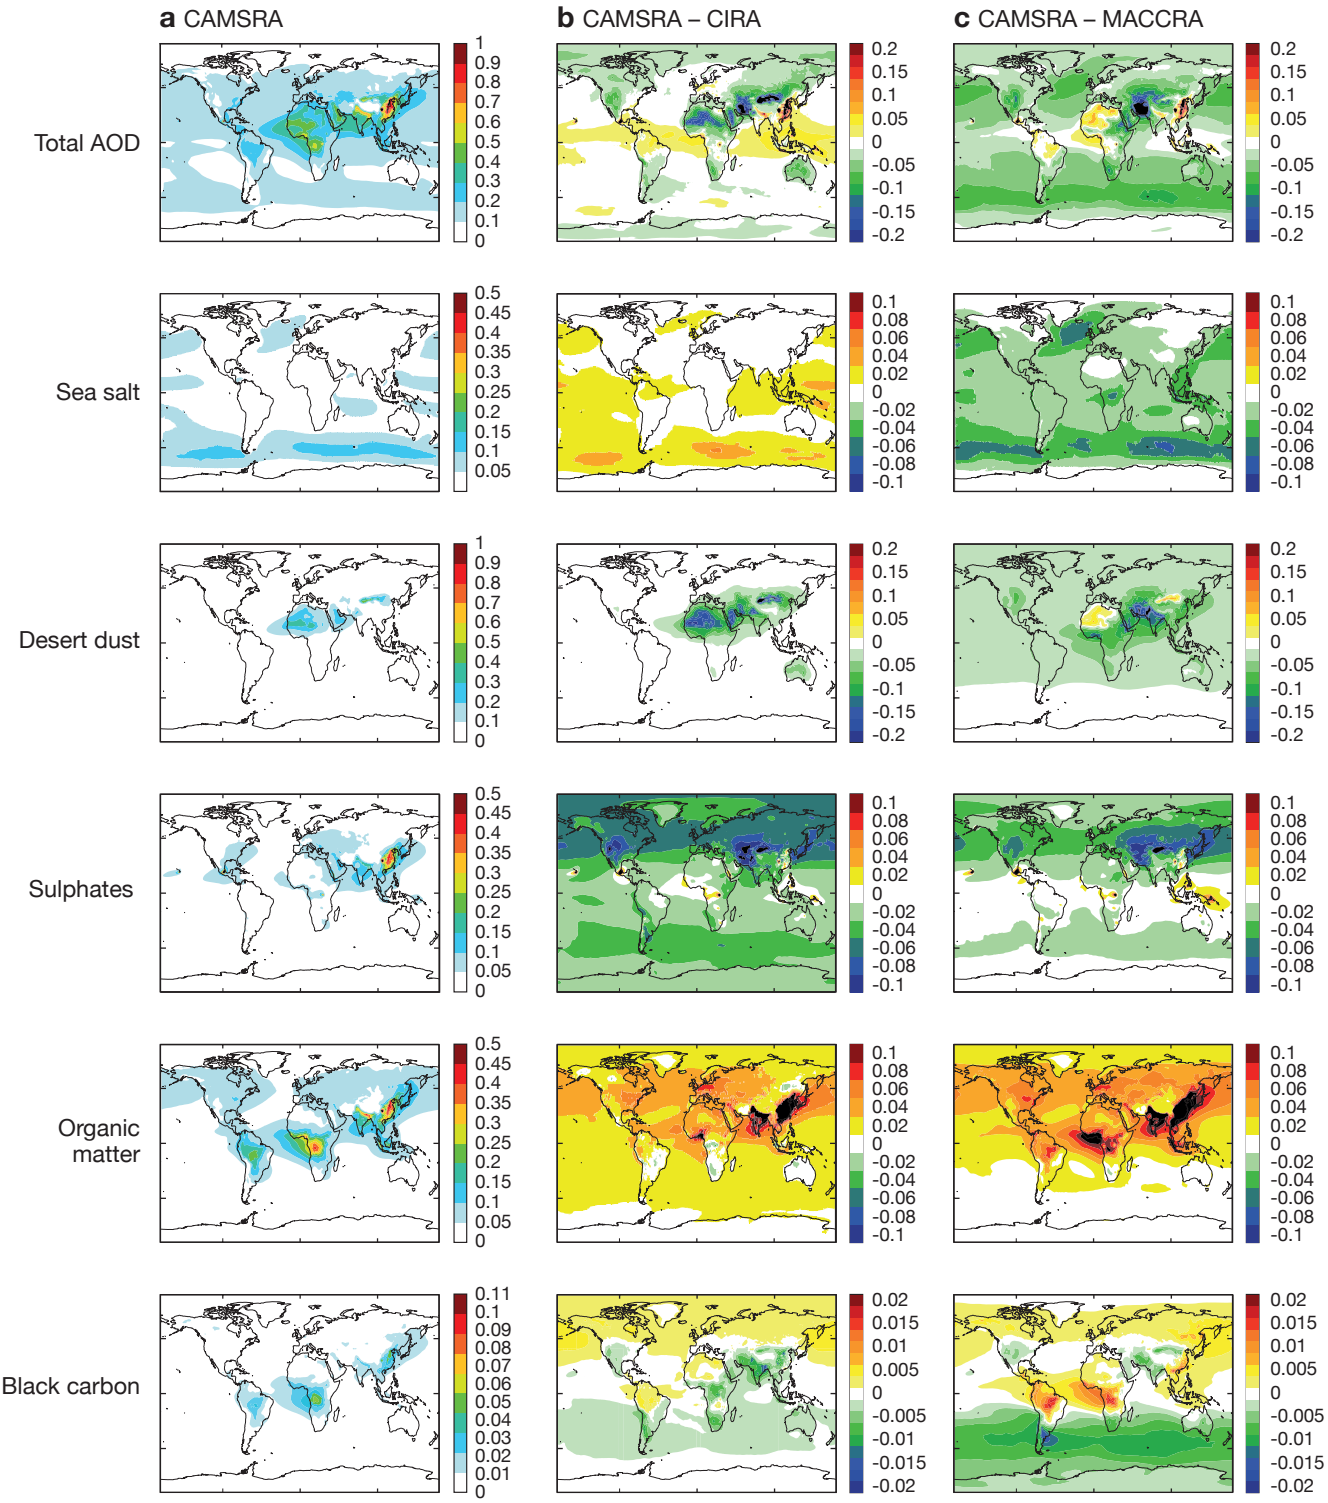

FIGURE 9 Annually averaged AOD species from (a) CAMSRA, (b) the difference between CAMSRA and CIRA and (c) the difference between CAMSRA and MACCRA (2003–2012 only). AOD is unitless.

the reason is that aerosol speciation is not well constrained by the assimilated AOD observations. Relative to the two earlier reanalyses, CAMSRA shows a reduction in desert dust, sulphates and black carbon in the southern hemisphere, compensated by an increase in organic matter and black carbon in the northern hemisphere. The reduction in sulphate globally is particularly strong relative to CIRA, where sulphate was greatly overestimated (Flemming et al., 2017b). More work is needed to validate the individual aerosol components.

Access to data and outlook

CAMSRA is of better quality and provides better temporal consistency than its predecessors. It also provides more chemical species, and data are available at a higher temporal and spatial resolution. In total, 56 tropospheric chemical species of the CB05 chemical mechanism, 12 aerosol components and many additional diagnostics, such as total columns and extinction coefficients, can be obtained from the CAMS reanalysis. The CAMS reanalysis data are freely available from *http://atmosphere.copernicus.eu* and can serve a multitude of users, ranging from small and medium-sized enterprises in the solar-energy sector to scientists and policy-makers. The data can be used to analyse the state of the atmosphere or to identify trends that have developed over the past years or decades. Furthermore, the CAMS reanalysis can be used to compute climatologies, evaluate models, benchmark other reanalyses or serve as boundary conditions for regional models for past periods.

The CAMS reanalysis is being continued, running shortly behind real time, and additional years will become available in the future, one year each year. A reanalysis for the greenhouse gases CH_4 and CO_2 is currently being produced separately.

The following institutes contribute to the global validation of CAMS products: Royal Netherlands Meteorological Institute (KNMI); Academy of Athens (AA); Aarhus University (AU); Aristotle University of Thessaloniki (AUTH); Institut d'Aéronomie spatiale de Belgique (BIRA-IASB); Barcelona Supercomputing Center and the State Meteorological Agency of Spain (BSC/AEMET); Laboratoire des Sciences du Climat et de l'Environnement (CEA-LSCE); Centre National de la Recherche Scientifique et Université Paul Sabatier – Laboratoire d'Aérologie (CNRS-LA); Deutscher Wetterdienst – Hohenpeissenberg Meteorological Observatory (DWD); Institute of Environmental Physics, University of Bremen (IUP-UB); Norwegian Meteorological Insitute (MET-NO); Max Planck Institute for Meteorology (MPG); Science and Technology (S&T); Uni Bremen Campus GmbH (UBC).

Further reading

Cariolle, **D.** & **H. Teyssèdre**, 2007: A revised linear ozone photochemistry parameterization for use in transport and general circulation models: multi-annual simulations. *Atmos. Chem. Phys.*, **7**, 2183–2196.

Eskes, **H.J.**, **Y. Bennouna**, **M. Schulz**, **Y. Christophe**, **S. Basart**, **A. Benedictow**, **A.-M. Blechschmidt**, **S. Habrillat**, **H. Clark**, **E. Cuevas**, **H. Flentje**, **K.M. Hansen**, **U. Im**, **J. Apsomenakis**, **B. Langerock**, **K. Petersen**, **A. Richter**, **N. Sudarchikova**, **V. Thouret**, **A. Wagner**, **Y. Wang** & **C. Zerefos**, 2018: Validation report of the CAMS global reanalysis of aerosols and reactive gases, years 2003–2016. *Copernicus Atmosphere Monitoring Service (CAMS) Report*.

Flemming, **J.**, **V.-H. Peuch** & **L. Jones**, 2017a: Ten years of forecasting atmospheric composition at ECMWF. *ECMWF Newsletter* **No. 152**, 5–6.

Flemming, **J.**, **A. Benedetti**, **A. Inness**, **R.J. Engelen**, **L. Jones**, **V. Huijnen**, **S. Remy**, **M. Parrington**, **M. Suttie**, **A. Bozzo**, **V.-H. Peuch**, **D. Akritidis** & **E. Katragkou**, 2017b: The CAMS interim Reanalysis of Carbon Monoxide, Ozone and Aerosol for 2003–2015, *Atmos. Chem. Phys.*, **17**, 1945–1983, doi:10.5194/acp-17-1945-2017.

Inness, **A.**, **A.-M. Blechschmidt**, **I. Bouarar**, **S. Chabrillat**, **M. Crepulja**, **R.J. Engelen**, **H. Eskes**, **J. Flemming**, **A. Gaudel**, **F. Hendrick**, **V. Huijnen**, **L. Jones**, **J. Kapsomenakis**, **E. Katragkou**, **A. Keppens**, **B. Langerock**, **M. de Mazière**, **D. Melas**, **M. Parrington**, **V.H. Peuch**, **M. Razinger**, **A. Richter**, **M.G. Schultz**, **M. Suttie**, **V. Thouret**, **M. Vrekoussis**, **A. Wagner** & **C. Zerefos**, 2015: Data assimilation of satellite-retrieved ozone, carbon monoxide and nitrogen dioxide with ECMWF's Composition-IFS, *Atmos. Chem. Phys.*, **15**, 5275–5303, doi:10.5194/acp-15-5275-2015.

Inness, **A.**, **M. Ades**, **A. Agusti-Panareda**, **J. Barré**, **A. Benedictow**, **A.M. Blechschmidt**, **J. Dominguez**, **R. Engelen**, **H.J. Eskes**, **J. Flemming**, **V. Huijnen**, **L. Jones**, **Z. Kipling**, **S. Massart**, **M. Parrington**, **V.-H. Peuch**, **M. Razinger**, **S. Remy**, **M. Schulz** & **M. Suttie**, 2018: The CAMS reanalysis of atmospheric composition. *Atmos. Chem. Phys.* Discuss., doi:10.5194/acp-2018-1078, under review.

Morcrette, **J.J.**, **O. Boucher**, **L. Jones**, **D. Salmond**, **P. Bechtold**, **A. Beljaars**, **A. Benedetti**, **A. Bonet**, **J.W. Kaiser**, **M. Razinger**, **M. Schulz**, **S. Serrar**, **A.J. Simmons**, **M. Sofiev**, **M. Suttie**, **A.M. Tompkins** & **A. Untch**, 2009: Aerosol analysis and forecast in the European Centre for Medium-Range Weather Forecasts Integrated Forecast System: Forward modeling, *J. Geophys. Res.*, **114**, doi:10.1029/2008JD011235.

ECMWF Council and its committees

The following provides some information about the responsibilities of the ECMWF Council and its committees. More details can be found at:

http://www.ecmwf.int/en/about/who-we-are/governance

Council

The Council adopts measures to implement the ECMWF Convention; the responsibilities include admission of new members, authorising the Director-General to negotiate and conclude co-operation agreements, and adopting the annual budget, the scale of financial contributions of the Member States, the Financial Regulations and the Staff Regulations, the long-term strategy and the programme of activities of the Centre.

President Prof. Miguel Miranda *(Portugal)*

Vice President Prof. Juhani Damski *(Finland)*

Policy Advisory Committee (PAC)

The PAC provides the Council with opinions and recommendations on any matters concerning ECMWF policy submitted to it by the Council, especially those arising out of the four-year programme of activities and the long-term strategy.

Chair Mr Rolf Brennerfelt *(Sweden)*

Vice Chair Mr Eoin Moran *(Ireland)*

Finance Committee (FC)

The FC provides the Council with opinions and recommendations on all administrative and financial matters submitted to the Council and exercises the financial powers delegated to it by the Council.

Chair Mr Mark Hodkinson *(United Kingdom)*

Vice Chair Dr Gisela Seuffert *(Germany)*

Scientific Advisory Committee (SAC)

The SAC provides the Council with opinions and recommendations on the draft programme of activities of the Centre drawn up by the Director-General and on any other matters submitted to it by the Council.
The 12 members of the SAC are appointed in their personal capacity and are selected from among the scientists of the Member States.

Chair Prof. Alan O'Neill *(United Kingdom)*

Vice Chair Dr Inger-Lise Frogner *(Norway)*

Technical Advisory Committee (TAC)

The TAC provides the Council with advice on the technical and operational aspects of the Centre including the communications network, computer system, operational activities directly affecting Member States, and technical aspects of the four-year programme of activities.

Chair Dr Philippe Steiner *(Switzerland)*

Vice Chair Dr Sarah O'Reilly *(Ireland)*

Advisory Committee for Data Policy (ACDP)

The ACDP provides the Council with opinions and recommendations on matters concerning ECMWF Data Policy and its implementation.

Chair Mr Søren Olufsen *(Denmark)*

Vice Chair Mr Francisco Pascual Perez *(Spain)*

Advisory Committee of Co-operating States (ACCS)

The ACCS provides the Council with opinions and recommendations on the programme of activities of the Centre, and on any matter submitted to it by the Council.

Chair Mr Taimar Ala *(Estonia)*

Vice Chair Mr Nir Stav *(Israel)*

ECMWF publications

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

Technical Memoranda

835 **Stockdale**, **T.**, **M. Alonso Balmaseda**, **S. Johnson**, **L. Ferranti**, **F. Molteni**, **L. Magnusson**, **S. Tietsche**, **F. Vitart**, **D. Decremer**, **A. Weisheimer**, **C.D. Roberts**, **G. Balsamo**, **S. Keeley**, **K. Mogensen**, **H. Zuo**, **M. Mayer** & **B.M. Monge-Sanz**: SEAS5 and the future evolution of the long-range forecast system. *November 2018*

834 **Haiden**, **T.**, **M. Dahoui**, **B. Ingleby**, **P. de Rosnay**, **C. Prates**, **E. Kuscu**, **T. Hewson**, **L. Isaksen**, **D. Richardson**, **H. Zuo** & **L. Jones**: Use of in situ surface observations at ECMWF. *November 2018*

833 **Dragani**, **R.**, **A. Benedetti**, **J. Flemming**, **G. Balsamo**, **M. Diamantakis**, **A.J. Geer**, **R. Hogan**, **T. Stockdale**, **M. Ades**, **A. Agusti-Panareda**, **J. Barré**, **P. Bechtold**, **A. Bozzo**, **H. Hersbach**, **E. Hólm**, **Z. Kipling**, **A. Inness**, **J. Letertre-Danczak**, **S. Massart**, **M. Matricardi**, **T. McNally**, **M. Parrington**, **I. Sandu**, **C. Soci** & **F. Vitart**: Atmospheric Composition priority developments for Numerical Weather Prediction. *December 2018*

832 **Massart**, **S.**: A new hybrid formulation for the background error covariance in the IFS: Implementation aspects. *October 2018*

831 **Haiden**, **T.**, **M. Janousek**, **J.R. Bidlot**, **R. Buizza**, **L. Ferranti**, **F. Prates** & **F. Vitart**: Evaluation of ECMWF forecasts, including the 2018 upgrade. *October 2018*

826 **Salisbury**, **D.**, **K. Mogensen** & **G. Balsamo**: Use of in situ observations to verify the diurnal cycle of sea surface temperature in ECMWF coupled model forecasts. *November 2018*

EUMETSAT/ECMWF Fellowship Programme Research Reports

48 **Weston**, **P.** & **N. Bormann**: Enhancements to the assimilation of ATMS at ECMWF: Observation error update and addition of NOAA-20. *October 2018*

ERA Report Series

27 **Hersbach**, **H.**, **P. de Rosnay**, **B. Bell**, **D. Schepers**, **A. Simmons**, **C. Soci**, **S. Abdalla**, **M. Alonso Balmaseda**, **G. Balsamo**, **P. Bechtold**, **P. Berrisford**, **J. Bidlot**, **E. de Boisséson**, **M. Bonavita**, **P. Browne**, **R. Buizza**, **P. Dahlgren**, **D. Dee**, **R. Dragani**, **M. Diamantakis**, **J. Flemming**, **R. Forbes**, **A. Geer**, **T. Haiden**, **E. Hólm**, **L. Haimberger**, **R. Hogan**, **A. Horányi**, **M. Janisková**, **P. Laloyaux**, **P. Lopez**, **J. Muñoz-Sabater**, **C. Peubey**, **R. Radu**, **D. Richardson**, **J.-N. Thépaut**, **F. Vitart**, **X. Yang**, **E. Zsótér** & **H. Zuo**: Operational global reanalysis: progress, future directions and synergies with NWP. *October 2018*

Contact information

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

Telephone National	0118 949 9000
Telephone International	+44 118 949 9000
Fax	+44 118 986 9450

ECMWF's public website *www.ecmwf.int/*

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

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

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

ECMWF Calendar 2019

Jan 30 – Feb 1	Training course: ecFlow
Feb 4–8	Training course: ecCodes
Feb 11–14	Training course: Use and interpretation of ECMWF products
Feb 25 – Mar 1	NWP training course: Predictability and ensemble forecast systems
Mar 4–8	NWP training course: Parametrization of subgrid physical processes
Mar 11–15	NWP training course: Data assimilation
Mar 18–22	EUMETSAT/ECMWF NWP-SAF training course: Satellite data assimilation
Mar 25–29	Training course: Advanced numerical methods for Earth system modelling
Apr 2–5	Workshop on predictability, dynamics and applications research using the TIGGE and S2S ensembles
Apr 9–11	ACDP and data policy meetings of ECOMET and EUMETSAT
May 2–3	Finance Committee
May 3	Policy Advisory Committee
May 13–17	Security Representatives' meeting
May 13–17	Online training week: Software and computing at ECMWF
Jun 3–6	Using ECMWF's Forecasts (UEF)
Jun 10–14	Workshop on observational campaigns for better weather forecasts
Jun 17–21	OpenIFS user workshop (University of Reading)
Jun 27–28	Council
Sep 2–5	Annual Seminar
Oct 7–9	Scientific Advisory Committee
Oct 7–10	Training course: Use and interpretation of ECMWF products
Oct 10–11	Technical Advisory Committee
Oct 14–16	Workshop on robust scientific developments with reproducible workflows
Oct 28–29	Finance Committee
Oct 29	Policy Advisory Committee
Nov 18–21	Workshop on stratospheric modelling, predictability and influence on the troposphere
Nov 25–29	H-SAF and HEPEX workshop
Dec 10–11	Council

ECMWF

Newsletter | **No. 158** | Winter 2018/19

European Centre for Medium-Range Weather Forecasts

www.ecmwf.int