ECMWF
European Centre for Medium-Range Weather Forecasts

Newsletter

No. 157 | Autumn 2018

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 The ECMWF Newsletter is not peer-reviewed.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/news-centre/media-resources*

More data

On 13 October, the UN-sponsored International Day for Disaster Reduction celebrated how people and communities around the world are reducing their exposure to disasters. Most natural disasters are weather-related, so weather prediction has an important role to play in disaster mitigation and preparedness. In Europe, this was illustrated once again this summer, marked by unusually high temperatures, dry conditions and deadly wildfires. Many people looked to weather forecasters in our Member and Co-operating States for information on what was to come as parts of Europe turned from green to brown.

ECMWF's extended-range predictions showed a signal of warm anomalies weeks in advance. There was also a low river flow signal in the seasonal hydrological outlook of the newly upgraded European Flood Awareness System (EFAS), of which ECMWF is the computational centre. However, the spatial and temporal variability of the heatwave was correctly predicted only up to two weeks in advance, and shorter-range forecasts reflected known biases in predicted maximum and minimum temperatures.

Improving forecast quality requires progress on many fronts. One example is the work on addressing biases in near-surface forecasts undertaken as part of ECMWF's initiative entitled 'Understanding uncertainties in surface–atmosphere exchange' (USURF), described in this Newsletter. Another is more data: more and better weather observations will help to correctly initialise global forecasts and thus to improve forecast quality. An important step was taken in August, when the European Space Agency (ESA) launched the Aeolus satellite to measure global wind profiles using ground-breaking laser technology. As the Level-2 Meteorological Processing Facility, ECMWF has a key role in assessing Aeolus data and providing wind products to other NWP centres through ESA and EUMETSAT. After a full assessment, we expect to start assimilating these data next year. Aeolus will fill a data gap that is particularly acute over the tropics. The new data is expected to have a significant impact on forecast quality in the tropics and a smaller but still important impact in the extratropics.

The excitement over Aeolus should not make us forget the day-to-day work that goes into obtaining other kinds of weather data and monitoring their quality. This is illustrated by two examples described in this Newsletter: encouraging results on radiosonde descent data, and the continuing EUMETNET programme to launch radiosondes from ships, made possible by the willingness of sailors to fit the launches into their workload.

Looking ahead, ECMWF will not just assimilate more observational data but will also produce more forecast data as we aim to increase the resolution of our ensemble forecasts. ECMWF's outgoing Lead Scientist Roberto Buizza clearly restates the rationale for ensemble forecasts in an interview in this Newsletter. Roberto played a key role in the development of ensemble forecasts at the Centre more than 25 years ago and we wish him well in his new venture as a Full Professor of Physics at the Scuola Superiore Sant'Anna in Pisa, Italy.

Florence Rabier
Director-General

Contents

Editor Georg Lentze • **Typesetting & Graphics** Anabel Bowen • **Cover image** Satellite image of Europe 25 July 2018 – ESA

Forecasting the 2018 European heatwave

Linus Magnusson, Laura Ferranti, Freja Vamborg

The late spring and summer of 2018 were among the warmest on record for northern Europe. ECMWF extended-range forecasts predicted warm anomalies weeks in advance, but the northerly extent and intraseasonal variability of the heatwave were only reflected in forecasts up to two weeks ahead.

A chart produced by the Copernicus Climate Change Service (C3S) implemented by ECMWF shows that the near-surface air temperature anomaly in Europe for the period April to August 2018 was far larger than in any previous year since 1979. The strongest anomalies occurred in the Baltic Sea region, while the countries around the Mediterranean experienced close to normal temperatures on average. The heatwave can be divided into three parts: the second half of April, mid-May to mid-June, and the second half of July to the beginning of August, with more normal temperatures in between. While southern Europe had more normal temperatures on average, south-western Europe experienced a surge of heat at the beginning of August, with temperatures reaching 45°C in Spain and Portugal on 4 August.

With these temporal and spatial variations, there are many aspects to verify in ECMWF's forecasts, and not all can be covered here. Focusing on the two main parts of the heatwave, as a first verification we use composites of weekly anomalies from extended-range forecasts covering the period 7 May to 12 August at different lead times (one week, two weeks, three weeks and four weeks). For example, the composite of week-one forecasts issued on Mondays uses forecast days 0 to 7, the composite of week-two forecasts uses forecast days 8 to 14, etc. The predicted anomalies in week-one and week-two forecasts resemble the spatial pattern of the anomalies in the analysis well. Warm anomalies are also present in the week-three and week-four forecasts, but they are weaker and the forecasts did not reflect their northerly extent. It is worth mentioning that, even at the longest time ranges, the

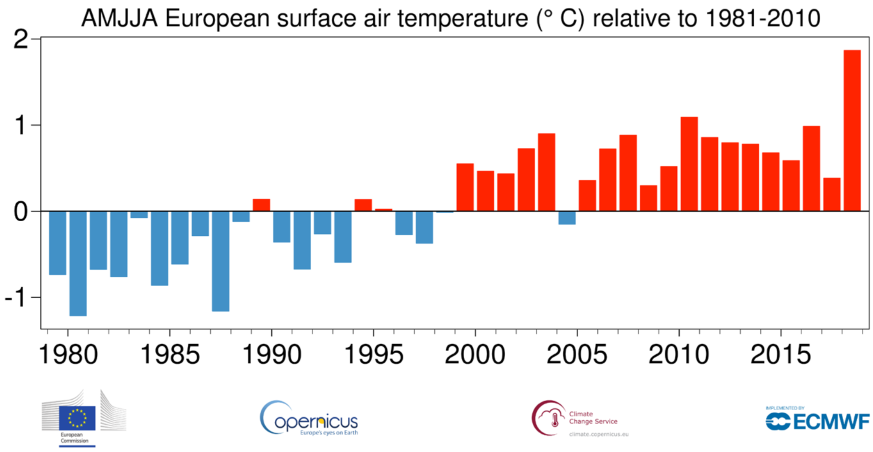

Evolution of near-surface air temperature anomalies. This chart produced by C3S shows that the near-surface air temperature anomaly in Europe in the period of April to August (AMJJA), calculated relative to the 1981–2010 average for those months, was much larger in 2018 than in any previous year since 1979.

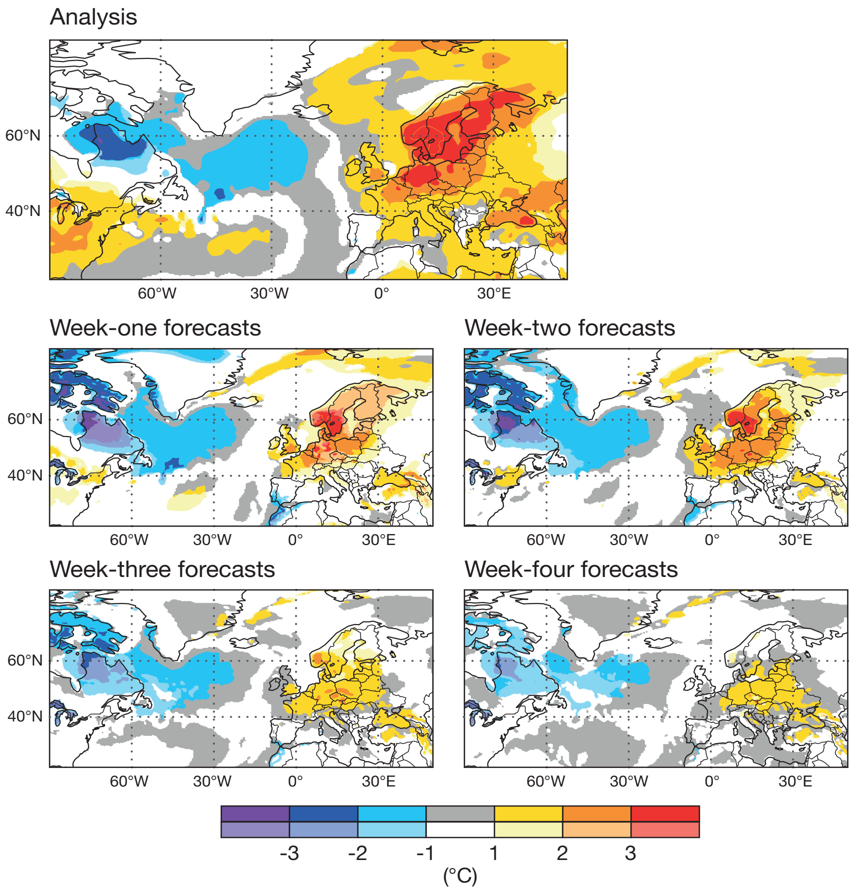

Analysis and forecasts of 2-metre temperature anomalies. The plots show ECMWF's analysis of the average 2-metre temperature anomaly 7 May to 12 August (top) and composites of weekly 2-metre temperature anomalies from extended-range forecasts valid for 7 May to 12 August, based on week-one forecasts (middle left), week-two forecasts (middle right), week-three forecasts (bottom left) and week-four forecasts (bottom right). Saturated colours indicate signficance at the 95% level.

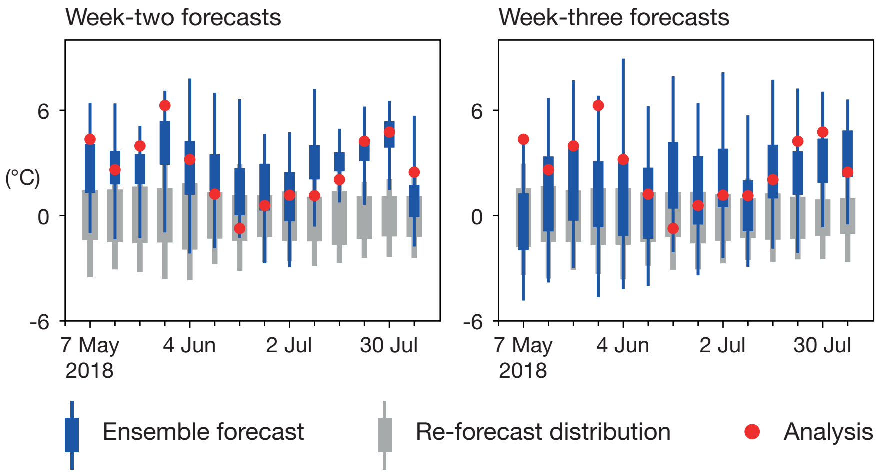

Ensemble forecast distribution of anomalies. Weekly ensemble forecast distribution of 2-metre temperature anomalies for a region in northern Europe (50°N–60°N, 10°E–20°E) for week-two forecasts (left) and week-three forecasts (right). The box-and-whisker symbols show the 5th, 25th, 75th and 95th percentiles.

cold anomaly over north-eastern Canada and the central-northern Atlantic was captured.

The results from this evaluation show that the extended-range forecasts predicted the warm anomaly fairly well on average, but they do not tell us how well the intra-seasonal variability was captured. We have therefore also visualised the week-by-week evolution of 2-metre temperature anomalies averaged over a region in northern Europe (50°N–60°N, 10°E–20°E) in week-two and week-three forecasts, the analysis (used for verification) and the model climate (re-forecast distribution). By definition, the anomalies in the re-forecasts are centred around zero. The plots show that week-two forecasts captured the intra-seasonal variations seen in the analysis reasonably well. Week-three forecasts, on the other hand, showed less variation in the predicted anomalies throughout the summer. They failed to give any indication of the warm peak at the end of May or of the break in the warm weather at the end of June, although they gave some indication of the warm period in the second half of July.

In short-range forecasts, when averaging over the period 7 May to 12 August 2018, a general tendency can be detected across Europe for daily maximum temperatures to have been underestimated and minimum temperatures to have been overestimated. The underestimation of the diurnal cycle in heatwave conditions is one of the topics being explored in ECMWF's USURF project. For more details on USURF, see the article on biases in near-surface forecasts in this Newsletter.

The European heatwave of 2018 poses questions about the driving mechanisms behind the strong anomalies. Evaluating the predictability of heatwaves is part of ECMWF's activities in the subseasonal-to-seasonal (S2S) project sponsored by the World Weather Research Programme (WWRP) and the World Climate Research Programme (WCRP). It will also be part of the newly approved EU-funded Horizon 2020 CAFE project, which starts in 2019 and in which ECMWF is a partner.

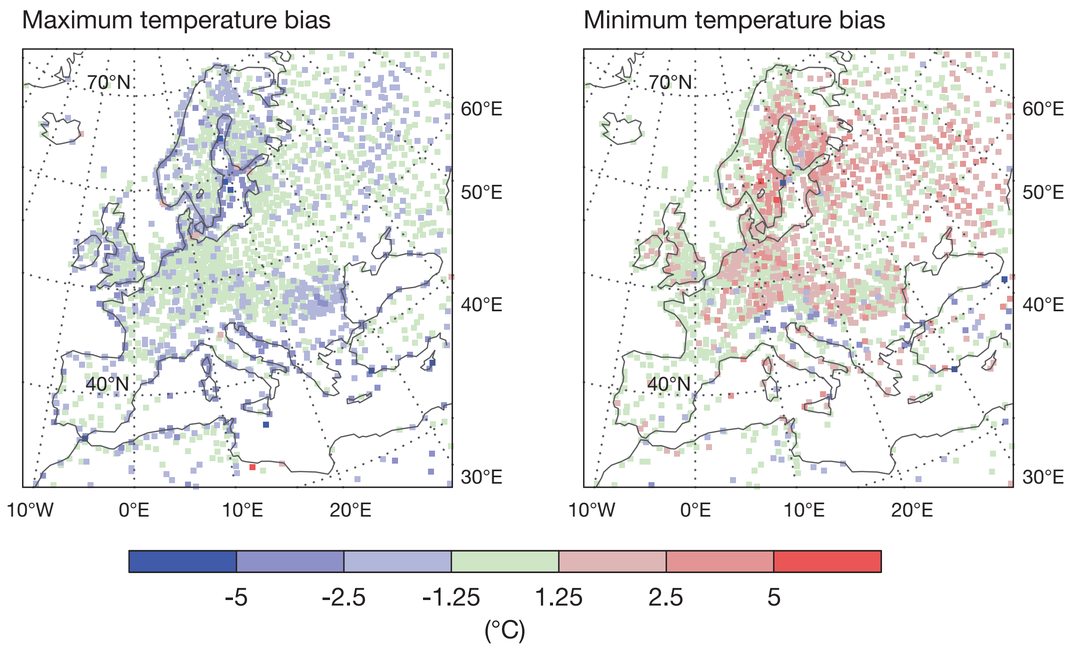

Underestimation of the diurnal cycle. The plots show the bias in maximum (left) and minimum (right) 2-metre temperature for day-two forecasts between 7 May and 12 August 2018, compared to SYNOP weather station observations. Blue colours indicate that on average the forecasts were too cold, red colours that they were too warm.

Low river flow signal during Europe's dry summer

Louise Arnal, Shaun Harrigan, David Lavers, Fredrik Wetterhall, Christel Prudhomme

The summer 2018 in Europe was remarkable from a hydroclimate perspective, with many regions experiencing persistent dry or drought conditions, a series of heatwaves with record-breaking temperatures and numerous wildfires. The European Flood Awareness System (EFAS) seasonal hydrological outlook, which is produced operationally at ECMWF as part of the Centre's work for the Copernicus Emergency Management Service (CEMS), managed to capture the intensifying low-flow signal throughout the summer with lead times up to two months.

A developing drought

The summer's extreme weather was caused by persistent blocking high-pressure systems over northern and central Europe, which were part of a summer variation of the North Atlantic Oscillation (NAO) pattern. These meteorological conditions had severe hydrological impacts. According to the European Drought Observatory (EDO), much of central and northern Europe was affected by drought, resulting in very dry soil and low river, groundwater and reservoir levels. A remarkable feature over this period was the growth in the spatial extent of low flows across the European river network. According to the daily EFAS discharge simulation, the low flow area covered about 10% of the extended EFAS domain river network at the beginning of June (primarily in southern Sweden and northern Africa) and increased to a maximum of 41% on 24 August. For details on the recently upgraded EFAS system used to produce the simulation, see the separate article on the EFAS upgrade in this Newsletter.

Low river flow prediction

The EFAS seasonal hydrological outlook issued on 1 July, valid from 1 July to the end of August, showed a low-flow signal in large parts of Europe up to two months in advance. The predictability of this event can be largely attributed to the dry initial hydrological conditions at the start of the forecast, carried over into the future through the land surface memory. The lack of predictability in western Europe is likely due to the wet anomaly predicted for August by ECMWF's SEAS5 seasonal forecast issued on 1 July. The next SEAS5 forecast, issued on 1 August, showed a drier signal in this part of Europe. Moreover, the hydrological model used to produce these outputs has known limitations on the Iberian Peninsula. Looking back at this event and similar ones (e.g. the summer of 2003 and of 2015) in more detail should help us understand the sources of predictability and uncertainty in the EFAS seasonal hydrological outlook and should lead to improved forecasts of future events.

A more detailed version of this article was published on 6 September 2018 as a blog post on the HEPEX website: https://hepex.irstea.fr/summer-2018-in-europe/

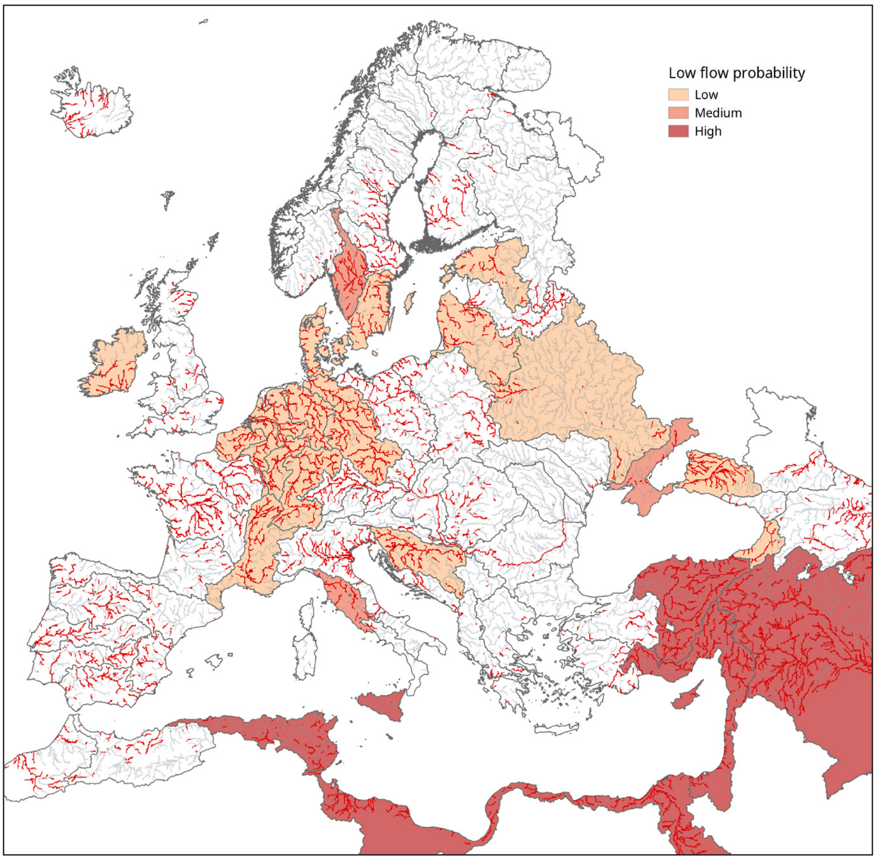

EFAS seasonal hydrological outlook for the last week of August 2018. The shading indicates different levels of elevated probabilities of low river flow, based on the EFAS seasonal discharge forecast issued on 1 July for the last week of August 2018. 'Low', 'medium' or 'high' probability indicates that 50–75%, 75–90%, or more than 90% of the ensemble distribution is below the 10th percentile of the EFAS discharge simulation weekly climatology. The red river network shows the low flow situation from the EFAS discharge simulation for 24 August as a proxy for the truth.

How many ensemble members are desirable?

Martin Leutbecher

A recent study has looked at the implications of changing the number of members in ensemble forecasting from two different perspectives: that of forecast users and that of scientists developing a numerical weather prediction (NWP) system. It found that users could benefit from increases in ensemble size even beyond the ensemble size of 50 used operationally at ECMWF, while scientists could become more productive by using ensembles with fewer than 10 members. However, the scientists' development work needs to be designed carefully to enable conclusions that are relevant for large ensemble sizes. The key elements for meaningful development work with few members are the use of 'fair scores' and an ensemble generation methodology that provides exchangeable members.

The users' perspective

Probabilistic forecasts are based on ensembles, and the number of members determines among other things how well a probability distribution can be estimated. The figure illustrates this for a simulated situation where ensembles of various sizes were obtained by repeatedly sampling members from a known distribution. Probability densities based on the different ensembles were estimated and plotted for each ensemble size. The differences among the realisations illustrate the magnitude of the sampling uncertainties that are to be expected for a range of ensemble sizes. They are not negligible for ensemble sizes used in operational weather forecasting.

The sensitivity of forecast quality to ensemble size varies depending on which verification metrics are considered. A number of scores converge with ensemble size M as 1 + 1/M to the score of an ensemble with an infinite number of members, provided the ensemble is reliable. An example is the continuous ranked probability score (CRPS), which is zero for a perfect deterministic forecast. Even if perfect reliability cannot be assumed, the convergence holds reasonably well, as results with the ECMWF ensemble for upper-air variables in the extratropics demonstrate for ensembles with up to 200 members. The convergence result implies that changing the ensemble size from M_1 to M_2 changes the score by $100 (M_1 - M_2)/M_2(M_1 + 1)$ per cent. Thus, skill measured with the CRPS would increase (decrease) by 1% (3%) if the ensemble size changed from 50 to 100 (20). For a score that measures the skill of quantiles in the tail of the distribution, the sensitivity to ensemble size is considerably larger. This is particularly relevant for users who make decisions for low cost-loss ratios.

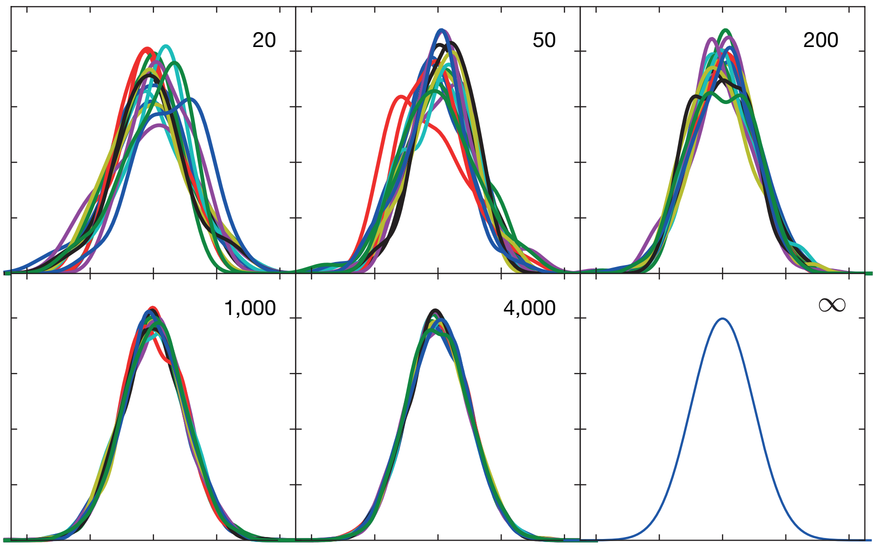

Probability densities for different ensemble sizes. Probability densities were estimated for ensembles with 20, 50, 200, 1,000 and 4,000 members. The underlying true distribution is also shown (bottom right). For each ensemble size, 16 different ensembles were generated from the underlying distribution. The variations among the densities illustrate the sampling uncertainty due to the finite ensemble size.

The scientists' perspective

For research and development, testing with a large ensemble size is a burden. The larger the ensemble, the longer scientists have to wait for the results to be generated, which means that progress will be slower. However, for some of the most relevant scores it is possible to make statistical corrections, which adjust the value of the score obtained from a small ensemble to the value of the score of a large ensemble or the underlying distribution even if the ensemble is not reliable. In a situation where ensemble members are exchangeable, it can be proven that the expected value of the adjusted score is identical to the score that would be obtained with infinite ensemble size. Ensemble scores that have been adjusted in this way are also referred to as fair scores. The study documents how well changes in fair scores computed with ensembles with as few as 2 to 8 members can predict the change in CRPS for the operational ensemble size of 50 members. These results could help to accelerate progress in NWP under the constraint of limited computing resources.

Further information can be found in an article published by Martin Leutbecher in the *Quarterly Journal of the Royal Meteorological Society* doi:10.1002/qj.3387.

Coordinated climate simulations using the IFS

Chris Roberts, Sarah Keeley, Franco Molteni, Retish Senan (all ECMWF), Torben Königk (SMHI)

Since 2016, researchers from ECMWF and the EC-Earth consortium have been collaborating with partners from 19 European institutes in the EU-funded PRIMAVERA project to develop a new generation of global climate models. The multi-decadal simulations performed for PRIMAVERA are providing constraints for model development and evaluation that are complementary to those available from short-term (re-)forecast datasets. These insights will be key to improving ECMWF forecasts on medium-range to seasonal timescales.

The core aim of PRIMAVERA is the development and process-based evaluation of a new generation of global climate models, with a focus on the impacts of increased horizontal resolution on simulations and predictions of regional climate. Ocean and atmospheric model resolution can affect many aspects of climate simulations, including climatological biases and the representation of key modes of climate variability, such as the El Niño–Southern Oscillation (ENSO). Understanding such sensitivities in ECMWF's Integrated Forecasting System (IFS) and the EC-Earth model, and their relative importance at different timescales, is crucial for improving the fidelity of the analyses and forecasts produced by ECMWF and the EC-Earth consortium.

The experimental backbone of the PRIMAVERA project is a multi-model ensemble of climate model simulations covering the period 1950 to 2014. Within the project, ECMWF and EC-Earth are providing simulations based on the IFS atmosphere coupled to the NEMO/LIM ocean–sea ice model. The EC-Earth consortium is providing configurations based on IFS Cycle 36r4 (as used in the previous seasonal forecasting system, S4) and researchers at ECMWF are providing configurations based on IFS Cycle 43r1 (as used in the latest seasonal forecasting system, SEAS5). The EC-Earth and ECMWF submissions differ in several ways, including the type of atmospheric grid (linear vs cubic octahedral), coupling strategy (single executable vs dedicated coupler), tuning of model components, ocean and sea ice model versions, and the inclusion of ocean waves. The PRIMAVERA project provides a unique opportunity for systematic comparisons of different versions of the IFS with each other, and with other climate models.

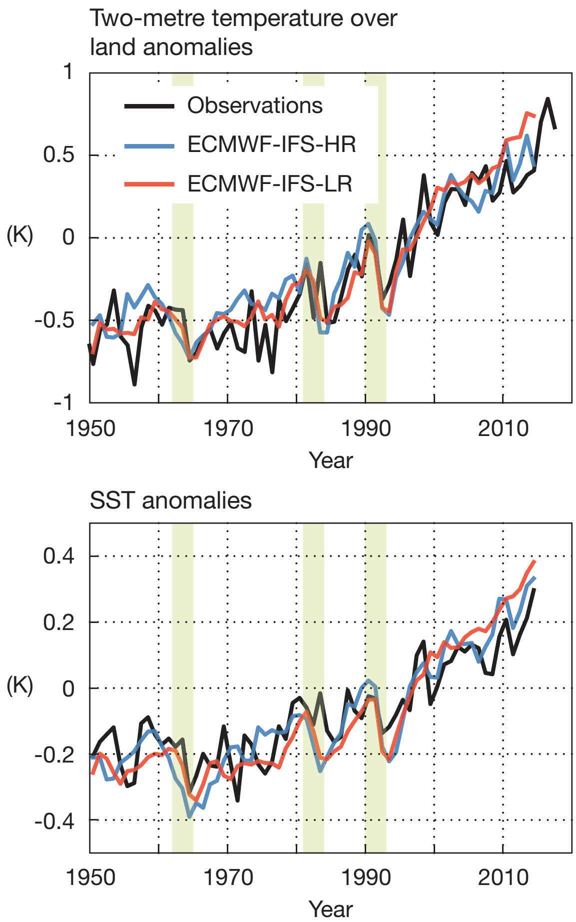

Temperature and SST anomalies in different IFS configurations. Global mean anomalies of 2-metre temperature over land using IFS Cycle 43r1 with a grid spacing of about 25 km (ECMWF-IFS-HR) and 50 km (ECMWF-IFS-LR) (top) and sea-surface temperature (SST) using an ocean model resolution of about 25 km (ECMWF-IFS-HR) and about 100 km (ECMWF-IFS-LR) (bottom), relative to the period 1981–2010. Model data are ensemble means from coupled configurations of ECMWF-IFS. Yellow shading indicates periods corresponding to the Mount Pinatubo (1991), El Chichón (1982), and Mount Agung (1963) volcanic eruptions.

Preliminary results concerning the impacts of increasing ocean and atmosphere resolution in climate experiments with IFS Cycle 43r1 can be summarised as follows:

- All configurations of ECMWF-IFS successfully reproduce the observed long-term trends in global mean surface temperature.
- Increasing the atmospheric resolution from 50 km to 25 km has little impact on climatological surface biases but increases the magnitude of a cold bias in the lower stratosphere.
- Increasing the resolution of the NEMO ocean model from about 100 km to about 25 km substantially reduces biases in North Atlantic sea-surface temperature (SST) and northern hemisphere sea ice extent, but it increases the magnitude of a warm bias in the Southern Ocean.
- Increasing the ocean resolution also improves the simulated magnitude and asymmetry of ENSO variability and improves the representation of associated non-linear SST–radiation feedbacks.
- Ocean coupling and increased atmospheric resolution seem to improve the representation of teleconnections between tropical Pacific rainfall and geopotential height anomalies in the North Atlantic, but the significance of this result needs to be assessed.
- Work is ongoing within the PRIMAVERA project to assess the

impact of increased resolution on climate variability and extremes.

Multi-decadal coupled experiments are not currently performed routinely at ECMWF. However, the asymptotic behaviour of the coupled model is becoming important for numerical weather prediction with the development of coupled approaches to data assimilation and reanalysis. In systems using such coupled approaches, the climatological attractor of the model is important because of its influence as a background field for periods and/or regions with limited observational constraints. The scientific and technical developments required for the ECMWF-IFS configurations used in PRIMAVERA will thus help to evaluate the representation of the more slowly evolving components of the Earth system as part of the model development process.

More details on the results presented here can be found in an article by Chris Roberts et al. in *Geosci. Model Dev.* doi:10.5194/gmd-11-3681-2018.

Can SMOS data support flood and fire forecasts?

Francesca Di Giuseppe, Calum Baugh, Christel Prudhomme, Daniel Thiemert

A new collaboration between ECMWF and ESA will explore the potential of Soil Moisture and Ocean Salinity (SMOS) satellite data to support fire danger and flood forecasts. The 18-month project will start on 1 December 2018.

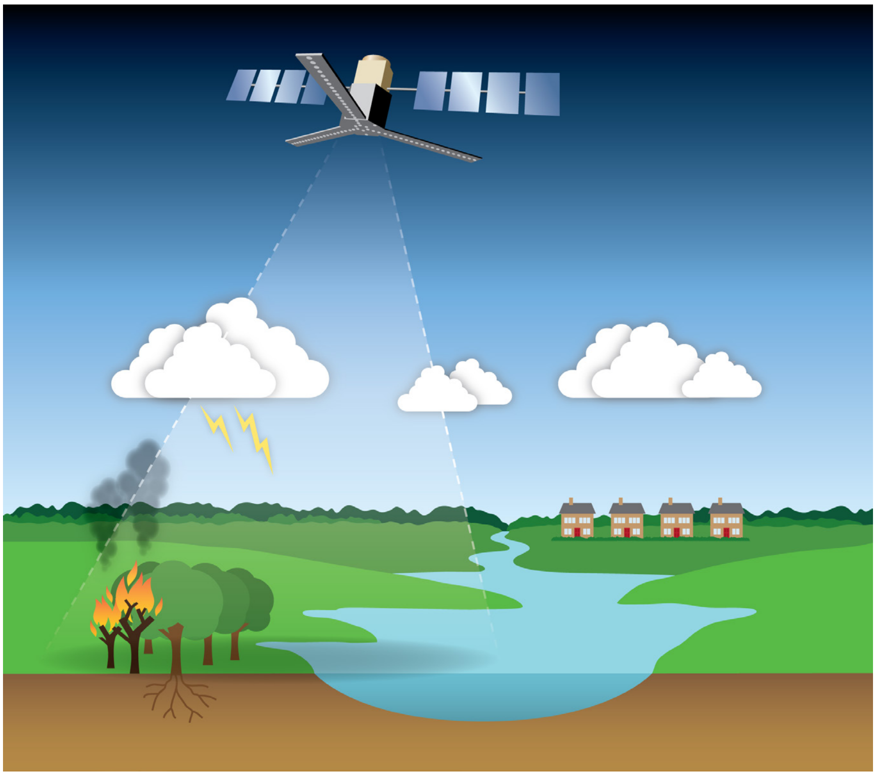

Possible SMOS data applications. The project will investigate how SMOS satellite observations of soil moisture can support fire danger and flood forecasting.

SMOS use in weather prediction

SMOS is the second Earth Explorer Opportunity mission developed as part of the European Space Agency (ESA) Living Planet programme. SMOS was launched on 2 November 2009 and provides two-dimensional interferometric radiometer measurements of L-band (1.4 GHz) brightness temperature from a satellite in polar orbit. At this frequency, the atmosphere is almost transparent, and surface emission is strongly related to soil moisture over continental surfaces and salinity over oceans.

The SMOS mission was launched to improve our understanding of the global water cycle and to contribute to advances in weather and seasonal/ climate prediction. ECMWF has played a major role in achieving those goals by developing methods to integrate SMOS brightness temperature into numerical weather prediction (NWP) models, with a focus on land-surface processes.

New uses for SMOS data?

Beyond the direct use in NWP, a number of scientific studies have shown the potential benefit of integrating SMOS observations into other applications. Since ECMWF contributes to the Copernicus Emergency Management Service as the computational centre for fire danger and flood forecasts, it is in a unique position to evaluate the potential of SMOS data to support applications in these areas.

SMOS data could, for example, help to estimate fuel moisture and thus the amount of biomass available for burning. In addition, soil moisture information could be combined with lightning forecasts to estimate the probability of ignition from lightning. Understanding the influence of fuel moisture on fire emissions could also lead to better predictions of air quality in downwind areas.

Soil moisture is also a key driver of flood risk, with saturated soils increasing susceptibility. River flow forecasts depend on an accurate estimate of initial soil moisture. SMOS data could improve the accuracy of such estimates as used by the European and Global Flood Awareness Systems (EFAS and GloFAS, respectively). Finally, SMOS data could be used to create a new standalone product to highlight areas susceptible to flooding.

ECMWF develops flash flood forecast system

Calum Baugh, Christel Prudhomme

As part of efforts to make weather predictions useful for emergency planning, ECMWF has developed a prototype flash flood forecast system based on the total precipitation Extreme Forecast Index (EFI) product. Results from tests in Europe show that areas susceptible to flash flooding may be identifiable up to five days ahead. The system has the potential to provide global flash flood awareness products for the medium range.

A widespread hazard

Flash flooding is a subset of flooding hazards not currently well captured in flood forecasting systems. It poses an especially significant risk in urban areas and in rapidly responding river catchments. Many countries regard flash flooding as one of the most important natural hazards in their territory. There is no universally agreed definition of what constitutes a flash flood. Meteorologically, flash floods are driven by extreme rainfall intensities which are confined both spatially (between tens and hundreds of square kilometres) and temporally (less than 24 hours). Such situations are often associated with convective activity or orographic enhancement. Hydrologically, the land surface converts most of the precipitation into surface runoff. Flooding then arises either through the rapid accumulation of surface runoff on urbanised or poorly drained surfaces, or the rapid rise of a river within a small, steep catchment.

ECMWF produces global forecasts of weather phenomena that contribute to flash flooding. The aim is to use those forecasts to support existing flood forecast systems, such as the European and the Global Flood Awareness Systems (EFAS and GloFAS) or the World Meteorological Organization Flash Flood Guidance System, which currently does not have continuous global coverage.

Forecast skill

A relevant forecast product is ECMWF's Extreme Forecast Index (EFI) for 24-hour total precipitation. The EFI indicates how extreme the value of a predicted variable is by integrating the difference between the cumulative distribution functions of the forecast and the model reforecast-derived climatology. High EFI values highlight areas liable to receive more extreme precipitation than would normally be expected at that time of year. Since the EFI is calculated relative to a model-derived climatology, it is less susceptible than precipitation forecasts to the underestimation of extreme rainfall totals in ECMWF's ensemble forecasts.

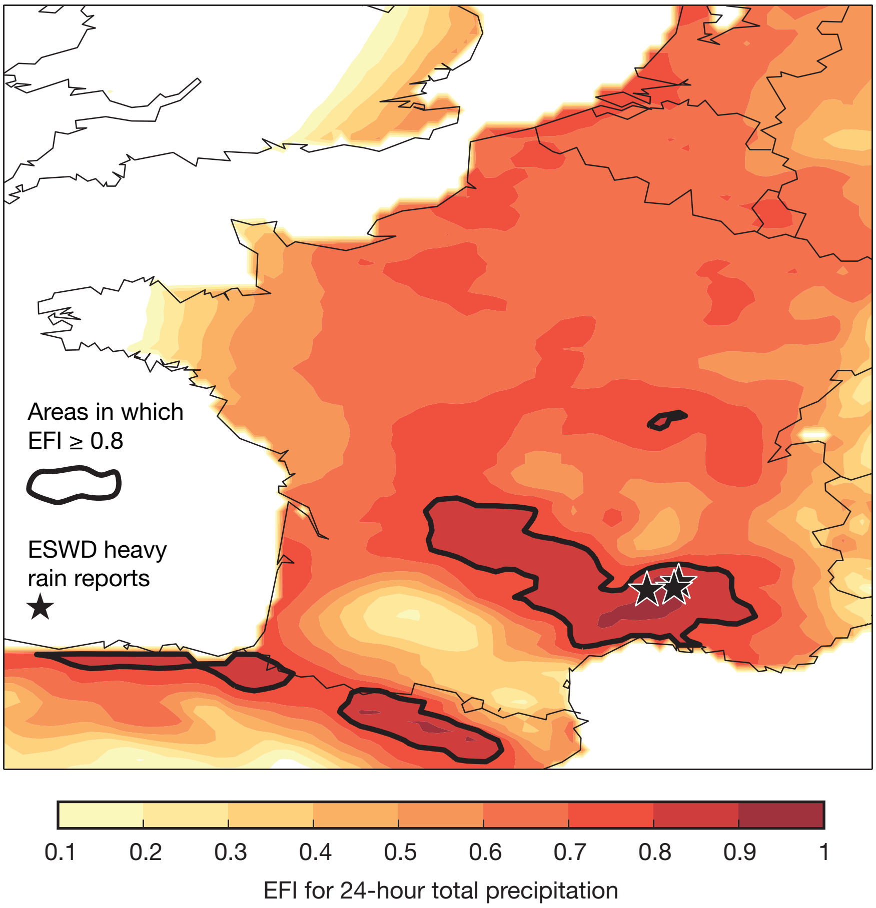

Flash floods in France in August 2018. The chart shows the Extreme Forecast Index (EFI) over land starting at 00 UTC on 7 August 2018 for 24-hour total precipitation on 9 August 2018. The stars indicate locations of observed heavy rain with flash flood impact from the European Severe Weather Database (ESWD).

Using the EFI product, a threshold could be applied to highlight flash flood susceptible areas which exceed this value. For example, on 9 August 2018 flash flooding occurred in the Gard, Ardèche and Drôme departments of southern France. Reported rainfall totals include 105 mm falling in 1 hour in Saint-Martin-d'Ardèche and 167 mm in 24 hours at Méjannes-le-Clap. Plotting the 24-hour total precipitation EFI from 00 UTC on 7 August 2018 for 9 August shows that locations where flooding was reported were in an area of high EFI values. Applying a threshold, such as EFI ≥ 0.8, could highlight broader-scale 'flash flood susceptible areas', which could be very useful for emergency planners.

To explore the range at which such forecasts can be considered skilful, we extracted all 00 UTC forecasts of 24-hour total precipitation EFI for Europe from March 2016 to

March 2017, for lead times up to 5 days. For each forecast at each lead time, different EFI thresholds were applied. Areas exceeding the threshold were compared against 2,663 heavy rain reports from the European Severe Weather Database (ESWD) which mentioned flood impacts. The number of hits, misses, false alarms and correct negatives were then collated over the entire one-year period. Forecast skill was computed as the area underneath the Relative Operating Characteristics curve (aROC). aROC values ranged from 0.78 at 0–24 hours lead time to 0.75 at 96–120 hours lead time. Values greater than 0.5 are considered to be skilful, so these are promising results warranting further analysis.

Outlook

A next step could be to apply the flash flood forecasts globally with a view to integrating them into GloFAS. This will require a repeat of the verification procedure at the global scale using flash flood reports from FloodList.com. The results may make it possible to identify regionally and seasonally specific EFI thresholds that would work best to identify areas at risk. Combining the warnings with exposure information, such as population density and critical infrastructure, could refine the identification of areas where floods would have the greatest impacts.

Further work could consider the added value of using more forecast variables, such as surface runoff and convective activity. The ecPoint rainfall product (see Newsletter No. 153) could be used alongside the EFI to estimate point-rainfall totals within a flash flood susceptible area.

ECMWF assesses pioneering Aeolus wind data

Lars Isaksen, Michael Rennie

ECMWF has begun to process wind profile data from the ground-breaking Aeolus satellite launched on 22 August 2018. First comparisons with ECMWF model data are very encouraging. ECMWF now expects to start assimilating Aeolus data operationally in 2019.

ECMWF is contracted by the European Space Agency (ESA) to produce Aeolus wind products suitable for use in numerical weather prediction. As the Level-2 Meteorological Processing Facility, it will produce wind products and auxiliary meteorological data products for ESA in an operational manner during the mission's lifetime. The Centre started to process Aeolus data within two weeks of the satellite's launch. Just three days after the instrument was switched on, Aeolus was delivering data showing clear features of the wind. Within two weeks, realistic wind profiles as shown in the figure could be produced from the data.

Aeolus is an ESA Earth Explorer satellite mission with an expected lifetime of three years. The polar-orbiting satellite uses ground-breaking laser technology to obtain vertical wind profiles in the troposphere and lower stratosphere. It aims an ultraviolet laser into the atmosphere and detects the Doppler-shift of the backscattered light from both molecules (clear air) and particles (clouds/aerosols). Since very few wind profile measurements are available from other observing systems, Aeolus could bring a significant improvement in the quality of weather forecasts around the globe. The expectation is that the near-real-time wind information it provides will lead to much-improved analyses of the state of the atmosphere, especially in the tropics, with a major impact on forecast quality in the tropics and a smaller but still important impact in the extratropics.

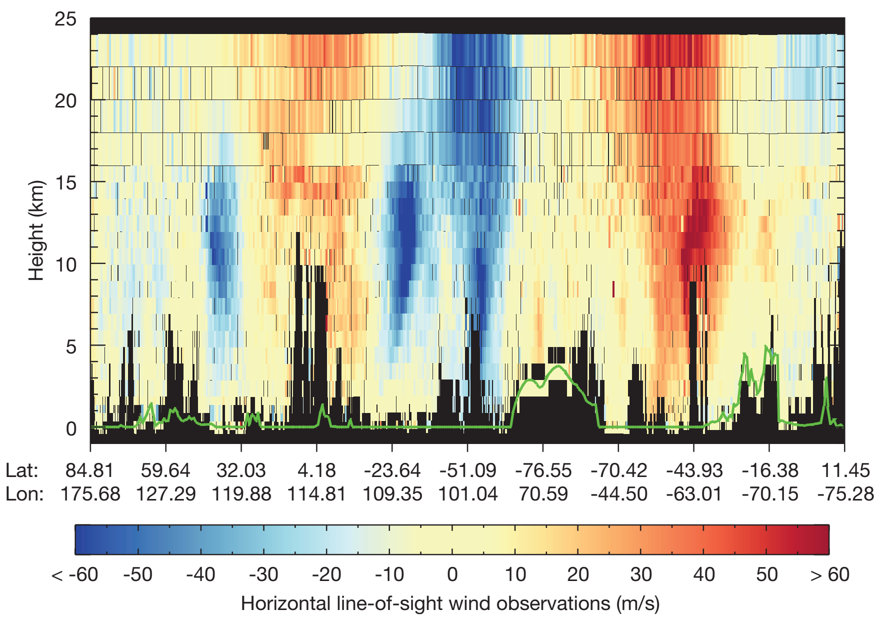

First wind data from the Aeolus satellite. These wind profile data from 4 to 5 September 2018 cover three quarters of an orbit from the Arctic (left) to the tropics and the Antarctic (middle) and back to the tropics (right). There are no data in the black areas either because the ground is elevated (green line) or because of the presence of thick clouds which the Aeolus laser cannot penetrate.

Radiosonde descent reports look promising

Bruce Ingleby (ECMWF), David Edwards (UK Met Office)

The radiosonde descent after balloon burst offers the possibility of an extra atmospheric profile at little or no extra cost. Quality checks suggest that it should be possible to start assimilating some of the descent data at ECMWF over the next couple of years.

The Vaisala RS41 radiosonde software can generate descent reports, currently in BUFR dropsonde format. Descent data from Germany, Finland and the UK for January and June 2018 have been processed and compared with the ECMWF short-range forecast (background). The descent rate is faster than the ascent rate, especially at upper levels. Over 95% of the radiosondes generating ascent reports also generate descent reports. Most of the descents reach 700 hPa, but the numbers available decrease below that level. This is because the signal is lost when the radiosonde goes below the horizon as seen from the launch station.

Encouraging results

The figure shows observation-minus-background (O–B) statistics for June 2018 for German radiosonde ascents/descents, which have the largest sample size and the smoothest O–B statistics. Mostly, the descent statistics look similar to the ascent statistics with two main differences: larger temperature biases at upper levels and smaller vector wind root-mean-square (rms) differences, especially at upper levels. It is known that the ECMWF background has a cold bias at lower stratospheric levels, but it seems that the ascent data here is closer to the truth than the descent data. The Finnish descent data looks even warmer (not shown), possibly due to higher descent speeds (Finnish radiosondes do not use a parachute), but further investigation is required. The descent wind rms differences are smaller than those for ascent winds because the measured descent winds are somewhat smoother. Vaisala apply a digital filter to the winds before producing the meteorological reports, to remove pendulum motion and for smoothing in general. The same time filter is applied to both ascent and descent data. It seems that the filtering should depend on the ascent/descent rate. The ascent data may contain residual noise from imperfect handling of the pendulum motion.

Steps towards operational use

The World Meteorological Organization (WMO) has recently approved a new BUFR sequence specifically for descent data. We hope to see this in use in 2019. For numerical weather prediction (NWP), any biased temperatures are problematic and we need to consider correcting or rejecting the data. As the stratospheric descent data is fairly close in space and time to the ascent, rejecting those temperatures would not be a major loss. Other aspects of processing and quality control would also need attention. Vaisala is changing the RS41 casing, which makes the radiosonde lighter, so assessment should be repeated on the new version. It should be possible to make fairly minor changes relatively soon and assimilate much of the descent data with hopefully positive impact, and at minimal cost. In the longer term, the interface to NWP could be rethought, with NWP using a raw version of the measurements.

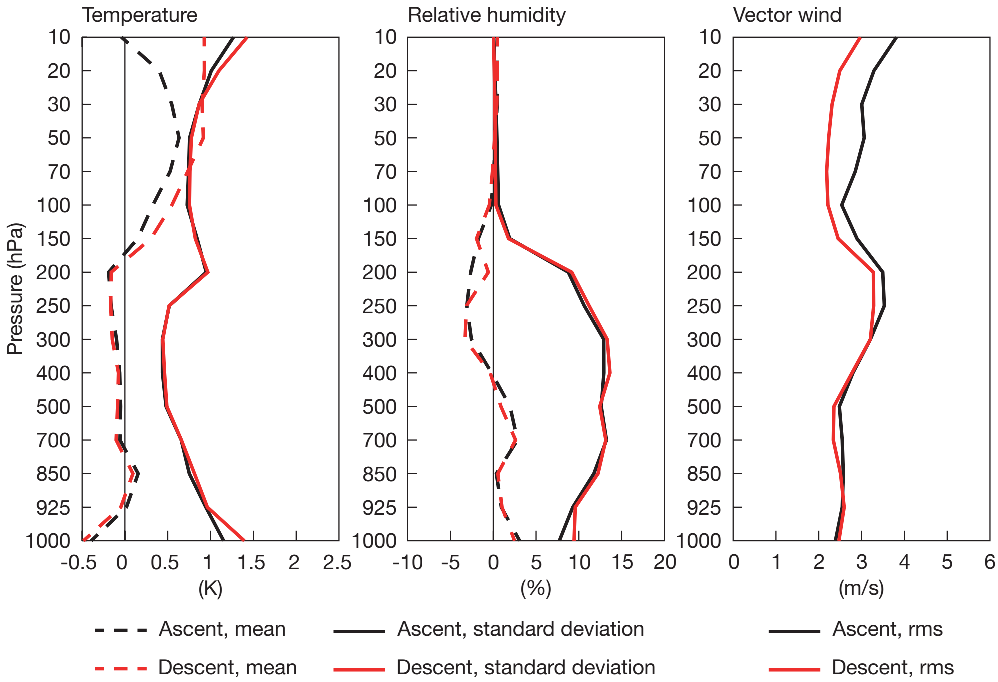

O–B statistics for Germany, June 2018. Temperature and relative humidity mean and standard deviation and vector wind root-mean-square (rms) O–B statistics.

Ship-launched radiosondes plug weather data gap

Rudolf Krockauer (DWD), Cristina Prates (ECMWF)

December 2018 is the last month of the current phase of a European programme to launch radiosondes from ships. Although all details have not yet been agreed by the EUMETNET Assembly, there will be no break or lack of data as of January 2019 as the EUMETNET Automated Shipboard Aerological Programme (E-ASAP) will continue to operate. E-ASAP is a purely operational programme with no or very little R&D. It is part of the global ASAP programme and makes a small but important contribution to the range of data assimilated into numerical weather prediction (NWP) systems, including at ECMWF.

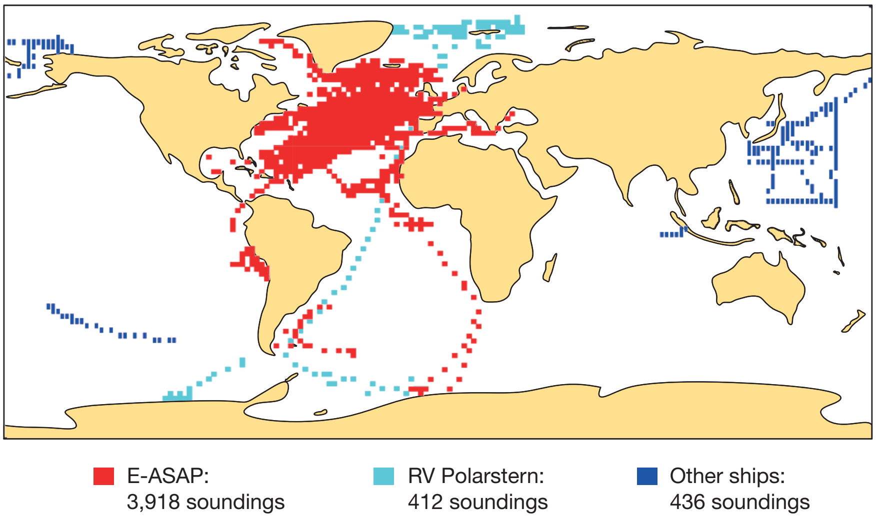

Global ASAP soundings in 2017. E-ASAP makes a very big contribution to the global ASAP programme.

A unique source of data

There is good coverage of radiosonde and ascending/descending aircraft (AMDAR) profiles over Europe and North America, but there are very few soundings over the data-sparse North Atlantic. ASAP radiosondes are the only source of upper-air soundings over the oceans. A look at the global distribution of ASAP soundings illustrates the size of the contribution made by the European E-ASAP fleet. Around 85% of all global ASAP soundings are provided by the 18 ships that make up the E-ASAP fleet. The remaining 15% are mainly provided by two Japanese ships and a German research vessel plus some occasional research campaigns. The unique characteristic of the E-ASAP fleet is that it mostly comprises merchant ships in regular service between Europe and North America. Three main sailing routes over the North Atlantic are reflected in the spatial distribution of the received soundings. The leading role of Europe's contribution to the ASAP programme is recognised by the World Meteorological Organization (WMO).

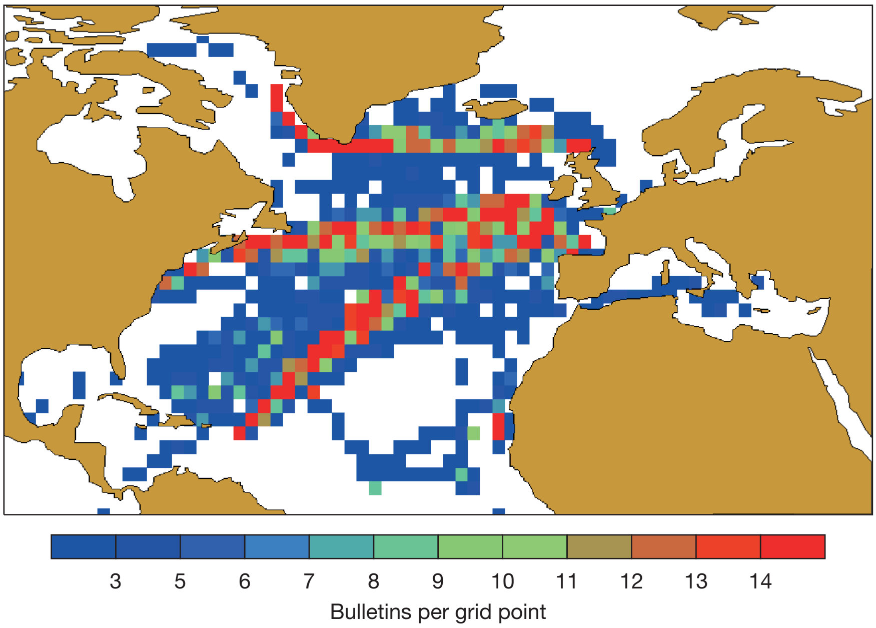

E-ASAP soundings in the North Atlantic in 2017. The sounding frequency plot highlights the three main sailing routes used by E-ASAP ships.

Manual operation

It is important to clear up a common misunderstanding regarding ASAP operations: ASAP stations are not automatic launchers but have to be operated by the ship's crew. Lack of experience, the workload of the sailors, sailing speeds of more than 20 knots, adverse weather conditions etc. produce higher failure rates than for land radiosonde stations. However, E-ASAP soundings have proved to be of high quality and have a positive impact.

The main components of an ASAP station are the balloon launcher and the sounding equipment. The launcher may be a semi-automatic container launcher with a pneumatic hatch or a solely manual deck launcher. Successful soundings depend on the experience of the operators on board. Another difference to land stations is the data transmission from the station to the Global Telecommunication System node. All messages have to be transmitted via satellite communication. The limited satellite communication quality requires small data files of less than 15 kilobytes to overcome communication problems due to low signal strength or interference with other radiation sources on board the ships.

To produce small data files, messages sent in the BUFR data format from the ASAP ships only contain information on measurements made at 10 or 20 second intervals plus at standard and significant levels. The BUFR messages are sent out after termination of the sounding. Prior to this, a small message containing only standard and significant levels is transmitted at 100 hPa. This ensures the transmission of basic profile data in case the transmission of the higher-resolution BUFR message fails.

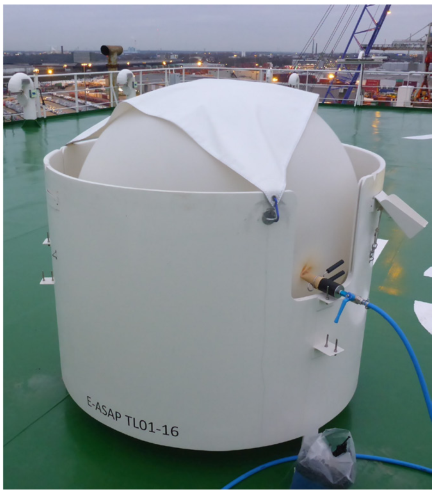

Types of launcher. Some ships carry container launchers (left), others use manual deck launchers (right).

Quality checks

ECMWF routinely performs an NWP-based quality check of E-ASAP soundings. The data assimilation system provides the analysis and background fields of temperature, humidity, wind and geopotential that are used as independent estimates against which observations are compared. These daily monitoring activities are important to check if the quality of the assimilated observations from the different vessels has deteriorated or if the quality of any blacklisted ones (observations from ships not used because their quality is systematically poor) has improved. Currently, most of the observations from E-ASAP reports are used in the ECMWF data assimilation system. In addition, ECMWF produces an annual report on the ASAP quality control monitoring to the Joint WMO-IOC Technical Commission for Oceanography and Marine Meteorology (JCOMM) Ship Observation Team (SOT). Overall, the quality of E-ASAP observations has improved considerably over the last few years to reach quality standards that are fully comparable to land-based radiosondes. The occurrence of reports giving wrong locations – an intermittent problem in the past – is almost non-existent now.

Outlook

E-ASAP started as a three-year pilot project in 1999 before it became a EUMETNET core programme in 2003. Like all EUMETNET programmes, E-ASAP is agreed for periods of five to six years. The current programme phase will end on 31 December 2018. The next phase, 2019 to 2023, has already been agreed by the EUMETNET Assembly. The German national meteorological service (DWD) will continue to manage E-ASAP in the next phase. In addition to technical challenges, such as improving satellite communications, there are also managerial challenges, such as implementing the requirements of the WMO Integrated Global Observing System (WIGOS). Crucially, the future of E-ASAP also depends on the spirit and motivation of the people on board since autolaunchers are not well suited to the harsh conditions on board seagoing ships.

Rudolf Krockauer is the E-ASAP Programme Manager.

New Director of Computing takes up his post

Dr Martin Palkovič took up his position as ECMWF's Director of Computing on 1 October 2018. Martin was previously the Vice President of Engineering at Codasip Ltd in Brno, Czech Republic, where he was responsible for the management of R&D as well as collaboration with academic communities and national and EU representatives.

Martin was previously Managing Director of IT4Innovations, where he established and managed the Czech National Supercomputing Centre and was responsible for building and operating the national HPC e-infrastructure and related R&D activities.

A Slovakian national, Martin holds a PhD in electrical engineering from the Technische Universiteit Eindhoven in the Netherlands and has authored or co-authored more than 50 international publications in the field of ICT. He has served on several advisory boards and is currently a member of the EU Horizon 2020 Future and Emerging Technologies Advisory Group.

Martin Palkovič. (*Photo: Sznapka Petr*)

Copernicus services hold general assemblies

Jean-Noël Thépaut, Vincent-Henri Peuch

The Copernicus Climate Change Service (C3S) and the Copernicus Atmosphere Monitoring Service (CAMS), which are both implemented by ECMWF on behalf of the European Commission, held their General Assemblies in September and October 2018.

C3S

C3S's General Assembly took place at the German Federal Ministry of Transport and Digital Infrastructure, Berlin, from 24 to 28 September. Participants reviewed a particularly exciting year for C3S and looked ahead to a new era.

Since 2015, C3S has moved on from a proof-of-concept phase to become a fully operational service that supplies climate information to tens of thousands of users, ranging from policy-makers and businesses to scientists and citizens. After consolidating links with users and data providers and developing initial service elements, such as the provision of monthly climate bulletins and sectoral climate impact indicators, 2018 saw the launch of the C3S Climate Data Store (CDS). The CDS is a gateway for accessing climate data and data processing tools, freely available and open for everyone to use.

Several sessions of the General Assembly focused on the CDS, including its data content and tools and how to ensure and improve service quality. Other sessions covered topics such as the Sectoral Information Systems, which aim to provide information to help specific sectors deal with climate change (see the separate article in this Newsletter). The meeting also included a training session for C3S scientists and contractors. The mix of plenary talks and thematic fairs enabled maximum interaction between C3S contributors and provided a comprehensive view of the progress achieved so far. Several smaller meetings dedicated to specific technical topics were organised with C3S contractors, who are instrumental in implementing C3S by providing data and creating a wide range of applications.

C3S General Assembly plenary session. The event took place at the German Federal Ministry of Transport and Digital Infrastructure in Berlin.

The General Assembly enabled C3S data providers and users to share their knowledge and experience of the various components of the service. It gathered user requirements and feedback on the service so far and, perhaps most importantly, it facilitated networking and brainstorming for the climate change community in the context of an operational C3S. Looking to the future, C3S is planning to integrate even more datasets into the CDS, which will cover a wider variety of climate variables and indicators.

CAMS

Participants in the CAMS General Assembly, which took place in Lisbon from 16 to 18 October, were also able to review a string of exciting developments over the past year. The Assembly was opened by Prof. Miguel Miranda, president of IPMA and President of ECMWF's Council. The agenda included updates on all aspects of the Service, from input data to user uptake and communications activities. The two highlights of the past year – the successful inclusion of Sentinel-5P observations in the CAMS global system (see the separate article in this Newsletter) and the completion of the CAMS global reanalysis of atmospheric composition – of course featured highly in the programme.

One of the days was dedicated to presentations and discussions with users and potential users of CAMS, with a special focus on Portuguese public and commercial entities. A poster and demonstration session enabled further exchanges between users and providers, with opportunities to showcase and discuss their activities related to CAMS. Participants also heard about a range of major developments planned for 2019 and beyond, in particular the ADS (Atmosphere Data Store), which will be built as a new instance of the CDS to serve all CAMS products.

C3S poster session. There were plenty of opportunities for networking and informal discussions at the C3S General Assembly in Berlin.

Sentinel-5P air quality data look promising

Antje Inness

Earlier this year, the first data from the new Earth observation satellite Sentinel-5P were released, providing a wealth of information about air quality. The results of quality checks carried out at ECMWF on ozone and nitrogen dioxide data from Sentinel-5P are very promising.

Sentinel-5P – the precursor to Sentinel-5 – is the first satellite mission in Europe's Copernicus Earth observation programme to be dedicated to monitoring atmospheric composition. Data from the mission will help to address global issues such as air quality, climate change and the ozone layer. Information from Sentinel-5P will be used for air quality forecasting by the Copernicus Atmosphere Monitoring Service (CAMS), which is implemented by ECMWF on behalf of the European Commission. Processed data products will better equip Europe to address some of the challenges affecting citizens in their daily lives.

Sentinel-5P hosts the state-of-the-art TROPOMI (Tropospheric Monitoring Instrument) spectrometer, which maps trace gases that affect our health and the climate. TROPOMI can also identify the location of significant emissions and their impact on atmospheric concentrations of key pollutants. This will help to better mitigate air quality problems and to provide better air quality forecasts.

Data evaluation

After its launch in October 2017, Sentinel-5P data were evaluated and instruments were calibrated during a nine-month commissioning phase. Early data were compared with the CAMS global forecasting system, which is based on ECMWF's Integrated Forecasting System (IFS). This made it possible to detect and solve teething problems.

Near-real-time (NRT) TROPOMI total column ozone and nitrogen dioxide retrievals were included passively in the CAMS system in July 2018 on the day they were officially released by the European Space Agency (ESA). Since then the data have been routinely monitored by ECMWF and plots can be found on the CAMS website (https://atmosphere.copernicus.eu/). 'Monitoring' in this case means that differences between the forecast model and the TROPOMI observations are calculated, allowing us to assess the quality of the data. However, at this early stage the data are not actively assimilated, hence they do not yet influence the CAMS forecasts.

Early impressions for ozone and nitrogen dioxide are that the data look very promising, and assimilation tests with the data are now beginning in parallel to the routine monitoring. If these tests are successful, assimilation of the data will be activated in the CAMS NRT system in the near future. Other retrievals from TROPOMI, including carbon monoxide, sulphur dioxide, formaldehyde and methane, will also be included in the CAMS system once the data become available later this year.

Outlook

We expect that Sentinel-5P will soon become one of the most important data sources underpinning the quality of CAMS information products. If all goes well, Sentinel-5P data will soon feed into daily global and European air-quality forecasts, ensuring that CAMS will be able to provide users with an improved reliable and quality-assured service.

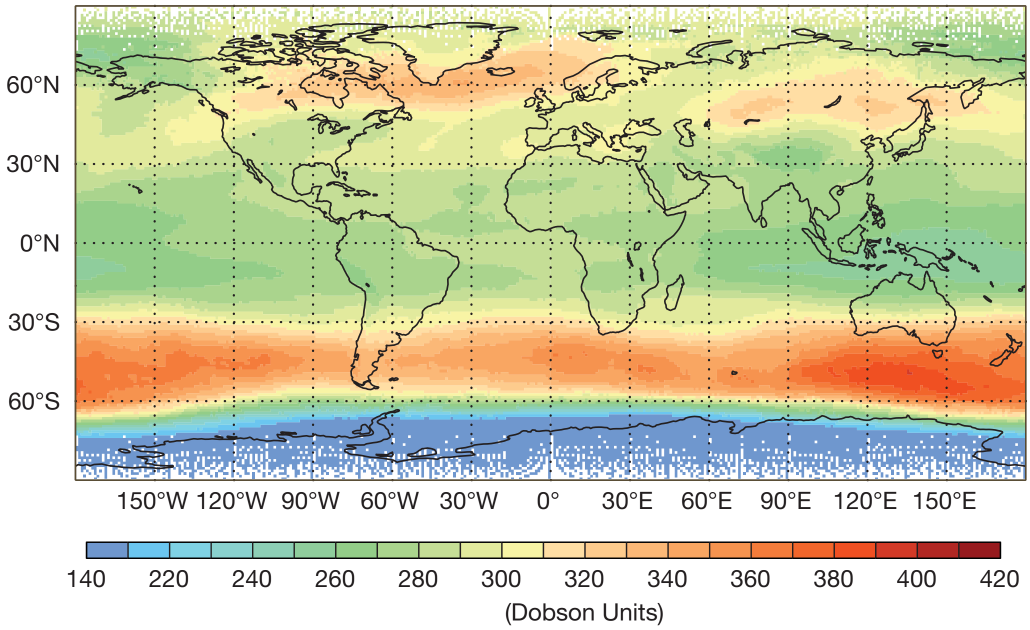

Sentinel-5P ozone data. Total column ozone retrieval from TROPOMI, averaged over the period 25 August to 24 September 2018. (*Source: Copernicus Atmosphere Monitoring Service, ECMWF*)

Sectoral Copernicus climate applications take off

Carlo Buontempo, Samuel Almond

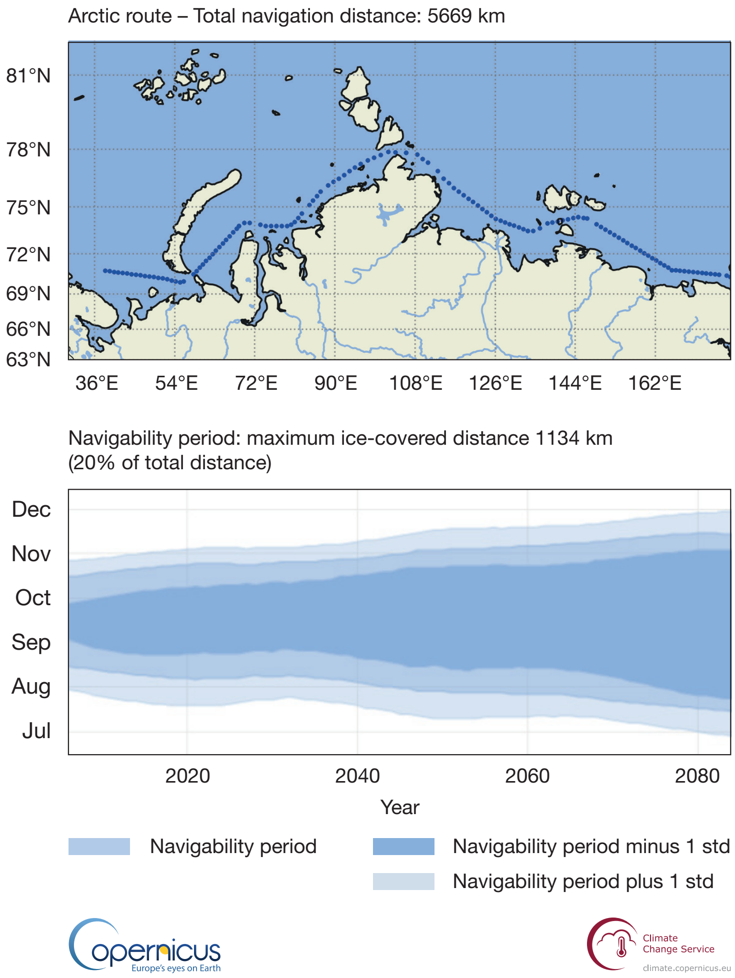

Charts from an Arctic route availability application. The dotted line in the top chart represents the standard Northeast Passage route. The plot below indicates the times of year when no more than 20% of the total navigation distance is projected to be covered by ice over the next few decades (medium blue shading), plus or minus one standard deviation (light and dark blue shading). (*Charts produced by B-Open for C3S*)

The EU-funded Copernicus Climate Change Service (C3S) implemented by ECMWF has made rapid progress in transforming the data it provides into applications tailored to different economic sectors. After a successful proof-of-concept phase, the Sectoral Information System (SIS) component of C3S has developed a range of demonstrator applications which are now becoming operational. SIS transforms the data contained in the Climate Data Store (CDS) and tailors it to the requirements of users in sectors such as water management, agriculture & forestry, tourism, insurance, energy, health, coastal management, transport, infrastructure, disaster risk reduction and biodiversity.

In 2015, ECMWF signed seven contracts with European research institutes, R&D centres and commercial companies to develop proof-of-concept climate applications addressing user requirements in five sectors: water, energy, agriculture, insurance and managing the urban environment. The applications were developed based on external datasets provided, for example, by the Earth System Grid Federation (ESGF), ECMWF's MARS archive and the European Climate Assessment and Dataset project (E-OBS dataset). These demonstrators have proved invaluable: they have shown that there is an appetite for trusted climate information and the provision of a sustainable climate service that can support downstream climate applications, impact assessments and policy development. The POC applications were completed prior to the launch of the CDS (June 2018) but have now been migrated to the CDS infrastructure.

Building on the experience gained, a series of European and global SIS applications, addressing user and policy requirements in ocean transport, coastal management, agriculture, tourism and marine fisheries, will be completed in spring 2019. Unlike the proof-of-concept applications, they are being developed directly on the CDS infrastructure. Like the migrated demonstrators, they will provide examples of how the data and tools available through the CDS can be used in specific user-relevant contexts.

Outlook

One of the top priorities for the next phase, up to 2020, is to increase the quality and robustness of the procedures that underpin the SIS applications (methodology, workflows, metadata etc.), the documentation (FAQs, technical documents, training resources, etc.) and user support (service desk, response to feedback, application evolution). This will be supported by an Evaluation and Quality Control activity that will independently assess the fitness for purpose of all SIS output, including demonstrators, against best scientific practice, computational efficiency and user requirements.

The SIS applications are key to ensuring that trusted climate information provided by C3S is disseminated to end users across multiple sectors, enabling them to make informed decisions to mitigate or adapt to the effects of climate change. They will serve as examples of how climate applications can be developed utilising the data, technologies and tools provided by the CDS and its toolbox. They will thus help intermediaries (service providers positioned between C3S and end users) to develop further climate applications to meet societal and market needs.

ECMWF's Summer of Weather Code

Julia Wagemann, Claudia Vitolo, Anna Ghelli

At the start of 2018, ECMWF launched the first edition of the ECMWF Summer of Weather Code (ESoWC) to find novel and creative ways to address weather-related software challenges. Applications opened in January 2018. ECMWF proposed 13 challenges aimed at streamlining and supporting daily activities at the Centre. External developers were asked to submit a proposal addressing a challenge of their choice. More than 170 people applied from all over the world and five teams were selected. Their open-source projects were released in September. They should prove useful to ECMWF and its Member and Co-operating States.

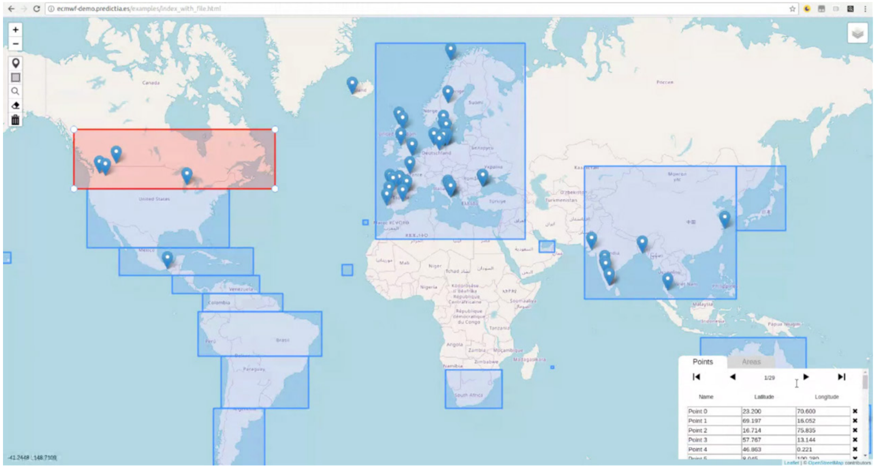

GIS widget. The team behind the GIS widget project demonstrated the functionalities of the widget based on Leaflet during the week of webinars.

Outcomes

The five selected teams, supported by ECMWF mentors, started coding in May. By August they had developed:

- a geographic information system (GIS) widget, based on the Leaflet JavaScript library, which simplifies ECMWF's data extraction workflow by selecting and displaying bounding boxes and points, performing geo-searches and inputting data on a map
- the migration of calibration software to downscale rainfall products to the Python programming language, including the development of a related graphical user interface (GUI)
- a web crawler to facilitate and automate a web-based search for new environmental data sources
- innovative visualisation techniques which make it possible to display the wealth of information contained in ensemble forecasts in an easy-to-understand and catchy way for the general public
- a virtual file system to represent NetCDF files to help users browse and manipulate these often complex file structures.

Mentors and developers had regular virtual meetings to discuss progress and potential issues. The programme ended with a week of webinars from 17 to 21 September 2018, during which the results of the challenges were presented. The proof-of-concept projects are open to contributions and can be further tailored to better support ECMWF's internal activities as well as those of ECMWF's Member and Co-operating States.

Benefits

The enhanced GIS widget is useful for many web applications across ECMWF, for example the real-time requirements-editing interface and the online costing tool, where users calculate the cost of their data and configure the requirements. Any other online applications where users need to specify latitude/longitude locations or bounding boxes can also benefit from this widget.

The migration of the software ecPoint-PyCal to Python provides a dynamic environment while the development of a GUI makes the software more user friendly. The software compares numerical model outputs against point observations to identify biases/errors at local scale.

The web crawler for environmental data obviates the need for a time-consuming manual search for new data sources. The software automates the discovery, analysis and assessment of candidate web pages. The resulting data can ultimately be used to improve global weather forecasting models.

The new visualisation design makes ensemble forecasts more accessible to a non-expert audience. This innovative visualisation is based on value-suppressing uncertainty icons. A larger range of icons is allocated to forecasts with low uncertainty and a smaller range when uncertainty is high.

The representation of NetCDF files as a virtual file system is useful for those working with weather and climate data in NetCDF format who would like to quickly explore and edit a dataset. The software is written in Python and allows users to easily mount, view and edit the contents of a NetCDF dataset using file-system operations and general-purpose Unix tools.

Useful links

More details, including links to the open-source project codes, are available on Julia Wagemann's science blog post on ESoWC: *www.ecmwf.int/en/about/media-centre/science-blog/2018/ecmwfs-summer-weather-code*

Recordings of the webinars are available on YouTube: *http://bit.ly/esowc18-webinars*

Re-architecture of Atlassian collaboration tools

Daniel Varela, Manuel Martins

The Atlassian suite is a family of web applications used at ECMWF to facilitate collaboration, both internally and with our Member and Co-operating States. An exponential increase in its use has led ECMWF to reconfigure the way in which the Centre provides these tools.

Growth in use

The suite includes Jira, a project and issue management system; Jira Service Desk, an incident and support system; Confluence, a wiki and collaboration platform; Bitbucket, a git server and code collaboration system; and finally, Bamboo, a continuous integration platform. These tools are all provided by the same vendor, Atlassian, and integrate with each other seamlessly. For example, it is possible to access Confluence knowledge base articles from Jira Service Desk, or to link together a Jira bug report with the code that was committed to Bitbucket to fix that bug and to access the results of the unit tests in Bamboo.

In 2010, ECMWF started using the Atlassian tools to help improve the way we supported our software packages. The initial user base was rapidly extended to ECMWF users beyond the software development and support groups, and from there to many external collaborators and partners across many projects.

As illustrated by the charts, over the last eight years we have seen exponential growth in the usage and amount of content stored in Confluence and Jira.

New architecture

The original architecture, which involved running all services in a single VMware virtual machine, was not designed for such an intense mission-critical usage. We needed to migrate to a ‘container’ solution where each tool is isolated into its own virtual system for increased performance, ease of recovery and future scalability. The migration to our new data centre in Bologna was also a concern. The new architecture provides an infrastructure/environment for running containerised applications on premises. This ensures that we can support the tools appropriately.

The container technology chosen for the Applications Group web deployment is Docker (https://www.docker.com). Docker enables us to package and version all Atlassian releases independently. An image for each version of each Atlassian product is packaged with all of ECMWF’s custom configurations and libraries and pushed to our internal Docker registry. By doing this we ensure consistency, guaranteeing that what is released and deployed in any of our stacks, pre-production and production, is the same.

Deployment

The current production versions of Jira and Confluence have already been migrated to new systems that are accessible at https://jira.ecmwf.int and https://confluence.ecmwf.int. However, we are aware that the old URLs (https://software.ecmwf.int/issues and https://software.ecmwf.int/wiki) are still referred to from many previous documents and websites, so we have provided an automatic redirection. The migration has not yet involved upgrading the software versions of Jira and Confluence, an exercise that will be performed separately once the new production infrastructure proves to be stable. The new upgrade procedure will involve equivalent automated redeployments in the development and pre-production environments, to be pushed to production only once unit and user tests have been passed.

The final steps in the re-architecture project will be to migrate our git server Bitbucket and finally Bamboo, the continuous integration system. We are confident that the new infrastructure will enable us to handle at least a few more years of increased usage of the Atlassian tools while maintaining the stability and good performance currently enjoyed by our users.

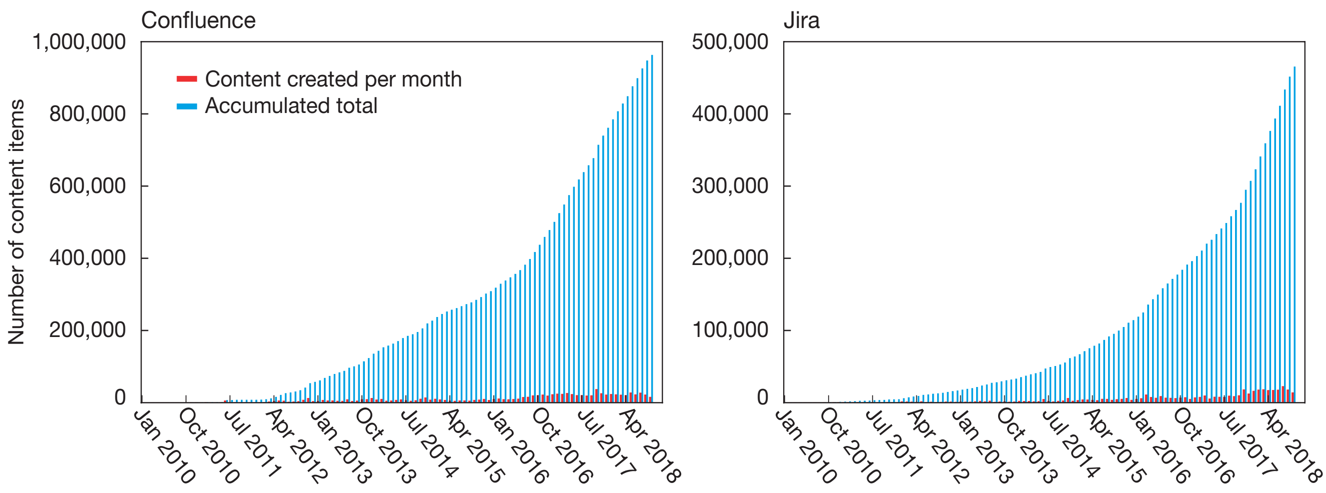

Content growth. Total accumulated content and content growth per month for Confluence and Jira.

Roberto Buizza talks about probabilistic forecasting and life after ECMWF

ECMWF Lead Scientist Roberto Buizza is leaving the Centre after 27 years to take up a position as Full Professor of Physics at the Scuola Superiore Sant'Anna in Pisa. Roberto has been one of the key players in the development of ensemble forecasting at the Centre and in promoting the use of a probabilistic approach to manage weather risk in a range of sectors. He has also worked on ocean reanalysis and coupled data assimilation, and he was the coordinator of the EU-funded ERA-CLIM2 project.

Roberto in Pisa in November 2017.

In the eighties, you completed a physics degree with a thesis on plasma physics. How did you come to move from there into numerical weather prediction?

R.B. After I completed my degree, the Centre for Thermonuclear Research of the Electricity Board of Italy (CRTN/ENEL) offered me a position as a scientist, to work on the diffusion of pollution, meteorology and climate. I accepted and it turned out to be a great choice: while working for CRTN, I visited ECMWF to extract data that we were using to assess the impact of weather on the energy sector. At CRTN, we also worked with a limited-area model, and I started to face the predictability questions that would then keep me busy for the rest of my career.

What were the main questions in numerical weather prediction when you joined the Centre in 1991?

R.B. One of the key questions we were facing was how to estimate forecast confidence. We knew that forecast failure was flow dependent, and we were wondering whether we could design an operational method that would give us a flow-dependent, reliable and accurate estimate of the confidence we can have in a forecast. Tim Palmer, Franco Molteni, Čedo Branković, Ernst Klinker and Robert Mureau were working on the development of a prototype ensemble which could be used to give us such an estimate. Joe Tribbia from NCAR was visiting ECMWF for six months. I worked mainly with Franco and Joe to implement a method to initialise an ensemble of forecasts. We decided to use the singular vector (SV) approach. The method was known and was used in simple barotropic or quasi-geostrophic models, but no-one had ever computed singular vectors with a complex primitive equation model with a few thousand degrees of freedom. I remember inspirational discussions, long days debugging the code, a great team spirit, a unique atmosphere – and then the successful implementation of the ECMWF ensemble at the end of 1992. Today, ensembles are used to provide more complete, reliable and accurate forecasts for all forecast ranges, from a few hours to months ahead, in practically all the world's national meteorological centres. Ensembles are also used to estimate the uncertainties of analyses.

As you leave the Centre 27 years later, what are the main challenges in predictability today?

R.B. I think there are two key challenges from the scientific point of view. The first is how to improve our simulation of model uncertainties. Currently our ensembles are under-dispersive for some key variables, e.g. near-surface variables, and the Ensemble of Data Assimilations underestimates uncertainties in the analysis. The second is that, although we are now using a coupled (atmosphere, land, three-dimensional ocean, ocean waves and sea-ice) model to generate all our forecasts, we still estimate the initial conditions (ICs) in uncoupled/weakly coupled ways. This makes both our analyses and the ICs for forecast ensembles suboptimal. From the user point of view, the key challenge that we are facing is to convince users that it is always more valuable to take decisions using probabilistic rather than single forecasts. Forecasters use ensembles to assess the range of possible future weather scenarios, but my feeling is that they are still not exploiting their full potential. End users are still mainly asking for categorical yes/no information and are often unable to use probabilities in decision-making. We need to work more closely with them and develop tools and products so that they can exploit the full range of ensemble-based information we generate daily.

What do you say to the charge that ensemble forecasts are a cop-out: they are almost never wrong because almost any outcome is liable to be predicted, if only with a small probability?

R.B. Ensembles can be wrong, and they are definitely not a 'cop-out'! Ensembles must be evaluated using metrics designed for probabilistic forecasts. Two key properties that these metrics assess are reliability and accuracy. In a reliable ensemble, a forecast probability of, say, 75% that a weather event will occur means the event should occur 75% of the time when such a forecast is issued. If one verifies such an ensemble forecast for a large enough sample of cases and finds that instead of 75% this event happens only 60% of the time, we can say that the ensemble is unreliable. Another property is accuracy, i.e. how close the ensemble forecast probability distribution is to observations. Using these metrics, we can make

doi: 10.21957/r206bvq3y7

quantitative assessments of ensemble forecast skill: poor ensembles will score low using these metrics, while good-quality ensembles will score high.

Today the rationale for ensemble forecasting is widely accepted. Nevertheless, at ECMWF we still produce a higher-resolution analysis and forecast, as well as ensembles. Is there still a case for a higher-resolution analysis and forecast?

R.B. We need to distinguish between the role of a single higher-resolution analysis and the role of a single higher-resolution forecast. In the ECMWF operational suite, the single high-resolution analysis provides the unperturbed state to which we add perturbations generated using the SVs and the Ensemble of Data Assimilations (EDA). We need this analysis since the 51 members of the EDA cannot be completed in time to generate the ensemble initial conditions: we do not have enough computing power to do this. Thus, until we can generate the 51 EDA analyses in time to initialise ensemble forecasts (ENS), the current setup with one analysis run at higher resolution (compared to the EDA and ENS) pays off.

Regarding the role of the single high-resolution forecast (HRES), the key question to ask ourselves is: which among the forecasts ECMWF produces is on average more accurate: the HRES, the ensemble mean, or the ensemble median? If one compares the skill of these forecasts, it turns out that on average HRES is the most accurate only in the short forecast range (the exact forecast time depends on the variable considered). So HRES provides valuable additional information in the short forecast range, up to about forecast day two. What is more, running at least one member at higher resolution than the ensemble is valuable because it helps us and our users to prepare for the future: in line with our Strategy, the horizontal resolution of our ensemble forecasts will at some point increase to about 9 km, and the experience gained with the 9 km HRES since 2016 will help us prepare for that. At the next resolution increase, ideally we should set the resolution of the HRES to what we will be able to afford for ENS at the following resolution upgrade. In other words, we should strengthen the link between the resolutions used in ENS and HRES to guarantee that ENS inherits a well-tested model version.

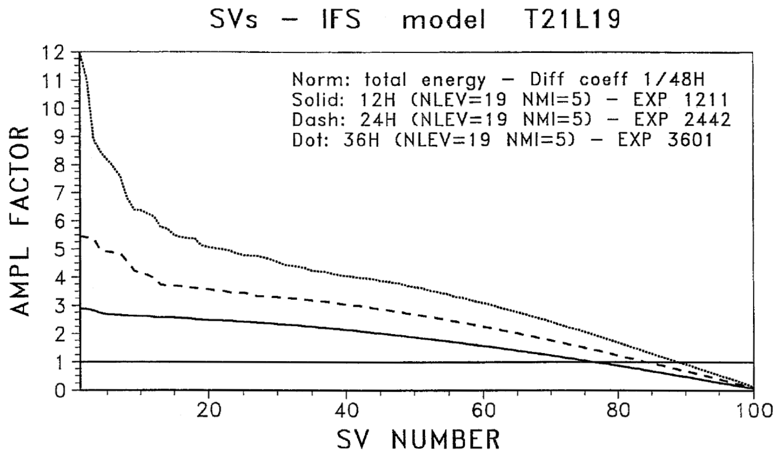

One of the first singular vector results obtained in 1992. The figure shows the amplification factor of 'well-behaved' leading singular vectors computed with ECMWF's Integrated Forecasting System (IFS) at T21L19 resolution with different optimisation times.

Beyond ensemble forecasting, what are the key areas in weather prediction where we can expect to see progress in the next 10 years or so?

R.B. As stated in our Strategy, one of the key aspects that ECMWF aims to understand is why it is so difficult to predict regime transitions over Europe in the forecast range from two to four weeks. Europe is also a difficult region for seasonal prediction. We aim to improve the skill of European forecasts of regime transitions in the sub-seasonal range, and of large-scale anomalies in the seasonal range. Recent upgrades to our model that have brought coupling to a three-dimensional ocean and a dynamical sea-ice model in all our forecasts were steps in the right direction. We need to continue to investigate which aspects we need to improve to increase forecast skill. For example, is there something missing in our treatment of stratospheric processes and in the way we propagate signals between the stratosphere and the troposphere? Do we need a finer vertical resolution to propagate these signals more realistically? Do we need a higher-resolution ocean capable of resolving the Gulf Stream better, thus improving the development and propagation of atmospheric systems over the Atlantic Ocean? How can we further improve the quality of the initial conditions exploiting all available observations, for example the ones coming from the newly launched Aeolus satellite?

Looking to the future, what will be the focus of your work at the Scuola Superiore Sant'Anna?

R.B. Scuola Superiore Sant'Anna of Pisa, together with Scuola Normale of Pisa and Scuola Superiore IUSS of Pavia, have agreed to start a new inter-university centre that will study the impact of climate change on society and sustainability. I will be working with experts in many different fields, from economics and management to agriculture and health care, from mathematicians and physicists to bioengineers and lawyers. Climate change is an extremely complex problem, which is having a huge impact on society: we want to work on some of the most urgent climate change-related problems using interdisciplinary approaches, and propose possible solutions and ways forward.

Let me close this interview by thanking everyone with whom I have worked during the past 27 years for their openness, kindness, ideas, constructive criticism, and challenging discussions. In particular, let me say a huge thank you to Franco Molteni, one of my closest friends and colleagues: Franco contacted me at the end of 1990 informing me that there was a vacancy in Tim Palmer's group, asking whether I would be interested in applying. This is how it all started. I also want to thank Tim, Adrian Simmons and Dave Burridge, who in 1991 decided to offer that position to me, thus giving me a great opportunity. ECMWF is a great, unique place, and I wish you all the best for the future.

Addressing biases in near-surface forecasts

Thomas Haiden, Irina Sandu, Gianpaolo Balsamo, Gabriele Arduini, Anton Beljaars

ECMWF's medium-range forecasts of near-surface weather parameters, such as 2-metre temperature and 10-metre wind speed, have become more skilful over the years, alongside improvements in upper-air scores. There are, however, persistent biases in these forecasts which have proved difficult to eliminate. In-depth investigations carried out at the Centre show that these biases are closely related to the coupling between the atmosphere and the land surface in the Integrated Forecasting System (IFS).

The biases are also related to other processes, such as turbulent mixing, radiation and clouds. In some cases, the representation of these other processes leads to errors which partially cancel some of the errors that can be attributed to the atmosphere–land coupling. See Box A for more details on such 'compensating errors'. A deeper understanding of the underlying causes is necessary to address biases in near-surface weather parameters in a way that ensures increased physical realism and reduces compensating errors. Because of atmosphere–surface feedback mechanisms, an improved representation of surface fluxes may also lead to an increase in medium- and extended-range predictive skill overall.

The investigations presented in this article are part of an ECMWF initiative entitled 'Understanding uncertainties in surface–atmosphere exchange' (USURF), which started in November 2017. USURF addresses the very useful feedback about near-surface issues which ECMWF receives from forecasters in the Member and Co-operating States via the Technical Advisory Committee (TAC), Using ECMWF's Forecasts (UEF) meetings, the 'Green Book' on verification, the 'Known Forecast Issues' page, and various other interactions with forecasters and customers. Key to making progress was the development of a conditional verification methodology, which helped to identify specific processes as likely causes of some of the biases. Work so far has focused on 2-metre temperature (T2m) biases in Europe. However, because of the physical links between T2m, 2-metre humidity and 10-metre wind speed, investigations have also included some aspects of humidity- and wind-related processes.

Biases can be removed statistically to some extent when forecasts are passed on to end users. However, there is considerable value to forecast users in having less biased direct model output as it provides a physically more consistent representation of the atmospheric state. Furthermore, a substantial fraction of the 'random' error may result from state-dependent systematic errors which need to be addressed in the forecast model itself.

a Compensating errors

Increasing the physical realism of surface processes in a model to reduce systematic biases may increase the root-mean-square error (RMSE) because different kinds of errors may no longer partially cancel each other. For example, 2-metre temperature (T2m) is computed diagnostically in the IFS from the temperature at the lowest model level and the skin temperature. There is a limiter in the computation of T2m which becomes active in very stable, low wind situations, and which prevents the T2m from deviating too strongly from the temperature at the lowest model level. Removing this limiter would be physically desirable, but tests have shown that doing so in the current model setup increases the RMSE. This is because errors introduced by the limiter and cloudiness errors partially cancel each other: if the limiter is removed, negative temperature errors at night are increased in cases where the forecast underestimates cloudiness.

In more general terms, trying to increase the realism in one process can leave the model more exposed to errors in other processes. Another example is the strength of the thermal coupling between the surface (the vegetation canopy) and the uppermost soil layer. Decreasing this coupling allows the surface to cool more strongly and produce stronger surface inversions, more in line with observations. However, it also makes T2m in the model more reactive and increases errors in cases where cloudiness is underestimated. The solution to the problem lies in identifying and properly attributing errors in all contributing processes, and then reducing these errors at the same time.

doi: 10.21957/eng71d53th

Two-metre temperature biases in Europe

Routine verification against SYNOP weather station observations shows that T2m biases in the IFS have diurnal and annual cycles (Figure 1) and a pronounced regional dependence. In winter, for example, there is a night-time cold bias of 0.5–1 K in large parts of Europe, and a warm bias of several K throughout the day in parts of Scandinavia (Figure 2). In summer, there is a general underestimation of the amplitude of the diurnal cycle of temperature and a daytime low-humidity bias. Over recent years, there have been some changes in these biases due to model changes, but they are relatively robust in terms of geographical patterns and annual and diurnal variations.

The results shown in Figures 1 and 2 are based on a subset of SYNOP stations. They include only those locations where the model orography differs by no more than 100 m from the actual one, and where the nearest model grid point is a land point. This excludes most stations in mountain areas, specifically those on peaks and in small valleys, and many coastal stations. The purpose of this filtering is to focus the verification on larger-scale bias patterns and on areas where the IFS can be expected to represent near-surface weather parameters reasonably well given the limitations imposed by grid resolution.

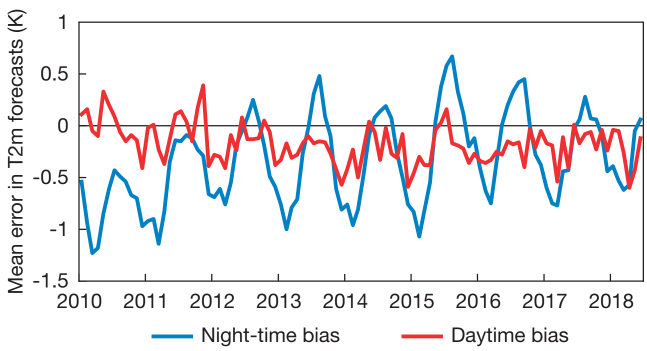

FIGURE 1 Night-time (00 UTC) and daytime (12 UTC) bias of the ECMWF high-resolution T2m forecast in Europe from verification against SYNOP observations. Stations for which the model elevation differs by more than 100 m from the true elevation, and stations where the nearest grid point is an ocean point have not been included. Lead times are 60 hours and 72 hours, respectively.

Night-time cold bias

Conditional verification can be used to quantify relationships between errors in different variables and to disentangle their sources. For example, if we only consider T2m forecasts for days which are (nearly) clear-sky, both in the forecast and SYNOP observations, then the wintertime night-time T2m bias in central Europe is negligible. This suggests that cloudiness plays a role in this bias. In addition to stratifying T2m forecasts according to a quantity like cloudiness, one can also stratify T2m errors according to the forecast error for cloudiness. The left-hand panel of Figure 3 shows that the night-time negative temperature bias in the IFS in central Europe in winter increases roughly linearly with the amount by which total cloud cover is underestimated (against SYNOP observations). However, when weighted by the frequency distribution of the cloud cover errors (shown as green bars in the plot), it turns out that cases where the total cloud cover is underestimated and cases where it is nearly correct contribute about equally to the negative T2m bias (Figure 3b). This indicates that the wintertime negative total cloud cover bias in the IFS over central Europe (on the order of 10% against SYNOP observations, not shown) does not fully explain the negative night-time T2m bias in this region. In cases when the total cloud cover is correctly predicted, the negative T2m bias could be due to other cloud errors, e.g. an

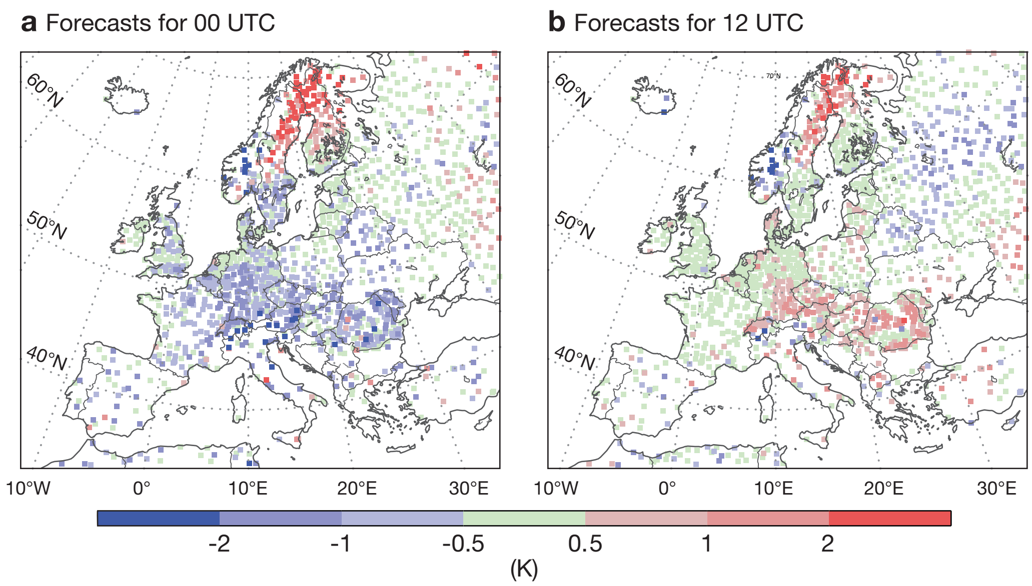

FIGURE 2 Mean error (bias) of the T2m forecast for day 3 in winter 2017/18 (December–January–February) at (a) 00 UTC and (b) 12 UTC. Verification was performed against the same subset of SYNOP observations as in Figure 1.

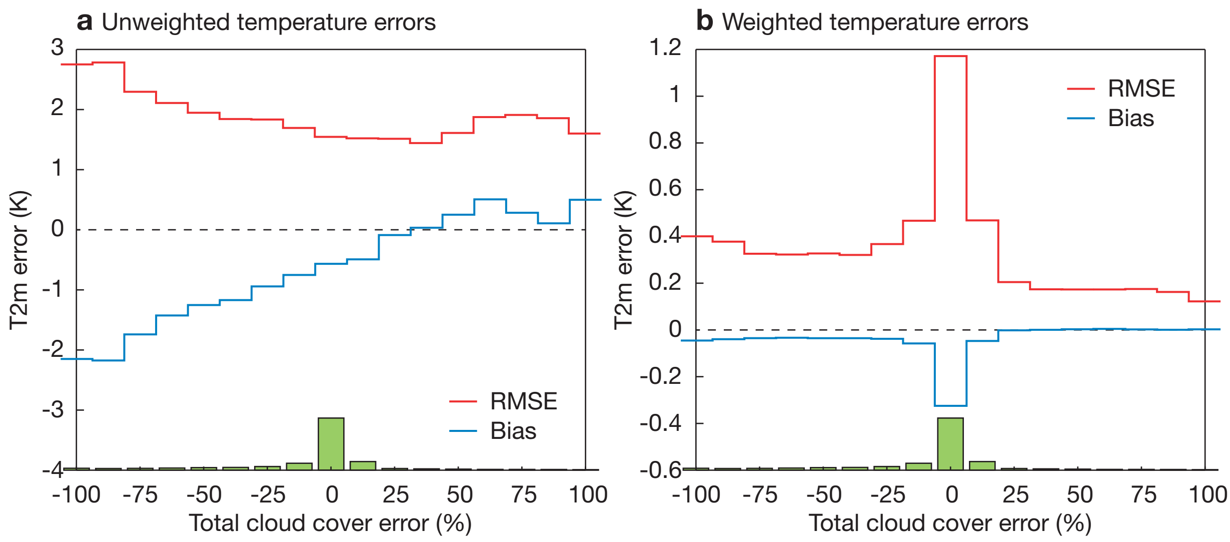

FIGURE 3 Root mean square error (RMSE) and mean error (bias) for T2m forecasts valid at 00 UTC as a function of the total cloud cover (TCC) error for December–January–February 2016/17 in a central European domain (48–55°N, 0–15°E) at a lead time of 12 hours (a) averaged for each TCC error bin and (b) averaged for each TCC error bin and weighted by the TCC error relative frequency of occurrence. Green bars show the TCC error frequency distribution (arbitrary vertical scale).

underestimation of cloud optical depth, erroneous cloud type or erroneous cloud base height. It could also be due to errors in processes not directly related to clouds, such as vertical mixing or coupling with the surface.

To gain further insight into the cloudiness errors, we have also verified downward solar radiation at the surface against both SYNOP observations and a satellite product (Figure 4). Both show an overestimation of downward solar radiation in the order of 5–10 W/m^2 during wintertime, which corresponds to a relative bias of about 5–10%. These daytime results are consistent with the total cloud cover underestimation against SYNOP observations. The fact that three independent observational datasets indicate very similar forecast biases means that we can have relatively high confidence in these results.

In order to distinguish between different types of cloudy situations, total cloud cover errors can also be stratified according to cloud top height derived from satellite data. Figure 5 shows that a large contribution to the negative cloud cover bias comes from low clouds. Since the satellite identifies the top of the uppermost cloud layer only, the full frequency distribution of cloud top height will be shifted towards lower levels. In central Europe, low stratus with cloud tops typically below 2 km is known to be the main contributor to the negative bias in total cloud cover (*Haiden & Trentmann*, 2016). Due to recent model upgrades this bias has, however, been

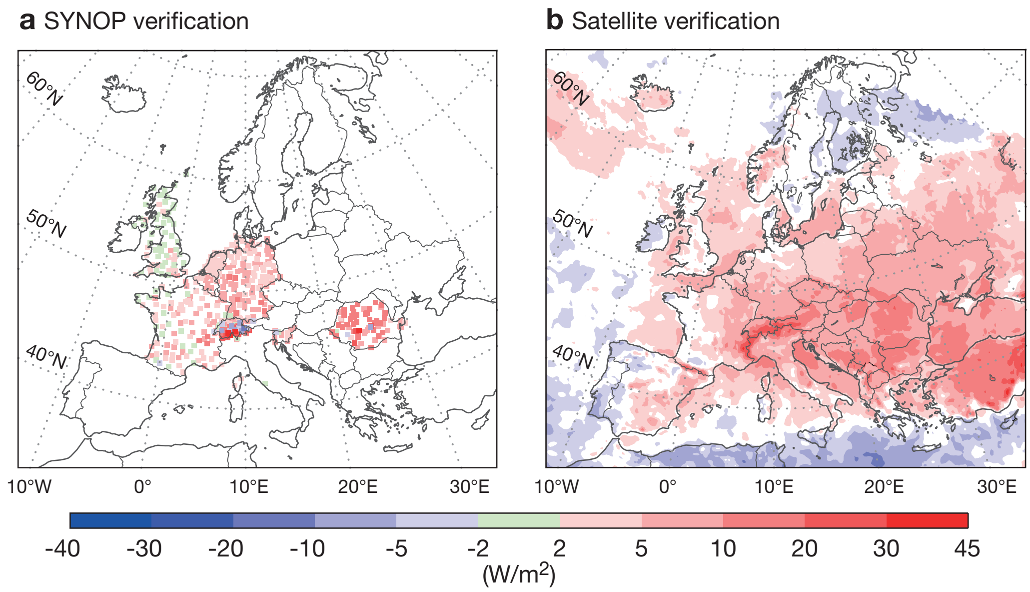

FIGURE 4 Bias in downward surface solar radiation (24-hour averages) at forecast day 2 in November–December–January 2017/18 from (a) verification against SYNOP and (b) verification against the corresponding satellite product from the Climate Monitoring Satellite Application Facility (CM SAF).

reduced significantly and cloud forecast skill has increased accordingly.

Warm bias in Scandinavia

Cloudiness errors are also a factor contributing to the wintertime positive T2m bias in parts of Scandinavia, although the sign of the errors is different: for reasons that are not fully understood, in this region cloudiness tends to be overestimated rather than underestimated. Another factor appears to be the representation of snow cover in the IFS. In clear, calm nights in the real atmosphere, the uppermost layers of the snow cool rapidly, and a strong vertical temperature gradient is established within the snowpack. The skin temperature of the snow drops substantially and T2m decreases accordingly. The single-layer snowpack used operationally in the IFS at present reacts more slowly, due to its larger thermal inertia. The result is a delay in the drop in skin temperature. This delay cannot be fully compensated for by reducing the coupling between the skin temperature and the snow layer because that would lead to an overestimation of the daytime warming of the snow surface. Preliminary results from tests performed with an experimental multi-layer snow scheme show substantial improvement in the form of 2–3 K stronger night-time T2m drops when conditions are undisturbed (Figure 6). However, some adverse effects on daytime T2m (increase of the warm bias) have been noted which require further study.

Summertime biases

The main systematic T2m forecast error in summer is an underestimation of the diurnal cycle by about 1–2 K, with a cold bias during the day and a warm bias during the night. Cloudiness errors do not appear to play a major role in this case. Another factor that influences interactions between the surface and the atmosphere is surface type, such as soil, vegetation, snow, water or ice. The sub-grid heterogeneity underneath the atmospheric model grid-point is represented by means of a mosaic or tiling that makes it possible to solve the energy balance for each component separately. This is necessary because each land/water element has different properties. Key land surface-related factors controlling the near-surface temperature are soil moisture and soil temperature in a warm climate and snow density and snow/ice temperature in a cold climate. Over land, forests, grassland, and bare soil interact differently with the atmosphere. The interaction also depends on local physiographic conditions, which influence surface drag, aerodynamic resistance, and canopy resistance to evapotranspiration.

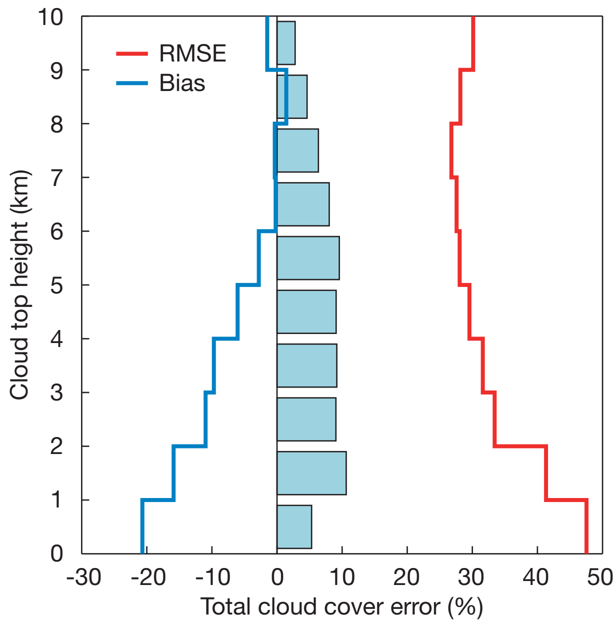

FIGURE 5 Dependence of short-range wintertime total cloud cover RMSE and mean error (bias) on cloud top height. The bars show the frequency distribution of cloud top heights. The results shown are for central and eastern Europe in December–January–February 2016/17.

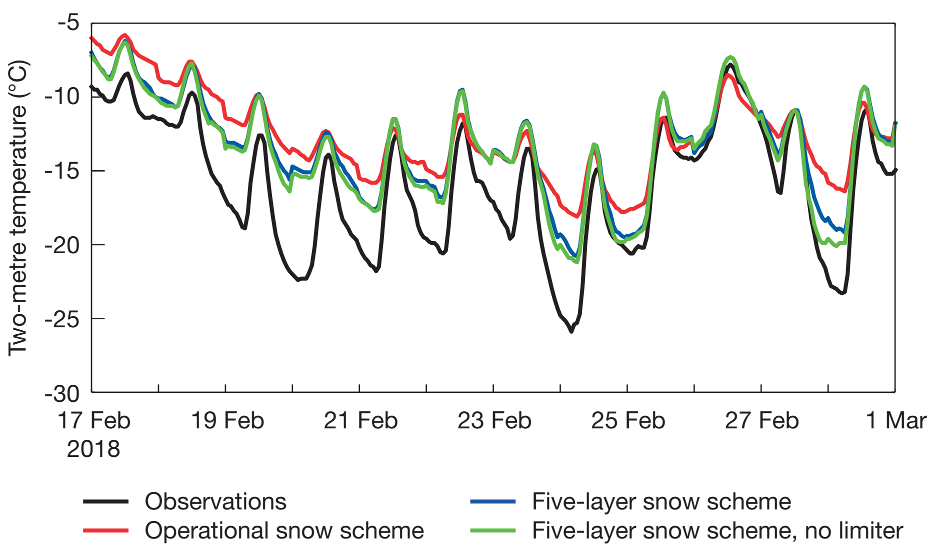

FIGURE 6 Observed and predicted T2m averaged over northern Scandinavia (64–70°N, 15–30°E). The forecasts are a control experiment with the operational snow scheme and two multi-layer snow scheme experiments in which a five-layer vertical discretization is used. In one of them the T2m limiter is active as in the operational model, in the other the T2m limiter is deactivated. Verification is against SYNOP stations for the period 17 February – 1 March 2018.

Looking first at the night-time warm bias in summer, it is worth noting that night-time cooling in the model is sensitive to the parametrized strength of the thermal coupling between the vegetation canopy and the soil. A reduction of this coupling reduces the night-time bias in summer, but it also leads to stronger night-time cooling in winter, potentially increasing negative T2m biases in central Europe. In IFS Cycle 43r3, implemented in July 2017, the coupling was slightly reduced, which led to a reduction in the night-time warm bias in summer. Reducing it further would have degraded the winter T2m. The night-time warm bias in summer can be reduced further only if the cloud bias in winter can be improved. This example illustrates how a combination of simultaneous changes can be required to reduce near-surface biases without worsening forecast scores.

Turning next to daytime biases in summer, daytime T2m and humidity forecasts for Europe have an overall cold and dry bias. This could be an issue just in the surface layer (up to about 50–100 m) or it could be due to problems in a deeper layer of the atmosphere or both. Figure 7 shows that the model underestimates the difference in temperature and humidity across the surface layer compared to radiosonde observations. This underestimation is particularly pronounced at lower latitudes and contributes to the negative biases there. It means that part of the daytime cool/low humidity bias in summer is likely due to the surface layer in the model being too strongly mixed.

A method that is complementary to conditional verification is to investigate the sensitivity of modelled near-surface weather parameters to changes in boundary-layer physics. Such experiments indicate the extent to which observed biases can be attributed to specific parametrization choices in the turbulent mixing in the boundary layer. Figure 8 shows that summertime T2m is quite insensitive to major changes in the mixing profile, whereas 2-metre dew point does show some sensitivity. Increased vertical mixing reduces the 2-metre dew point by about 0.5 K. Combined with insights from conditional verification of the dew point for clear-sky and cloudy cases, these experiments suggest that the overall low humidity bias during summertime is partly associated with an overestimation of mixing in cloudy cases associated with summertime shallow or deep convection. If only clear-sky cases are evaluated, the forecast has a slight moist bias.

Summary of findings

Near-surface weather parameters such as T2m are governed by a range of processes, such as advection, boundary-layer turbulent mixing, the strength of the land–atmosphere coupling, radiation fluxes, the state of the soil and vegetation, and the presence of snow or orography. The large number of factors involved complicates forecast error attribution. Significant progress has recently been made by using conditional verification and by running sensitivity experiments to

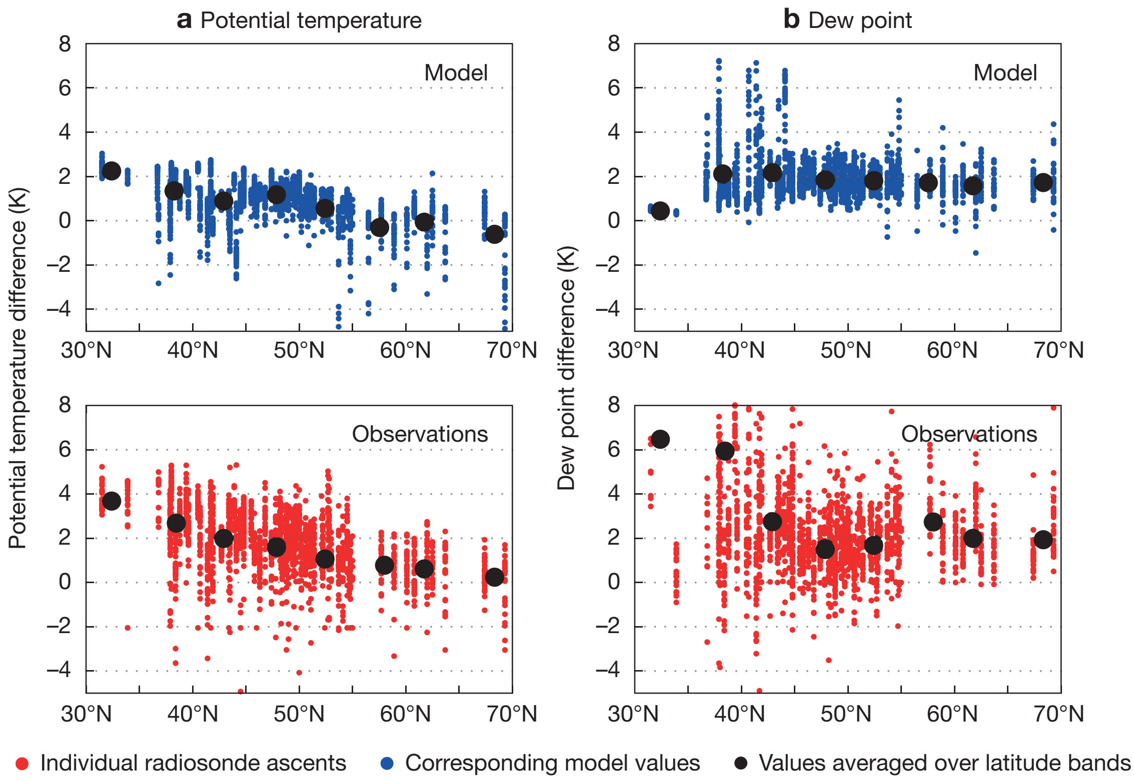

FIGURE 7 Differences in (a) potential temperature and (b) dew point between the heights of 2 m and 200 m above ground in the ECMWF model and in radiosonde observations as a function of latitude in Europe (10°W–28°E). The verification period is July 2015.

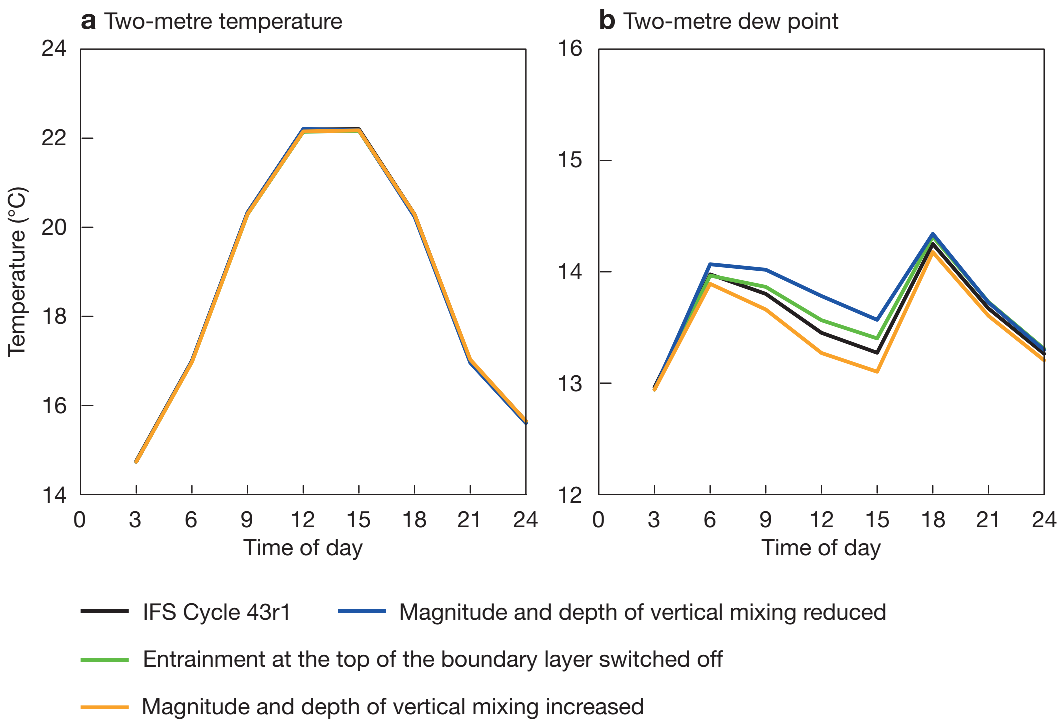

FIGURE 8 Effect of different degrees of mixing across the daytime planetary boundary layer on (a) T2m forecasts and (b) 2-metre dew point forecasts in a central European domain as a function of the time of day. Results for increased strengths of turbulent vertical mixing show negligible sensitivity for temperature but a noticeable effect on dew point. Black lines correspond to IFS Cycle 43r1, operational from November 2016 to July 2017. The forecasts are short-range forecasts at 25 km resolution, aggregated over the month of July 2016.

explore the impact of parametrization changes on near-surface parameters. The main findings of the ongoing ECMWF project focusing on these biases are that (a) biases are easier to address if one focuses first on non-coastal stations outside major mountain areas; (b) the night-time cold bias in most of Europe in winter is partly related to an underestimation of the cloud cover, but some of it is present even when the cloud cover is correct; (c) the warm bias in Scandinavia in winter is partly due to the use of a single layer in the snow scheme; (d) the underestimation of near-surface temperature and humidity in summer over land is at least partly due to an insufficient temperature and dew-point gradient in the lowest 200 m; (e) daytime near-surface temperature in the model is resilient to changes in atmospheric mixing, while humidity is moderately sensitive to atmospheric mixing; and (f) the low humidity bias in summer appears to be mostly related to an overestimation of turbulent mixing in cloudy boundary layers.

Further work

One of the next steps will be to perform a more in-depth verification against datasets from meteorological masts, such as the Lindenberg site mast (run by the German national meteorological service, DWD), which is now available to ECMWF in near-real time, and the Cabauw mast (run by the Dutch national meteorological service, KNMI). This will show to what extent biases in near-surface temperature and dew point are representative of biases over a deeper layer and how this changes with the time of day and with season. It will also make it possible to concurrently examine errors in temperature in the lowest 100 m of the atmosphere and in the soil, as well as errors in the surface energy budget. It is hoped that this will help to further pin down the cause of biases in the operational forecast. The reasons for the different kinds of cloud errors found in forecasts for Scandinavia and central Europe will also be investigated, notably using data from the Sodankylä mast in Finland.

Work towards the operational implementation of a multi-layer snow scheme will continue. The scheme will be calibrated on in-situ measurements within the ESM-snowMIP experiment to optimise the underlying parameters so that the observed snow depth and density are reproduced. Evaluation of the scheme will particularly focus on its ability to reproduce the observed near-surface temperature amplitude of the diurnal cycle.

A simpler framework for moist processes is being developed, relying on consistent assumptions and improved coupling between the turbulent diffusion, convection and cloud schemes and the dynamics. This work, together with planned improvements to the representation of warm- and cold-phase microphysical processes, should help to further reduce systematic errors in cloudiness and precipitation and thereby reduce biases in near-surface weather parameters.

Further reading

Haiden, T. & **J. Trentmann**, 2016: Verification of cloudiness and radiation forecasts in the greater Alpine region. *Meteorol. Z.*, **25**, 3–15.

Progress in using single precision in the IFS

Peter Düben, Michail Diamantakis, Simon Lang, Sami Saarinen, Irina Sandu, Nils Wedi, Tomas Wilhelmsson

Research carried out at the University of Oxford, Météo-France and ECMWF has shown that it is possible to significantly reduce the arithmetic precision of many of the calculations performed in numerical weather prediction models without compromising the quality of weather forecasts. 'Single precision' forecasts have the advantage of being computationally less expensive than traditional 'double precision' forecasts. Such efficiency savings will greatly facilitate the introduction of higher-resolution ensemble forecasts and other model improvements in line with ECMWF's Strategy to 2025.

A lot of work has gone into enabling the use of single precision in ECMWF's Integrated Forecasting System (IFS), with the result that the quality of single precision forecasts at the operational resolution is now comparable to that of double precision forecasts. The main remaining difference between single and double precision forecasts is a larger mass conservation error when single precision is used. The reasons for this have been found to be complex, but the error can be mitigated by using a global mass fixer. At ECMWF, single precision simulations have already proved useful to reduce computational cost significantly in research experiments. Work to prepare for the use of single precision operationally in the IFS is under way.

What is single precision?

When a weather forecast model runs on a supercomputer, physical parameters are represented as strings of bits that can be either 0 or 1. The precision at which a number can be represented depends on the number of bits that is used per variable. In the IFS, the default number of bits per real number has been 64 for the last few decades. This level of precision is called 'double precision'. It makes it possible to represent real numbers to a precision of at least 15 significant decimal digits. Numbers as large as 10^{308} and as small as 10^{-308} can be represented. In single precision, the number of bits to represent real numbers is reduced to 32. Precision for real numbers is reduced to seven significant decimal digits, with a number range between 10^{-38} and 10^{38}. In general, the use of single precision instead of double precision speeds up simulations since less work needs to be done by the supercomputer. For uncoupled IFS simulations, this leads to a reduction in computing time of approximately 40%.

In the future, we aim to run standard forecasts with the IFS in single precision to improve computational efficiency while in principle keeping double precision for 4D-Var data assimilation. This includes forecasts for research purposes but also operational weather forecasts produced on ECMWF's next high-performance computing facility. See Box A for details on how single precision has been implemented in the IFS.

Why is single precision faster?

There are four reasons why single precision simulations are faster compared to double precision simulations:

1. Since data volume is reduced, more data can be stored closer to the processing unit (in memory and cache), and less waiting time and costly data transport is required.

2. The processing unit can perform more operations, with a speed-up by a factor of up to two. However, the size of any benefit depends on the extent to which the code is vectorised. Vectorisation is a style of computer programming in which operations are applied simultaneously to whole arrays instead of individual elements, with the number of parallel operations increased by a factor of two for single precision. Whether a significant ratio of the code is vectorised depends heavily on the compiler used.

3. Future supercomputers will use more and more processing units in parallel for a single simulation. The amount of information that needs to be shared between processors represents one of the most important bottlenecks for simulations. If single precision is used instead of double precision, the data volume that needs to be communicated between processors and compute nodes is halved.

4. For very large simulations, load balancing between compute nodes can be improved if overall data volume is reduced thanks to the use of single precision.

The IFS is rather complex with many different components that can have a very different computational footprint. This makes it difficult to carry out a reliable performance analysis. It is therefore not possible to make general statements on which of the four reasons listed above are the most important. Speed-ups depend on the hardware used; the model resolution; the MPI/OpenMP configuration and the number of processing units; the blocking of the code (using 'nproma'); the compiler; and

doi: 10.21957/ps2y9gfa2d

a Implementation of single precision in the IFS

In Fortran, KIND values define the number of bytes used to represent real numbers: 4 bytes = 32 bits correspond to single precision, 8 bytes = 64 bits correspond to double precision. KIND values are specified when real numbers are initialised at the beginning of programs, sub-routines or modules. In the IFS, precision is adjusted using a few global integer variables that define the KIND values for large groups of real numbers. There are three integer variables that are used to define the precision of most real number variables in the IFS:

- JPRB: This is the working precision that can be either double or single precision depending on the settings selected by the user.
- JPRM: These variables are always initialised in single precision.
- JPRD: These variables are always initialised in double precision.

JPRB is used for the overwhelming number of real numbers throughout the IFS model code. To change a simulation from double to single precision is, in principle, as easy as switching JPRB from 8 to 4 and changing some compiler options to define the use of single precision as default precision for variables that are initialised with no explicit specification of the KIND value. Starting from IFS Cycle 45r2 (a non-operational, technical cycle), it has been possible to choose the numerical precision and to start single precision simulations straight from prepIFS (using the 'Numerical precision' tab, under which users can pick a default precision and a precision for the 'FC' standard forecast job) with no need for any changes of the IFS branch. Jobs in single precision will automatically switch on the mass fixer.

If a local area in the code shows problematic behaviour if single precision is used, local variables can easily be upgraded from JPRB to JPRD to restore double precision locally. However, things become more complicated if relevant parameters are shared between subroutines, since this requires that the precision of information that is sent fits the precision of information that is received.

To make a special rule for single or double precision within the IFS code, an IF statement can be used to check whether JPRB is equal to JPRD. JPRB is equal to JPRD for double precision simulations and different for single precision simulations. This is, for example, useful if subroutines from libraries (such as LAPACK or BLAS) that are precision dependent are linked. The use of single precision to read GRIB input files or to write to GRIB output files, as well as MPI communication, is handled via interface blocks that pick the correct precision level automatically. The use of single precision will not change the precision of the GRIB data that is used for model I/O and data storage.

other factors. It is possible for some model components to speed up by more than a factor of two if single precision is used (for example if expensive data operations suddenly fit into the limited size of the cache which stores data very close to the processing unit). For other model components, speed may hardly change (for example if performance is limited by the time certain information needs to travel from processor A to processor B, rather than the data volume). In general, for the IFS we have found a reduction in model-run time of approximately 40% for uncoupled simulations if single precision is used (*Váňa et al.*, 2016 and 2017). This result is consistent with results from other models (see for example *Nakano*, 2018).

Making single precision work in the IFS

The idea to use single precision in the IFS emerged from a research project at the University of Oxford carried out by Peter Düben and ECMWF Fellow Tim Palmer. Single precision was tested in the OpenIFS model, a portable version of the IFS that can be used for research projects at universities. It was shown that single precision simulations are possible and that the results are reasonable since differences between single and double precision simulations were smaller than the spread between ensemble members in ensemble simulations (see *Düben & Palmer*, 2014).

As a next step, single precision was introduced as an option into the IFS. A more extensive model evaluation revealed that single precision produced comparable results to double precision for ensemble simulations at about 50 km horizontal resolution (*Váňa et al.*, 2017). This work was carried out in close collaboration between ECMWF and Météo-France, which also successfully tested the use of single precision in global simulations.

Since then, a number of improvements to the single precision configuration at ECMWF have been made, and differences between the single precision configuration and the operational double precision configuration have been removed. This includes fixes in the Legendre transforms, the IO server, the coupling to the wave model WAM and the ocean model NEMO, the new radiation scheme, post-processing, the lake scheme, the vertical integration scheme of the dynamical core, the trigonometric grid information for simulations at very high resolution, and the fixing of several bugs that were not related to single precision but were identified thanks to the use of a different data layout in single precision. Many researchers at ECMWF have been involved in this process. As a result, today no additional code changes are required to run single precision forecast experiments in the IFS.

It should be noted that obtaining comparable results for single and double precision requires the use of double precision for some model components, including the pre-computation of Legendre polynomials and the operators of the vertical integration scheme before the time step loop is started. See Box B for details on which parts of the model code are sensitive to the use of single precision.

Quality of single precision forecasts

In general, if single precision is implemented as described above, forecasts produced using single and double precision are very similar. The main difference between a single and a double precision simulation is a larger error in the conservation of total air mass for single precision simulations. While we have been able to reduce the magnitude of the error, the mass conservation error in single precision simulations still has an effect on

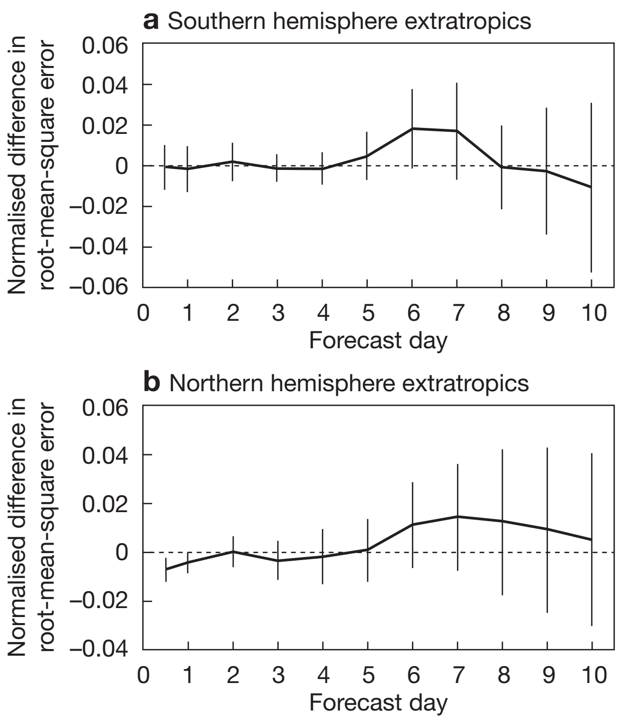

FIGURE 1 Normalised difference in the root-mean-square error for geopotential height at 500 hPa for a set of simulations in double and single precision at 9 km horizontal resolution (TCo1279) with 137 vertical levels for (a) the southern hemisphere extratropics (20°S to 90°S) and (b) the northern hemisphere extratropics (20°N to 90°N). For single precision simulations, the mass fixer was switched on. The difference was calculated as ‘double precision’ minus ‘single precision’ so that positive values indicate better results for single precision. The figures are based on the average of 45 simulations during January and February 2018. Vertical bars indicate the 95% confidence range.

b When do I need to be careful when using single precision?

There are a couple of common operations that can cause problems with numerical precision. They should therefore be avoided whenever possible or fixed using locally enforced double precision (via JPRD, see Box A).

- If large numbers are multiplied or if a number is divided by a very small number, results may become larger than the largest number that can be represented at a certain level of precision (for example >10^{38} for single precision). This will cause a number overflow and a crash of the model run. Rearranging the order of operations is often sufficient to avoid the multiplication of large numbers (for example $X^4/Y^6 \longrightarrow (X/Y)^4/Y^2$). If a divisor is very small and if there is a risk that the divisor may actually become zero, a number overflow can be avoided by adding a very small value (an epsilon) to the divisor.
- If numbers that are very similar in magnitude are subtracted from each other, several digits of precision can be lost in a single operation.
- It is possible for very small numbers that are added to large numbers to be rounded to zero. Even if the contribution of each summand on its own is not essential within a large sum, errors may be introduced when many small contributions are rounded to zero. It is therefore useful to begin sums over many numbers with the smallest number and to increase the size of the numbers that are added (a sum over pressure at model levels should, for example, start at the top of the model).
- If a very large or a very small number for which the exact value is not essential is required in the model code, it is important to use the intrinsic functions huge(1.0_JPRB) and epsilon(1.0_JPRB) to generate those numbers. This will adjust the value of numbers to the precision level that is used. Hard-coded numbers such as 10^{100} or 10^{-100} can cause a number overflow or will be rounded to zero if single precision is used.

forecast scores. We have identified three sources of the mass conservation error: the Legendre transformation, the vertical integration scheme and the semi-Lagrangian part of the model. The error in mass is fluctuating and can be both positive and negative with a small mean bias. Investigations have shown that reducing the error in mass conservation is not straightforward as it depends on a number of factors: it varies with resolution, is different for different initial conditions and even shows an annual cycle. However, the impact of the error in mass conservation can be mitigated by using a global mass fixer, which is cheap and easy to apply in a spectral model such as the IFS. If the mass fixer is switched on, differences in root-mean-square error between single precision and double precision geopotential height forecasts at the highest operational resolution (HRES) are mostly insignificant (Figure 1). However, even with the mass fixer switched on, there is currently still a degradation in some ensemble scores if single precision is used instead of double precision (Figure 2). Further investigations aimed at removing the remaining degradations are under way.

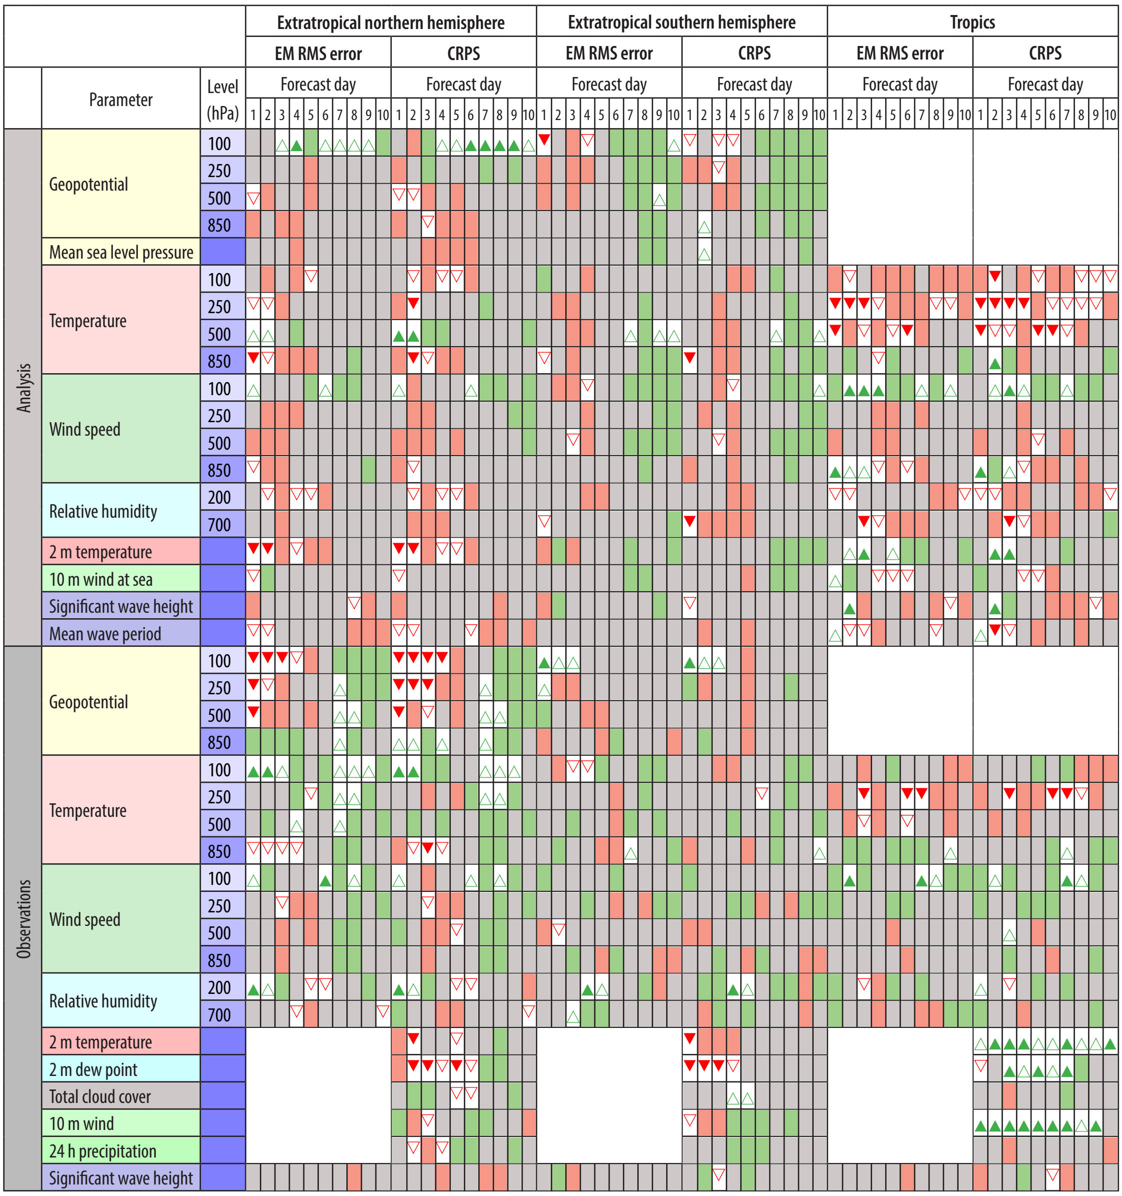

Symbol legend: for a given forecast step...

- ▲ Single precision better than double precision statistically significant with 99.7% confidence
- △ Single precision better than double precision statistically significant with 95% confidence
- Single precision better than double precision statistically significant with 68% confidence
- No significant difference between single precision and double precision
- Single precision worse than double precision statistically significant with 68% confidence
- ▽ Single precision worse than double precision statistically significant with 95% confidence
- ▼ Single precision worse than double precision statistically significant with 99.7% confidence

FIGURE 2 Ensemble score card comparing single precision to double precision (both with mass fixer). Results are based on 45 ensemble simulations up to forecast day 10 with 10 ensemble members during June, July and August 2017 at 18 km (TCo639) resolution with 137 vertical levels.

Single precision research experiments

Before making use of single precision in research experiments, it is important to verify that single precision simulations respond to changes in model configuration in the same way as double precision simulations. To test this, we have performed a set of simulations during winter that use double precision or single precision in the standard IFS model configuration on the one hand and in a model configuration in which the orographic gravity wave and low-level blocking parametrization was switched off on the other. This parametrization accounts for interactions of mountains with the flow that cannot be resolved explicitly on a given grid.

Figure 3 shows the differences between the simulations. Differences between single precision and double precision with the parametrization switched on are very small, while differences between the simulations with and without orographic gravity wave and blocking are virtually identical in single and double precision.

ECMWF aims to run ensembles at 5 km resolution by 2025. As the resolution increases beyond 10 km, some processes, such as convection or orographic drag, become resolved and the handover between parametrized and resolved processes poses a number of challenges. In order to prepare for future resolution upgrades and to make the most of the next supercomputer, it is therefore necessary to start testing the IFS at higher resolutions. However, testing the performance of the IFS at horizontal resolutions higher than that of the highest-resolution operational forecasts (9 km) is a very expensive exercise: the amount of information that needs to be available is very large and a very large amount of memory is needed to store the model state. Since the amount of memory that is available per compute node is limited, a large number of nodes is required for a single simulation. The overall performance of simulations is reduced since more data needs to be shared between processors and since the model needs to scale efficiently to a large number of processors.

Single precision will effectively reduce memory requirements by a factor of two, and this will have a very beneficial impact on the performance of simulations at very high resolution. We are therefore using single precision for tests in which the horizontal resolution is increased beyond the resolution of the deterministic operational forecasts. A series of two-day forecasts at horizontal resolutions ranging from 9 km to 1.25 km was performed for a day in August 2016 in the framework of the ESiWACE EU Horizon 2020 project, with both the IFS and the ICON model used by the German national meteorological service (DWD). The aim of these runs was to investigate scalability aspects in both models, but they are also very useful for understanding the challenges related to the representation of processes such as clouds, convection and precipitation in the 'grey zone'. For example, Figure 4 shows that, at 9 and 5 km with ECMWF's deep convection parametrization, the band of tropical rain over the Atlantic is too wide, but it has similar magnitude to the observed precipitation. In the 5, 2.5 and 1.25 km runs without the deep convection parametrization, the precipitation features become more realistic (the tropical band is narrower) but the rain is too intense. This suggests that work is needed both on the convection parametrization and its coupling with the dynamics in order to make the most of future resolution upgrades.

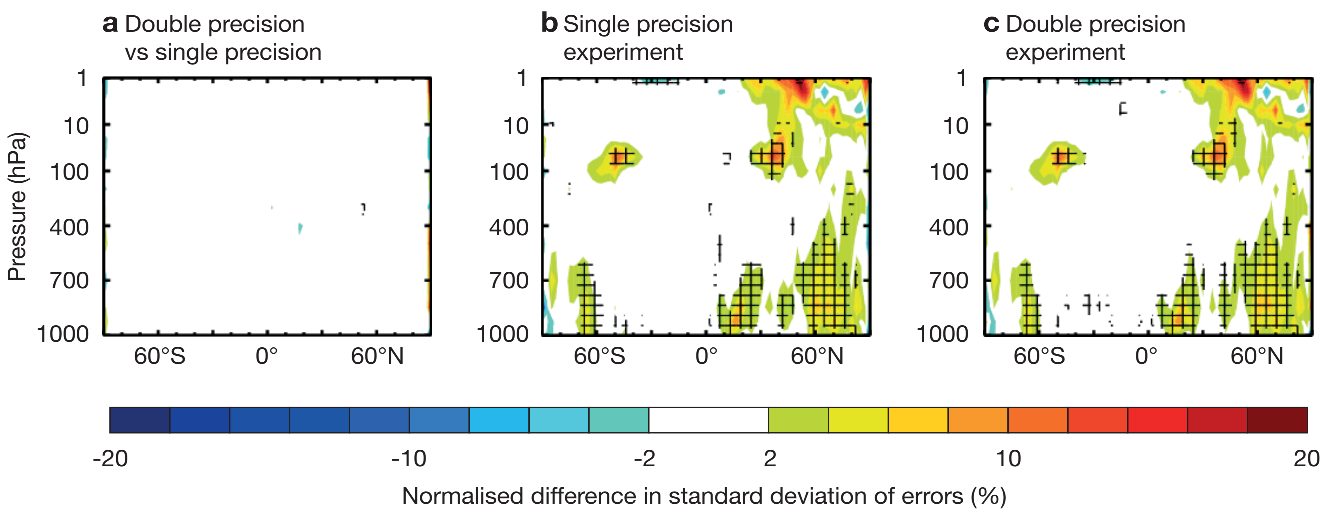

FIGURE 3 Normalised difference in standard deviation of errors against own analysis for vector components of horizontal winds for simulations at 25 km resolution (TCo399) and with 137 vertical levels after 24 hours, comparing (a) simulations with double precision with single precision simulations with mass fixer, (b) single precision simulations with and without gravity wave and low-level blocking parametrization and (c) double precision simulations with and without gravity wave and low-level blocking parametrization. The results are based on 30 forecasts with different starting times in December 2017. Cross-hatching indicates differences significant at the 95% level.

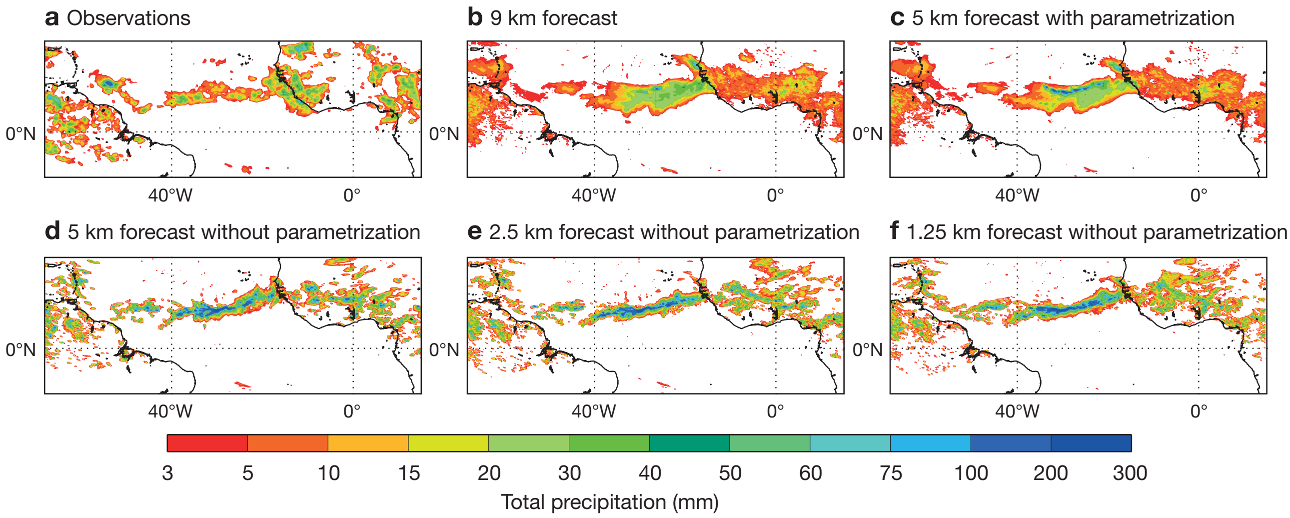

FIGURE 4 Total precipitation for 12 August 2016 over the tropical Atlantic according to (a) satellite observations (NASA's TRMM-3B42RT product), (b) a two-day 9 km horizontal resolution forecast (TCo1279), (c) a two-day 5 km horizontal resolution forecast (TCo1999) with deep convection parametrization, (d) a two-day 5 km horizontal resolution forecast without deep convection parametrization, (e) a two-day 2.5 km horizontal resolution forecast (TCo3999) without deep convection parametrization and (f) a two-day 1.25 km horizontal resolution forecast (TCo7999) without deep convection parametrization. The forecasts at TCo1999 and above were produced using single precision.

The future of single precision at ECMWF

The submission of ensemble simulations and analysis experiments that use the single precision version of the IFS will be possible straight from prepIFS from IFS Cycle 46r1, without any need for additional changes to the model configuration. However, the minimisation part of the 4D-Var data assimilation in the IFS and the tangent linear and adjoint models will remain in double precision since this part of the model has shown strong sensitivity to rounding errors in the past. To achieve optimal performance in standard research simulations using single precision, it may be useful to adjust the MPI/OpenMP configuration as well as code blocking (nproma).

In the near future, single precision will be tested for monthly and seasonal predictions as well as for atmospheric composition simulations. It is difficult to identify all differences between single and double precision in all aspects of the IFS. It will therefore be important for domain experts to have a more detailed look at the quality of single precision simulations in their specific area of expertise (such as land surface, ocean coupling, cloud physics, radiation, convection, stochastic parametrization schemes...) to identify any remaining differences.

In a second step, the use of single precision should also be tested in the NEMO ocean model. Preliminary tests on the use of single precision in NEMO at the Barcelona Supercomputing Centre are promising but more work is required.

In close collaboration with Tim Palmer's group at the University of Oxford, a reduction in numerical precision for weather forecasts beyond single precision is being investigated, such as the use of half precision (16 bit) arithmetic when calculating the Legendre transforms within the IFS, or a reduction in numerical precision when calculating dynamics at small spatial scales (*Hatfield et al.*, 2018; *Thornes et al.*, 2018).

The operational use of single precision will be a key element in moving towards the target of a 5 km ensemble set by ECMWF's Strategy to 2025. It will free up vital computational resources for forecast production and will thus maximise the benefits from the investment in ECMWF's next high-performance computing facility in Bologna from 2021.

Further reading

Düben, **P.D.** & **T.N. Palmer**, 2014: Benchmark Tests for Numerical Weather Forecasts on Inexact Hardware. *Mon. Wea. Rev.*, **142**, 3809–3829.

Hatfield, **S.**, **A. Subramanian**, **T. Palmer** & **P. Düben**, 2018: Improving Weather Forecast Skill through Reduced-Precision Data Assimilation. *Mon. Wea. Rev.*, **146**, 49–62.

Nakano, **M.**, **H. Yashiro**, **C. Kodama** & **H. Tomita**, 2018: Single Precision in the Dynamical Core of a Nonhydrostatic Global Atmospheric Model: Evaluation Using a Baroclinic Wave Test Case. *Mon. Wea. Rev.*, **146**, 409–416.

Thornes, **T.**, **Düben**, **P.** & **Palmer, T.**, 2018: A Power Law for Reduced Precision at Small Spatial Scales: Experiments with an SQG Model. *Q J R Meteorol Soc*, doi: 10.1002/qj.3303.

Váňa, **F.**, **G. Carver**, **S. Lang**, **M. Leutbecher**, **D. Salmond**, **P. Düben** & **T. Palmer**, 2016: Single-precision IFS, *ECMWF Newsletter* **No. 148**, 20–23.

Váňa, **F.**, **P. Düben**, **S. Lang**, **T. Palmer**, **M. Leutbecher**, **D. Salmond** & **G. Carver**, 2017: Single Precision in Weather Forecasting Models: An Evaluation with the IFS. *Mon. Wea. Rev.*, **145**, 495–502.

Major upgrade for European flood forecasts

Cinzia Mazzetti, Christel Prudhomme

On 16 May 2018, ECMWF implemented a major upgrade of the European Flood Awareness System (EFAS), which produces flood forecasts for 57 countries. The new model cycle, informally named 'Extended Domain' and officially EFAS version ER15, has been co-developed by ECMWF, which is the EFAS computational centre, and the European Commission's Joint Research Centre (JRC), which manages EFAS.

The upgrade has increased the geographic coverage of EFAS and has improved the observation processing and hydrological modelling. Technical changes include a revision of file formats and the geographic projection to make the products more user-friendly. The model computer code has been modernised, new meteorological and static input maps were created, the hydrological model was recalibrated, and a new climatology and reforecasts were generated.

A full evaluation of forecast performance for EFAS vER15 is in progress, and summary results will be provided to EFAS partners when the evaluation is complete. Initial findings show improvements in hydrological model performance against observations when the new model calibration is used. This is expected to improve forecast performance as it means that the recalibrated model is better at representing the physical behaviour of catchment areas.

What is EFAS?

EFAS is an operational pan-European flood forecasting system funded by the European Commission through its Copernicus Programme. The aim of EFAS is to support preparatory measures before major flood events strike, particularly in large transnational river basins and throughout Europe in general.

Following the major floods that hit parts of central Europe in 2002, EFAS was developed and tested at the Joint Research Centre (JRC) between 2003 and 2012, in close collaboration with national hydro-meteorological services across Europe, the European Commission's Emergency Response and Coordination Centre (ERCC),

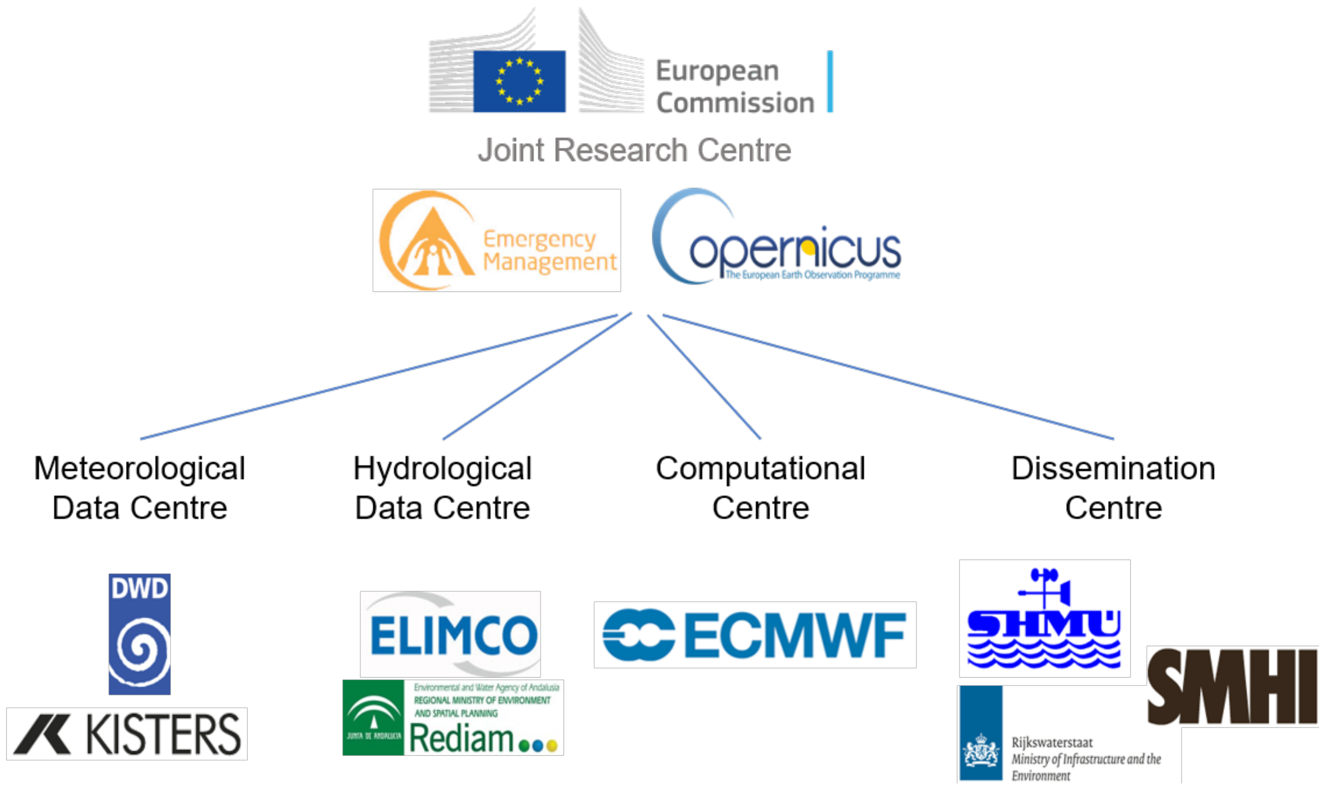

FIGURE 1 The EFAS organisational structure.

doi: 10.21957/3er1zcm8f6

and other research institutes. In 2012, EFAS became the Early Warning System for flood hazards of the Copernicus Emergency Management Service (CEMS), and ECMWF became the computational centre for CEMS-floods (EFAS-COMP). The Centre also took over responsibility for the web interface and archiving.

EFAS provides coherent medium-range flood forecasts and related information, including short-range flash flood products, flood impact assessment and hydrological seasonal outlooks. It keeps the ERCC informed about ongoing floods and about the possibility of upcoming floods across Europe. Today, EFAS delivers forecasts to more than 70 hydro-meteorological and civil protection services in Europe.

EFAS is managed by the JRC, and it comprises four centres run by different consortia (Figure 1):

- Computational centre (EFAS-COMP): the European Centre for Medium-Range Weather Forecasts produces forecasts and hosts the EFAS-Information System platform
- Dissemination centre (EFAS-DISS): the Swedish Meteorological and Hydrological Institute, the Dutch Rijkswaterstaat, and the Slovak Hydro-Meteorological Institute analyse EFAS on a daily basis and disseminate information to EFAS partners and the ERCC
- Hydrological data collection centre (EFAS-HYDRO): the Environmental and Water Agency of Andalucía (REDIAM) and Soologic Technological Solutions SL collect historical and real-time river discharge and water level data across Europe
- Meteorological data collection centre (EFAS-METEO): KISTERS AG and the German national meteorological service (Deutscher Wetterdienst, DWD) collect historical and real-time meteorological data across Europe

How does EFAS work?

EFAS forecasts are generated by cascading an ensemble of meteorological forecasts (from ECMWF, DWD and the COSMO Limited-Area Ensemble Prediction System consortium), meteorological and hydrological observations, land surface information and model parameters (static maps) through a deterministic hydrological model (LISFLOOD) and algorithms for flash floods. The resulting flood forecasts are then post-processed to produce all EFAS products, including flood alerts of different severity levels. Three alert levels are highlighted, corresponding to forecasts of floods expected to exceed flood peaks with return periods of 2, 5, and 20 years (a return period indicates the average number of years expected to pass between two floods of the predicted magnitude or greater). After the flood forecasts are produced at ECMWF, EFAS-DISS duty officers evaluate the forecasts and issue email notifications to EFAS partners. The schematic overview in Figure 2 summarises the EFAS production and dissemination chain.

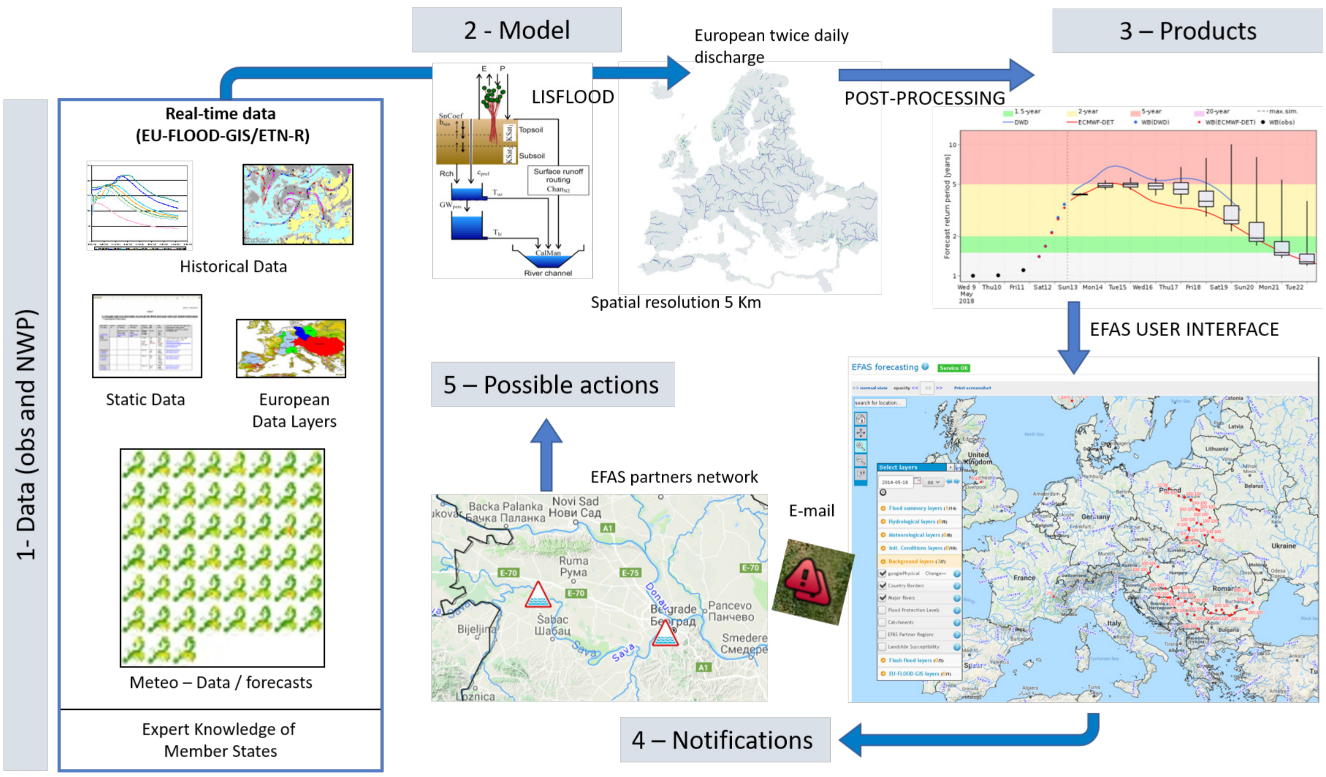

FIGURE 2 The EFAS production and dissemination chain.

For more information on EFAS, see: *www.efas.eu/user-information.html*

EFAS vER15: a major upgrade

Like other operational forecasting systems, EFAS is always evolving, but the May 2018 upgrade combined a large number of changes in a single development cycle. The sections that follow provide an overview of what has changed. Further details on the upgrade can be found in *Salamon et al.* (2018).

EFAS expands to the east and the south

To reflect the growing number of EFAS partners, in version ER15 geographic coverage has been increased by more than two million km^2 to the east and the south. It now includes about 5,300 new river basins located near the eastern and southern borders of the EU and 14 new countries (Turkey, Cyprus, Georgia, Azerbaijan, Armenia, Israel, Lebanon and parts of Syria, Iraq, Egypt, Libya, Tunisia, Algeria and Morocco) (Figure 3).

The change in the geographic domain was accompanied by a change in the geographical projection. EFAS now uses the INSPIRE-compliant ETRS89 Lambert Azimuthal Equal Area Coordinate Reference System (ETRS-LAEA). ETRS-LAEA is the recommended projection for pan-European statistical mapping at all scales or where true area representation is required.

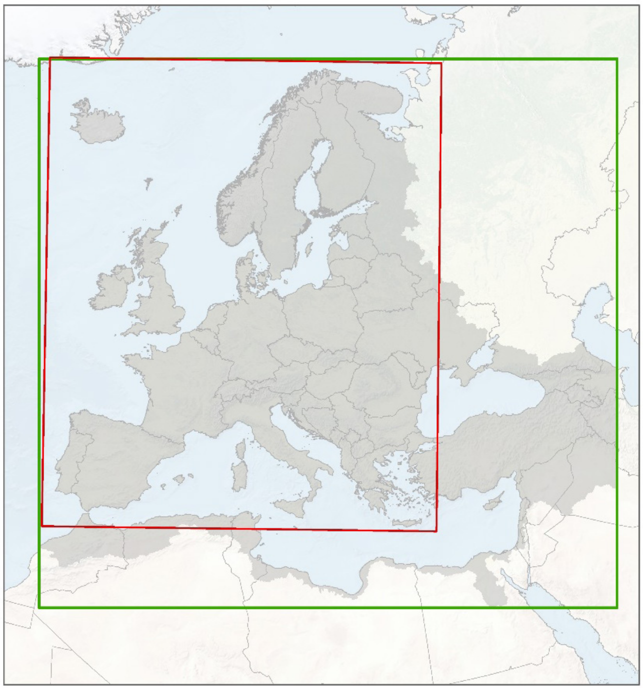

FIGURE 3 EFAS domain limits: the green box shows the new Extended Domain area, the red box the original domain area, and the grey shading the limits of the hydrological modelling for EFAS vER15.

New meteorological observation processing

Meteorological inputs are crucial to EFAS to run long-term simulations and to estimate the hydrological initial conditions of the system's forecasts. EFAS-METEO collects datasets of historical and real-time in-situ meteorological observations on a 24/7 basis and interpolates them to the 5 km hydrological model grid. The data is then sent to ECMWF each day for inclusion in the EFAS modelling chain. The number of contributors and observations (both in near-real time and historical records) provided to EFAS keeps growing. Today about 5,600 rain gauges and 8,900 temperature stations routinely provide data in real time for EFAS vER15. As part of the upgrade, a new interpolation scheme was introduced, called SPHEREMAP. The scheme gives different weights to observations depending on the number of nearby stations and their distance from the interpolated point and from each other, to avoid giving too much weight to stations in close proximity to each other. All meteorological observation forcing maps are EFAS products. They are included as summary maps in the EFAS web interface, which is currently only accessible to EFAS partners.

Quality-assured hydrological data

Near-real-time river discharge observations are used to adjust (post-process) EFAS forecasts. The EFAS hydrological data collection centre oversees collecting, validating and filling gaps in discharge and water level data for both near-real-time and historical records. As part of the recent upgrade, EFAS-HYDRO increased the number of data providers and it now collects data for about 1,700 near-real-time stations and about 2,900 historical stations from 60 different data providers.

Hydrological model upgrade

Most EFAS flood forecast products are based on outputs from LISFLOOD, a rainfall-runoff and routing model. LISFLOOD has been developed at the JRC since the year 2000 and has been used for operational flood forecasting at the pan-European scale since the early days of EFAS. Driven by meteorological forcing data (precipitation, temperature, potential evapotranspiration, and evaporation rates for open water and bare soil surfaces), LISFLOOD calculates a complete water balance at a 6-hourly or daily time step for every 5 km grid cell within the EFAS domain. Processes simulated for each grid cell include snowmelt, soil freezing, surface runoff, infiltration into the soil, preferential flow, redistribution of soil moisture within the soil profile, drainage of water to the groundwater system, groundwater storage, and groundwater base flow. Runoff produced for every grid cell is then routed through the river network using a kinematic wave approach. The model also includes options to simulate lakes and reservoirs (Figure 4).

In EFAS vER15, a new parametrization of water infiltration into the soil using three soil layers instead of two was introduced. Evaporation from open water was also included, to improve the realism of modelled losses from reservoirs and lakes. A water abstraction component was added, making it possible to define the amount of water resources removed for the purposes of agriculture including livestock farming, the manufacturing industry, energy production (cooling water needs), and public-sector water use. Finally, crop irrigation and paddy-rice irrigation components were added to dynamically simulate the water needs associated with the cultivation of land.

The upgrade also includes an improved representation of the storage and release of water by different land surface components. The new scheme better reflects how much water is already present and how much is added/removed by the atmosphere and vegetation.

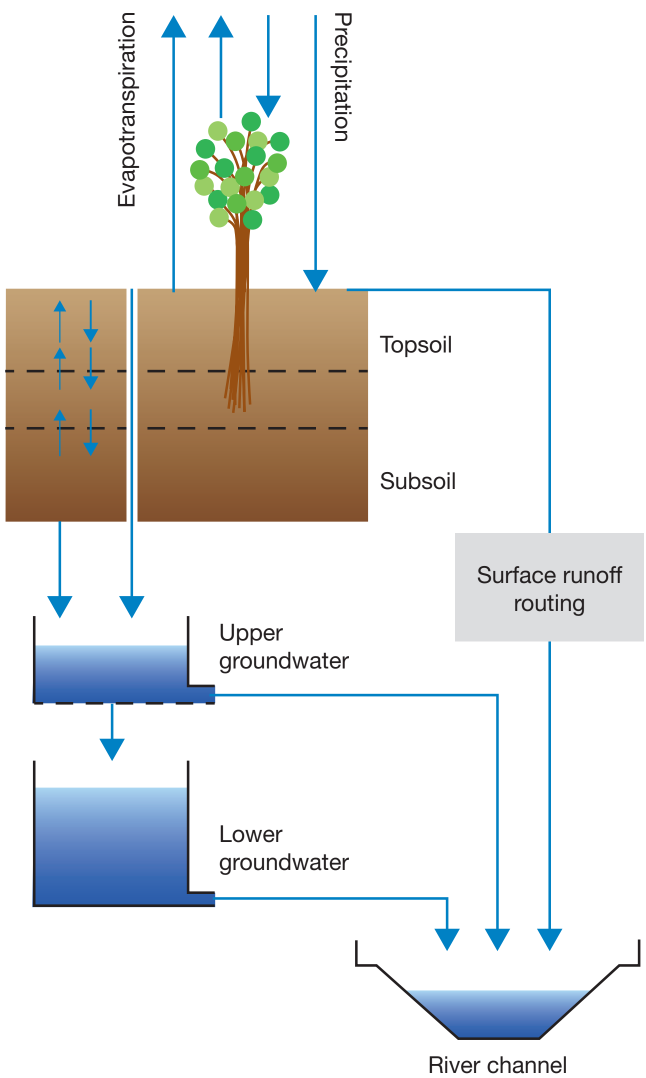

FIGURE 4 The LISFLOOD main model structure for a single grid cell. Parameters such as the snowmelt coefficient, soil infiltration rate and surface runoff coefficient help to model the processes that determine the magnitude of the flow in a river channel.

This part of the water cycle has an impact on future river discharge.

Change in LISFLOOD code

Since the first operational implementation of EFAS in 2012, the LISFLOOD code has been completely rewritten by the JRC. Earlier LISFLOOD versions were coded using PCRaster, a software package for environmental dynamic modelling including convenient hydrological and hydraulic routing functions. The LISFLOOD version implemented in EFAS vER15 was recoded using the Python programming language and PCRaster Python extension, which enables the use of the PCRaster modelling engine from Python. The structure of the code was made more modular and flexible. The NetCDF file format was adopted for inputs and outputs, to replace the previous PCRaster format. The internal file structure and naming convention was also changed to comply with World Meteorological Organization conventions.

Upgrade in static input maps

The EFAS pan-European setup of LISFLOOD uses a 5 km grid and spatially variable input parameters, also called static maps. Sub-grid information on land use (100 m resolution), soil type (1 km resolution) and elevation (100 m resolution) is incorporated into the model using land use fractions. EFAS vER15 uses a revised range of land use types, including forest, built-up area, water surface, irrigated land and paddy rice land. Leaf Area Index maps are used to help simulate evapotranspiration. Elevation bands within each 5 km grid based on 100 m digital elevation model maps are used to model snow accumulation and snowmelt. Soil information at 1 km is used to determine hydraulic soil parameters, aggregated to 5 km. River channel network and related parameters as well as water demand, water abstraction and water consumption maps are used for all water-balance and river routing processes. All maps were upgraded by the JRC to match the new projection, spatial coverage and process representation of EFAS vER15, with final checks and integration into the production chain performed by ECMWF.

Finally, the number of reservoirs accounted for in EFAS vER15 is 1,454, up from 34 in the previous version. Location and total reservoir volume are the primary metrics available from external sources. Other reservoir parameters were estimated or included in model calibration.

Recalibration of the hydrological model

Like most rainfall-runoff models, LISFLOOD hydrological model equations include a range of parameters. Some of these can be determined from physical data, such as reservoir storage-elevation curves or the drainage area of watersheds. Others vary from one area to another

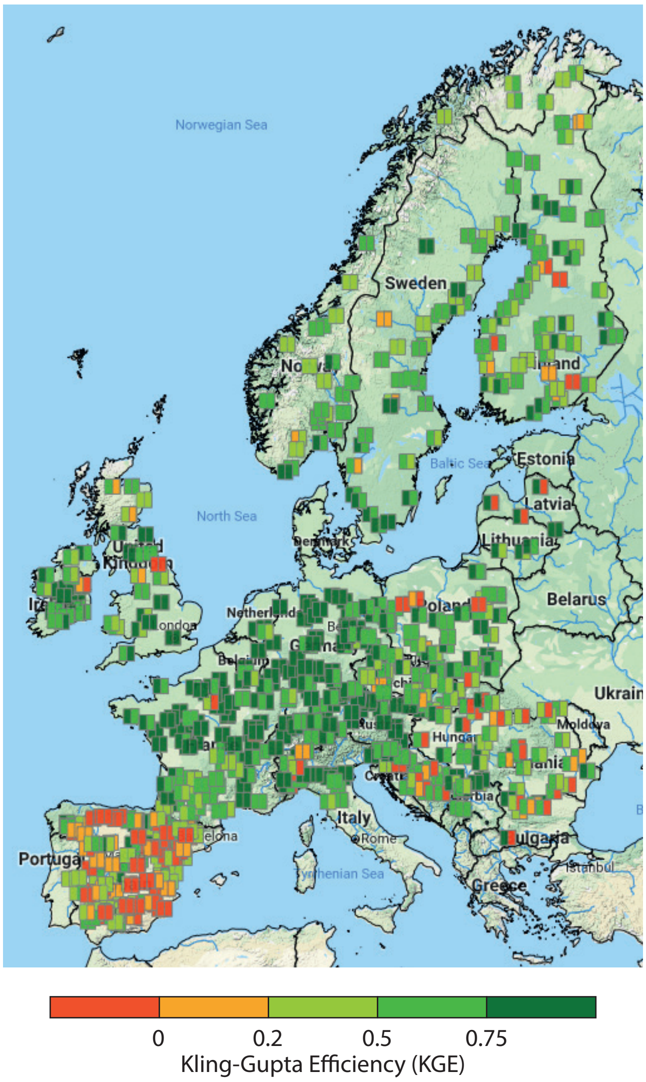

FIGURE 5 Hydrological skill of EFAS at the calibration locations in the calibration period (left-hand half of each square) and the validation period (right-hand half of each square).

based on changes in climatology and physical factors, e.g. hydraulic soil properties. Some model parameters require calibration, which is generally obtained by tuning parameter values based on a comparison between simulated and observed daily discharge (Q) at catchment outlets. This tuning or optimisation process generally aims to minimise errors in the volume and timing of simulated flow over a multi-year period.

With new processes introduced and increased geographic coverage, LISFLOOD was recalibrated by the JRC, which produced pan-European parameter maps based on calibration over 717 catchments. The optimisation procedure was carried out using Kling-Gupta Efficiency criteria (KGE). KGE is a goodness-of-fit measure varying between 1.0 (perfect simulation) and negative infinity (lower limit), developed to diagnose differences in correlation, bias and variability between observations and simulations. The 13 model parameters optimised by calibration were first estimated by prior expert knowledge. They relate to snowmelt, infiltration, preferential bypass flow through the soil matrix, percolation to the lower ground water zone, percolation to deeper groundwater zones, residence times in the soil and subsurface reservoirs, river routing and reservoir/lake simulation. The 717 calibration points were selected according to spatial and data availability criteria. To evaluate the temporal transferability of the calibrated parameter set, observed meteorological and river discharge records were split into a ‘calibration’ and an ‘evaluation’ period.

Following calibration, 75% of stations scored a KGE higher than 0.5 for the calibration period and 57% scored more than 0.5 for the evaluation period. The skill varies across geographical domains, with higher skill scores in large parts of central Europe and lower skill scores mostly in catchments where the influence of reservoirs is stronger, such as in Spain (Figure 5). Note, however, that low hydrological simulation skill associated with calibration/evaluation periods is not necessarily an indicator of poor forecast performance as EFAS forecasts are compared with model-derived thresholds.

New flood magnitude thresholds

After all model implementation and calibration tasks were completed, LISFLOOD reference simulations were generated for the period 1990 to 2016 using observed meteorological data as input. This reference simulation, also called ‘simulation from observations’, was used to estimate new discharge thresholds associated with floods with return periods of 1, 2, 5 and 20 years, and threshold anomalies for the seasonal outlooks. Those new thresholds now form the basis for flood notifications.

In addition, the European Runoff Index based on Climatology (ERIC) algorithm, which is used to derive flash flood products from COSMO-LEPS rainfall forecasts and near-real-time LISFLOOD soil moisture simulations (Figure 6), was also upgraded to account for the new three-layer soil scheme in LISFLOOD.

Hydrological forecast post-processing

Where near-real-time river discharge observations are available, LISFLOOD flood forecasts can be adjusted via statistical post-processing to minimise errors in the timing, the volume and the height of the peak when compared to observations, but also to derive more accurate probabilistic forecasts (Figure 7). The EFAS vER15 system now includes about 600 adjusted (post-processed) forecast points, covering all major European rivers. This has been made possible thanks to the strong partnership between hydro-meteorological agencies across Europe and the EFAS consortium, as a result of which access to automatic river gauge measurements was granted through EFAS-HYDRO.

Seasonal forecast upgrade

As part of the upgrade, the EFAS seasonal forecast production chain was modified to now use ECMWF's most recent seasonal forecasting system (SEAS5). SEAS5 benefits from a much higher atmospheric horizontal resolution than its predecessor S4 (36 km instead of 80 km). It also includes a better ocean model and a 25-member re-forecast ensemble (up from 15 members in S4). The EFAS seasonal release date has also been brought forward, from the 7th to the 5th day of each month.

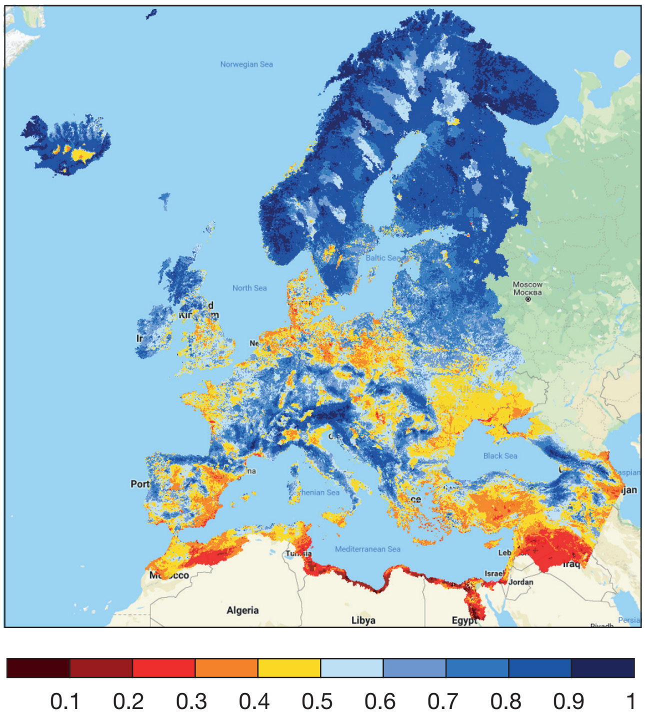

FIGURE 6 Simulated relative soil moisture (soil moisture content divided by soil moisture capacity) over the top two soil layers based on output from the LISFLOOD model with observed meteorological input, on the EFAS vER15 domain on 27 May 2018 at 00 UTC.

Flood hazard and impact assessment

The EFAS system uses a rapid risk assessment procedure to estimate flood extent as well as potential flood impacts based on EFAS forecasts. Every time a flood event is forecast in EFAS, the procedure identifies river sections where the magnitude of predicted peak discharge is expected to exceed the local flood protection level. For these river sections, the procedure identifies flood-prone areas using a catalogue of flood hazard maps covering the entire EFAS river network. These event-based flood hazard maps are combined with exposure information to assess several categories of impact, such as affected population, roads and cities, the total extent of urban and agricultural areas affected, and direct economic losses.

The EFAS vER15 upgrade has integrated the FLOPROS global dataset and information on design protection levels, where available, into the EFAS database of flood protection levels, and the global set of flood damage functions provided by *Huizinga et al.* (2017) into the vulnerability functions.

Evaluation

Although a full evaluation of forecast performance for EFAS vER15 is still in progress, some results comparing the new calibration with the previous calibration are already available. In six catchment areas, hydrological simulations using the previous calibration and the new

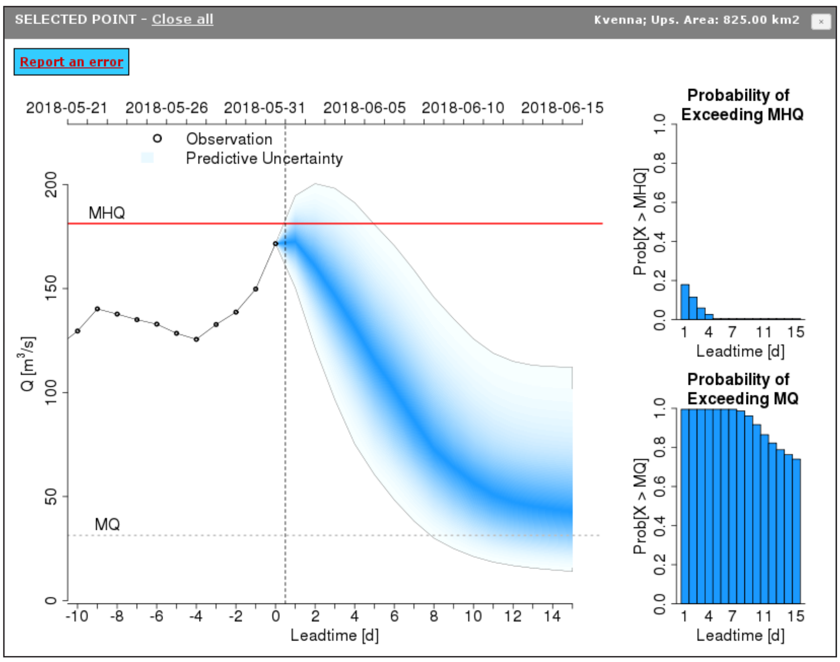

FIGURE 7 Example of a 15-day post-processed forecast. To the left of the dashed line, the main chart shows the observed discharge for the last 10 days. The top-right chart shows the probability of exceeding the average yearly maximum value (MHQ), and the bottom-right chart shows the probability of exceeding the average discharge (normal condition, MQ).

Catchment	Station	Area (10^3 km^2)	Change in bias (%)	Change in KGE (%)
Danube	Bazias	569	2.0	5.0
Rhine	Duesseldorf	147	2.9	0.0
Oder	Hohensaaten	112	-1.6	4.2
Po	Pontelagoscuro	71	13.4	26.1
Elbe	Dresden	52	0.5	2.4
Ebro	Zaragoza	41	5.4	2.5

TABLE 1 Differences between the new and the previous calibration for six catchment areas. Positive values mean a smaller bias and better KGE score when using the new calibration. The bias values were calculated in per cent by comparing simulations with observations.

calibration were compared against discharge observations to determine differences in bias and KGE scores. The results, summarised in Table 1, show a general improvement in scores when using the new calibration. This is expected to lead to better forecast performance as it means that the recalibrated hydrological model is better able to represent the physical behaviour of the catchments.

EFAS verification helps to assess the performance of ECMWF's forecasts, for example precipitation forecasts. It also helps to evaluate the realism and performance of land surface modelling, in particular for land surface variables such as runoff and soil moisture. As part of the verification of EFAS forecasting performance, the EFAS team works closely with other teams at ECMWF to develop appropriate verification software.

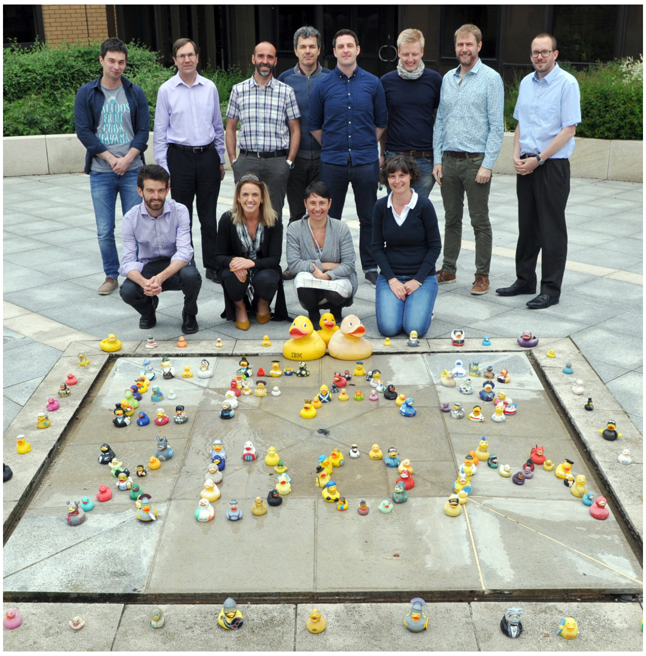

Team effort. The development and implementation of EFAS vER15 was a joint ECMWF–JRC effort over nearly three years. The EFAS team at ECMWF marked the implementation by a rearranged display of ducks outside the ECMWF building.

Further reading

Burek, **P.**, **J. van der Knijff** & **A. de Roo**, 2013: LISFLOOD Distributed Water Balance and Flood Simulation Model. Revised User Manual. *JRC Technical Report* **No. JRC78917**.

Huizinga, **J.**, **H. de Moel** & **W. Szewczyk**, 2017: Global Flood depth-damage functions. *JRC Technical Report* **No. JRC105688**.

Salamon, **P.**, **L. Arnal**, **S. Asp**, **C. Baugh**, **H. Beck**, **B. Bisselink**, **A. De Roo**, **J. Disperati**, **F. Dottori**, **M. Garcia-Padilla**, **R. Garcia-Sanchez**, **E. Gelati**, **G. Gomes**, **M. Kalas**, **B. Krzeminski**, **M. Latini**, **V. Lorini**, **C. Mazzetti**, **M. Mikulickova**, **D. Muraro**, **C. Prudhomme**, **A. Rauthe-Schöch**, **K. Rehfeldt**, **C. Schweim**, **J. Skoien**, **P. Smith**, **E. Sprokkereef**, **V. Thiemig**, **F. Wetterhall** & **M. Ziese**, 2018: EFAS upgrade for the extended model domain, EUR 29323 EN, Publications Office of the European Union, Luxembourg, doi:10.2760/806324.

Smith, **P.**, **F. Pappenberger**, **F. Wetterhall**, **J. Thielen**, **B. Krzeminski**, **P. Salamon**, **D. Muraro**, **M. Kalas** & **C. Baugh**, 2016: On the operational implementation of the European Flood Awareness System (EFAS). *ECMWF Technical Memorandum* **No. 778**.

ECMWF publications

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

Technical Memoranda

830 **Vitart**, **F.** & **M. Alonso-Balmaseda**: Impact of sea surface temperature biases on extended-range forecasts. *October 2018*

829 **Buizza**, **R.**, **M. Alonso-Balmaseda**, **A. Brown**, **S.J. English**, **R. Forbes**, **A.J. Geer**, **T. Haiden**, **M. Leutbecher**, **L. Magnusson**, **M. Rodwell**, **M. Sleigh**, **T. Stockdale**, **F. Vitart** & **N. Wedi**: The development and evaluation process followed at ECMWF to upgrade the Integrated Forecasting System (IFS). *October 2018*

828 **Bonavita**, **M.**, **P. Lean** & **E. Hólm**: Nonlinear effects in 4D-Var. *September 2018*

827 **Laloyaux**, **P.**, **S. Frolov**, **B. Ménétrier** & **M. Bonavita**: Implicit and explicit cross-correlations in coupled data assimilation. *August 2018*

825 **Yepes-Arbos**, **X.**, **M. Acosta**, **G. van den Oord** & **G. Carver**: Computational aspects and performance evaluation of the IFS-XIOS integration. *September 2018*

823 **Zuo**, **H.**, **M. Alonso-Balmaseda**, **K. Mogensen** & **S. Tietsche**: OCEAN5: The ECMWF Ocean Reanalysis System and its Real-Time analysis component. *August 2018*

EUMETSAT/ECMWF Fellowship Programme Research Reports

47 **Burrows**, **C.P.**: Assimilation of radiance observations from geostationary satellites: First year report. *June 2018*

ESA Contract Reports

Fielding M. & **M. Janiskova**: Operational assimilation of space-borne radar and lidar cloud profile observations for numerical weather prediction. Observation quality and pre-processing. *2018*

Janiskova, **M.** & **M. Fielding**: Operational assimilation of space-borne radar and lidar cloud profile observations for numerical weather prediction. Feasibility demonstration of 4D-Var assimilation system using CloudSat and CALIPSO observations. *2018*

Janiskova, **M.** & **M. Fielding**: Operational assimilation of space-borne radar and lidar cloud profile observations for numerical weather prediction. Conclusions and recommendations. *2018*

Fielding, **M.**, **M. Janiskova** & **R. Hogan**: Operational assimilation of space-borne radar and lidar cloud profile observations for numerical weather prediction. EarthCARE data handling and testing. *2018*

Janiskova, **M.**, **M. Fielding**, **M. Crepulja**, **D. Vasiljevic**, **T. Kral** & **P. Lean**: Operational assimilation of space-borne radar and lidar cloud profile observations for numerical weather prediction. Assimilation system development for cloud radar and lidar observations. *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 2018/19

Dec 4–5	Council
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 29–25	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	Advisory Committee for Data Policy and data policy meetings of ECOMET and EUMETSAT
May 2–3	Finance Committee
May 3	Policy Advisory Committee (PAC)
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 27–28	Council
Sep 2–6	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	Advisory Committee of Co-operating States (in Estonia)
Oct 29	Policy Advisory Committee
Nov 4–8	Workshop on the use and diagnostics of ECMWF forecasts during meteorological campaigns
Dec 10–11	Council

Newsletter | **No. 157** | Autumn 2018

European Centre for Medium-Range Weather Forecasts

www.ecmwf.int