ECMWF Newsletter

Number 153 – Autumn 2017

European Centre for Medium-Range Weather Forecasts
Europäisches Zentrum für mittelfristige Wettervorhersage
Centre européen pour les prévisions météorologiques à moyen terme

25 YEARS OF ENSEMBLE PREDICTION

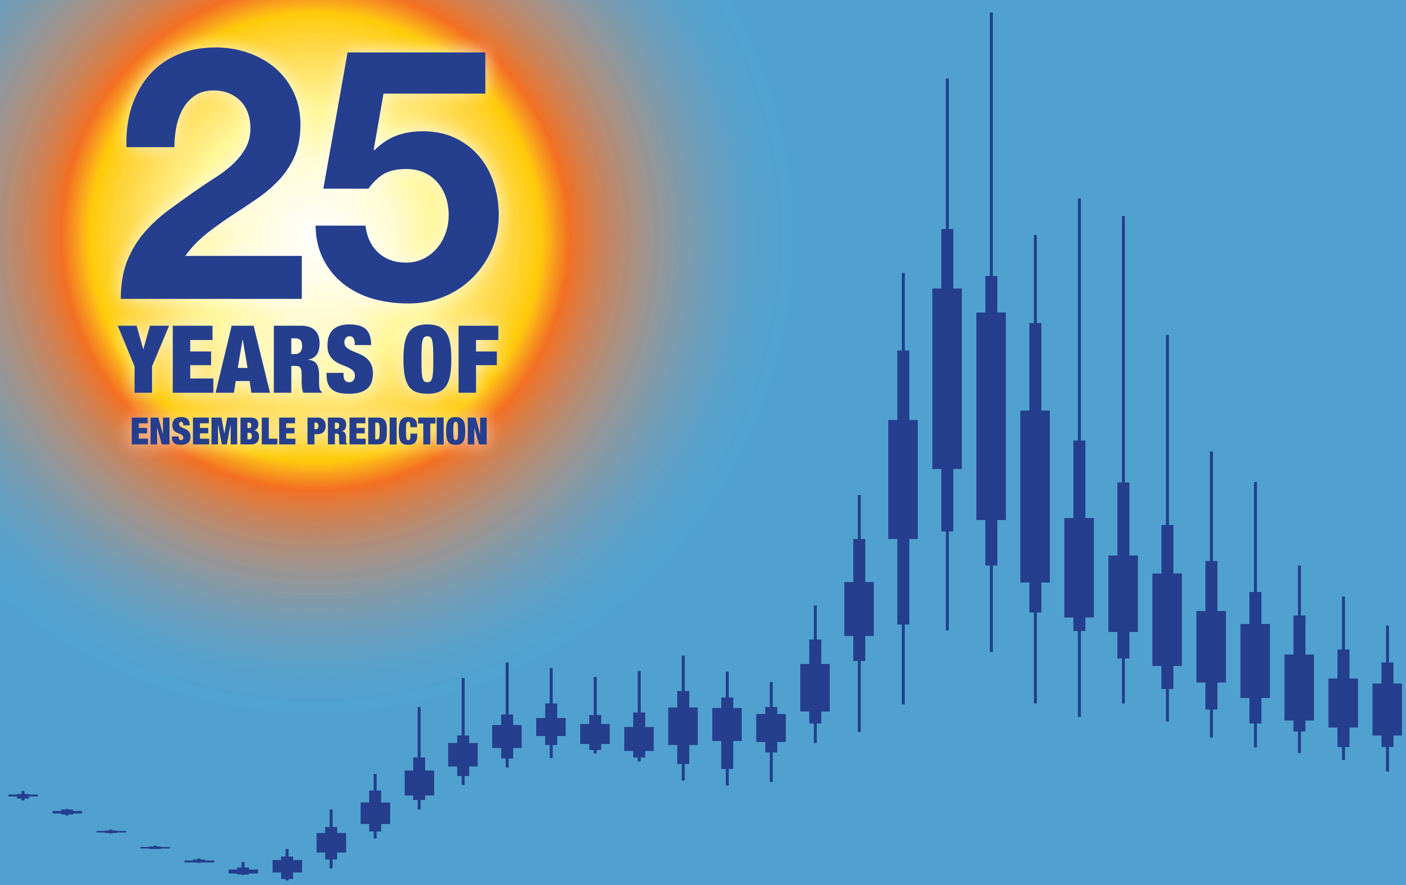

www.ecmwf.int/en/about/news-centre/media-resources

CONTENTS

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 The *ECMWF Newsletter* is not peer-reviewed.

Editor: Georg Lentze

Typesetting and Graphics: Anabel Bowen with the assistance of Simon Witter

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

CONTACTING ECMWF

Shinfield Park, Reading, Berkshire RG2 9AX, UK

Fax:		+44 118 986 9450
Telephone:	National	0118 949 9000
	International	+44 118 949 9000
ECMWF website:		www.ecmwf.int

Foresight

One year into the implementation of our 2025 Strategy, a triple anniversary highlights some of the science that plays a key role at ECMWF: 25 years of ensemble prediction, whose future was discussed at our Annual Seminar in September; 20 years of 4D-Var data assimilation, to be celebrated with a symposium on 26 January 2018; and 20 years of seasonal forecasts, which we are marking in style by releasing the new seasonal forecasting system SEAS5 on 5 November 2017.

Advances in these three areas as well as others have enabled substantial improvements in forecasts of severe weather, such as the heat wave that affected southern Europe in August or the tropical cyclones that have hit parts of the Caribbean and the US over the last few weeks. Better modelling of coupled processes, such as ocean–atmosphere interactions, is part of our strategic move towards an Earth system approach. Progress in this area has had a significant impact on the skill of ECMWF's extended-range forecast. The improvement has been particularly big in the tropics, especially in predicting the Madden–Julian Oscillation (MJO). As a result, the sub-seasonal forecast skill of a wide range of high-impact weather events, including tropical cyclones, has increased. Extended-range forecasting over Europe has also benefited through the impact of the MJO on the North Atlantic Oscillation. Further improvements in Europe are the result of including a sea-ice model.

Ensemble prediction, 4D-Var data assimilation and modelling coupled processes all carry high computational costs. These must be judged against the benefits they bring. As the science progresses, so does our need for more computing power. Our challenge today is to find ways to enable continued advances by our scientists by providing the required computing power whilst respecting legitimate financial and environmental constraints. Our ambitious Scalability Programme helps us in this process by optimising our use of computing resources.

We are only one piece of the puzzle, and it is not our sole responsibility to push the limits of science and technology to improve numerical weather prediction. Many others are involved, and we are working closely with them on the challenges that lie ahead. But a lot is at stake, and experience shows that success comes through each of us bringing our best to the table. The European heat wave of August 2017 was well predicted, and so were the recent tropical cyclones, but there is still room for much improvement, for example in predicting the intensification of hurricanes and severe convection events and in quantifying the uncertainty of our predictions.

Former NOAA Administrator Dr Kathryn Sullivan has described numerical weather prediction as the ability to give humankind foresight. This is an elegant and inspirational way to describe our job. Let's prove her right and get as close as we can to 20/20 foresight.

Florence Rabier
Director-General

New point-rainfall forecasts for flash flood prediction

FATIMA PILLOSU (ECMWF and University of Reading), **TIM HEWSON**

ECMWF has developed a probabilistic point-rainfall product which could support the prediction of flash floods across the globe. The product, which is based on an innovative post-processing method, aims to bridge the gap between the relatively coarse resolution of today's global forecasting models and the higher-resolution limited-area models needed to describe localised heavy rainfall. The methodology is based on physically relevant statistical relationships between the larger-scale weather features well represented by ECMWF forecasts and local realisations represented by point observations. These relationships make it possible to compute statistically based (rather than raw-ensemble-based) probabilities for point rainfall. This includes extremes, which can be used to infer the likelihood of flash floods for use on platforms such as the European and Global Flood Awareness Systems.

The post-processing method blends together information from different locations whenever they experience similar rainfall generation mechanisms, assuming that these physical mechanisms are universal and dependent on key atmospheric and geographic properties. This means that:

- one year of global rainfall observations is adequate because it can equate to hundreds of years used in locally-calibrated techniques, and
- extremes can be successfully predicted even when they do not exist in a local record.

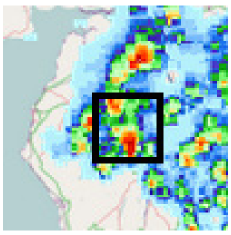

Example of large sub-grid variability due to slow-moving convection. The chart shows radar-derived totals (*source: netweather.tv*). The black square is an ENS gridbox, colours denote 0–80 mm.

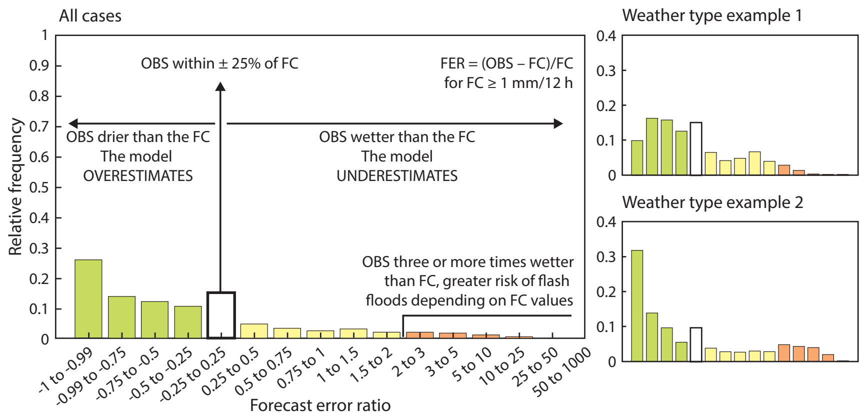

Forecast Error Ratio (FER) frequency distributions. The distributions are for 12-hourly post-processed short-range rainfall forecasts (FC) based on a nine-month training period, representing 1.6 million global rainfall SYNOP observations (OBS). The panels on the right show how frequency distributions change for two weather type examples: winter insolation (top) and summer insolation (bottom), both with mainly convective rainfall, moderate steering winds and medium CAPE.

Moreover, the reliance on physics means that forecasts can be confidently produced for anywhere in the world, even places without observations. The post-processing system has been fully automated and requires minimal computing resources to run compared to high-resolution numerical models.

Methodology

Radar-derived totals show that very different geometries of rainfall total variability can arise within regions of a global model grid box of typical size. Global models include standard output parameters such as convective precipitation fraction, mid-tropospheric wind speed, convective available potential energy (CAPE), total daily clear-sky solar radiation and total precipitation that, on the basis of physical reasoning, make it possible to anticipate sub-grid variability and account for biases in grid-scale rainfall. Those parameters are used to define subsets, hereafter called 'weather types', which share the same meteorological/geographical conditions.

The amount by which observed rainfall totals at a point differ from short-range forecasts for the equivalent model gridbox can be expressed as a ratio, hereafter called Forecast Error Ratio (FER). By compositing many cases a general frequency distribution can be derived, which can be used to transform ensemble gridbox rainfall forecasts into probabilistic point-rainfall forecasts. The true utility of the new approach, however, lies in creating and using separate frequency distributions for weather types that differ from one another in significant and physically realistic ways, as illustrated by the small panels in the FER frequency distribution figure.

Outputs and example forecasts

The current pre-operational version produces probabilistic rainfall forecasts for points, in the form of percentiles from 1 to 99, once a day (00 UTC run), up to day 5, for 06–18 and 18–06 UTC validity times. The operational version (available soon via ecCharts) will produce forecasts twice a day (00 and 12 UTC), up to day 10, for 6-, 12- and 24-hour accumulations.

The output can be used by forecasters in different ways. For example, a user may be interested in identifying the wettest places in the world in two days' time in order to anticipate which areas are vulnerable to flash floods. They could plot point rainfall (e.g. in mm/12h) for high percentiles, e.g. the 95th or 98th percentile, and could then say that there is 1/20 or 1/50 chance for the amount displayed to occur at a given location in a given grid box, albeit without being able to pinpoint exact locations within a grid box where extremes are most likely.

For instance, between 2 and 4 a.m. on 10 September 2017 more than 200 mm of rain fell in the vicinity of Livorno in Italy, causing widespread flooding and six fatalities. On the 98th percentile point-rainfall map for that night from two days earlier, parts of northern Italy, including Livorno, stand out as being especially vulnerable to very high totals. On the equivalent raw ensemble plot (not shown) this region does not stand out. On the basis of the point-rainfall forecast, one could have put parts of northern Italy on alert, albeit at a low risk level. Other considerations related to vulnerability (e.g. population density, topographic slopes, land use, soil moisture, flood defences) could have facilitated refinement of warning areas.

Where there is local knowledge regarding flash flood trigger levels, users may alternatively be interested in the probability of exceeding a certain threshold, as shown in the UK case study charts. This example also shows how the post-processing does not simply amount to amplifying signals of more extreme events.

Verification

A retrospective global verification for the current pre-operational system covers the period April 2016 to March 2017, for several sets of lead times (days 1, 3, 5), focusing on the reliability and resolution of the forecast ('resolution' here refers to how well the system is able to distinguish between occasions when events are more/less likely to occur).

Evaluation of the raw ensemble (ENS) shows an under-representation of small and large values. A much flatter rank histogram for the point-rainfall product is indicative of a much more reliable forecast. Finally, using the ROC area metric, the point-rainfall product exhibits greatly improved resolution. This is notably the case for higher thresholds, where the point-rainfall product at day 5 is about as good as the raw ensemble at day 1, illustrating clear added value for users.

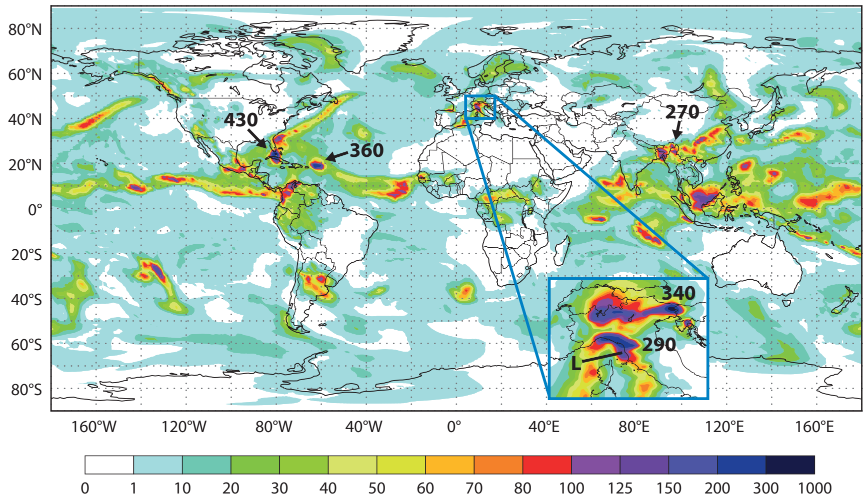

Point-rainfall forecast. 98th percentile in mm/12h from 00 UTC on 8 September 2017 for the period t+42 to t+54. The top five substantive maxima are labelled. The Florida/Atlantic peaks relate to hurricanes Irma and Jose. L points to Livorno.

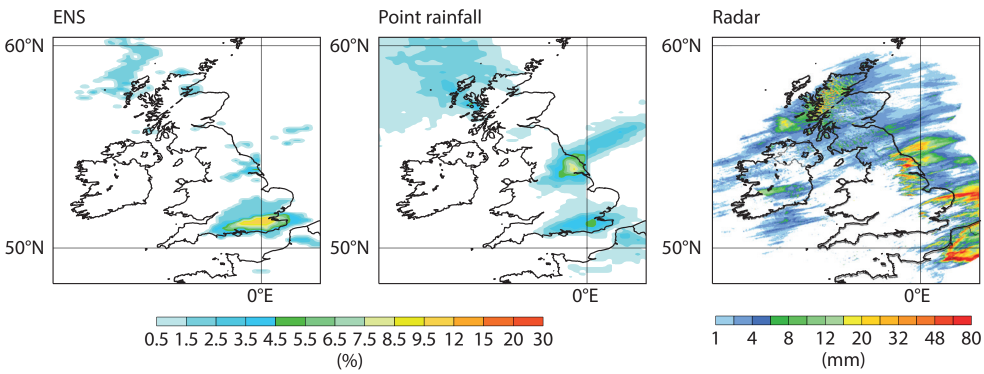

UK point-rainfall case study. The charts show the raw ENS probabilities (left) and point-rainfall probabilities (middle) of exceeding 20 mm/12h (30 to 42-hour forecasts valid 06–18 UTC 6 July 2017). The right-hand chart shows UK radar-derived totals (*source: netweather.tv*).

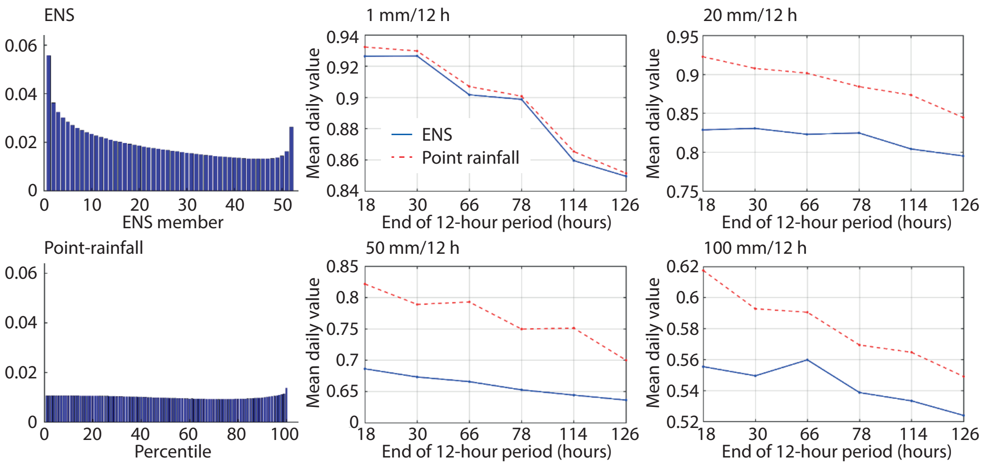

One year of point-rainfall forecast verification (global). Rank histograms (t+102 to t+114 and t+114 to t+126) for the raw ensemble where each rank is an ENS member (1 to 51 + outliers) (top left) and for point rainfall where each rank is a percentile (1 to 99 + outliers) (bottom left). A flatter histogram indicates better reliability. The other charts show the area under the ROC curve for exceeding different thresholds. Higher values indicate better resolution. The verification period is April 2016 to March 2017.

Predictions of tropical cyclones Harvey and Irma

LINUS MAGNUSSON,
IVAN TSONEVSKY,
FERNANDO PRATES

At the end of August and the beginning of September 2017, two major hurricanes, Harvey and Irma, were the first in a series to hit the Caribbean and the US. ECMWF forecasts predicted their paths fairly well. In the case of Harvey this helped to predict large amounts of rainfall over Texas. As is common in tropical cyclone (TC) forecasts, the intensity of the hurricanes was less well predicted.

Harvey

On 26 August (25 August local time) TC Harvey made landfall in Texas. Subsequently the cyclone became quasi-stationary and produced more or less continuous rainfall for three to five days. The rainfall totals reached more than 1,000 mm in the worst-affected areas around Houston, where unprecedented flooding occurred.

The cyclone formed from a tropical disturbance east of the West Indies on 18 August and propagated westward as a fairly weak system. On 22 August it made its first landfall on the Yucatán Peninsula (Mexico) and was downgraded to a tropical depression. After reaching the Gulf of Mexico, the system regained its status as a tropical cyclone. Over the next few days, the cyclone rapidly intensified and became a Category 4 hurricane before making landfall in Texas. After landfall, Harvey became quasi-stationary while gradually weakening to a tropical storm. On 29 August it moved out over the Gulf of Mexico again and a day later it made landfall for the third time, in Louisiana.

As early as 18 August, the ensemble forecast was confident about the propagation of Harvey towards southern Mexico. It also indicated that the system might enter the Gulf of Mexico. That risk temporarily decreased on 19 and 20 August when the cyclone was very weak in the central Caribbean Sea. After 21 August, the ensemble was confident that the system would turn into a tropical storm over the Gulf of Mexico but there was considerable uncertainty about where it might make landfall. Between 21 and 22 August, the risk of the cyclone becoming quasi-stationary over Texas increased, and with that came a risk of extreme rainfall. The high-resolution forecast (HRES), in particular, highlighted the risk of extreme precipitation. However, in the short range the predicted area of the worst rainfall was shifted to the southwest compared to the outcome. On 24 and 25 August the cyclone rapidly intensified. This was not captured well by the forecasts before the start of the intensification. However, the total accumulation of rainfall predicted by HRES over Texas was still in the same range as the observed amount.

Irma

TC Irma hit several countries along its path in the Caribbean. On 5 and 6 September, the Category 5 cyclone made its first landfall on some of the Leeward Islands. The first to be hit was Barbuda, followed by Saint Barthélemy, Saint-Martin/Sint Maarten and Anguilla. All these islands were crossed by the eye of the cyclone and wind gusts up to 70 m/s were reported on Barbuda. The cyclone later hit the Virgin Islands and the Turks and Caicos Islands. It also affected Puerto Rico and Hispaniola. On 8 September the cyclone hit the Bahamas. It made landfall on Cuba on 8 to 9 September as a Category 5 cyclone. Finally the cyclone made landfall on the southern tip of Florida on 10 September. ECMWF Member States with territories in the area have given positive feedback on the Centre's forecasts. Here we will focus on ECMWF's predictions for the Leeward Islands and Florida.

The cyclone formed on 30 August west of Cape Verde in the tropical Atlantic. The cyclogenesis was predicted about a week beforehand. The ensemble from 31 August showed a high risk of Irma passing the Leeward Islands six to seven days later. The ensemble was confident that the group of islands would be hit, but there was some uncertainty about the exact track. However, there were large forecast errors in the intensity and wind speed prediction.

The landfall in Florida was much more unpredictable than the landfall on the

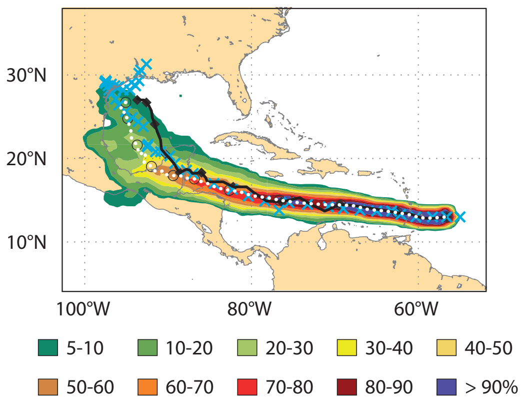

Strike probability map for TC Harvey. The chart shows the probability that Harvey will pass within a 120 km radius during the next 240 hours, according to the forecast from 18 August. The solid line is the HRES, the dotted line the ENS mean, and the crosses show the path as subsequently observed.

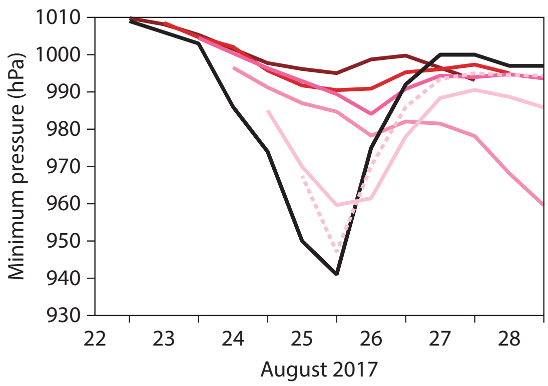

Central pressure forecasts for TC Harvey. The chart shows the evolution of central pressure for TC Harvey during the phase of rapid intensification according to 'best track' data (black) and HRES with different starting times (red).

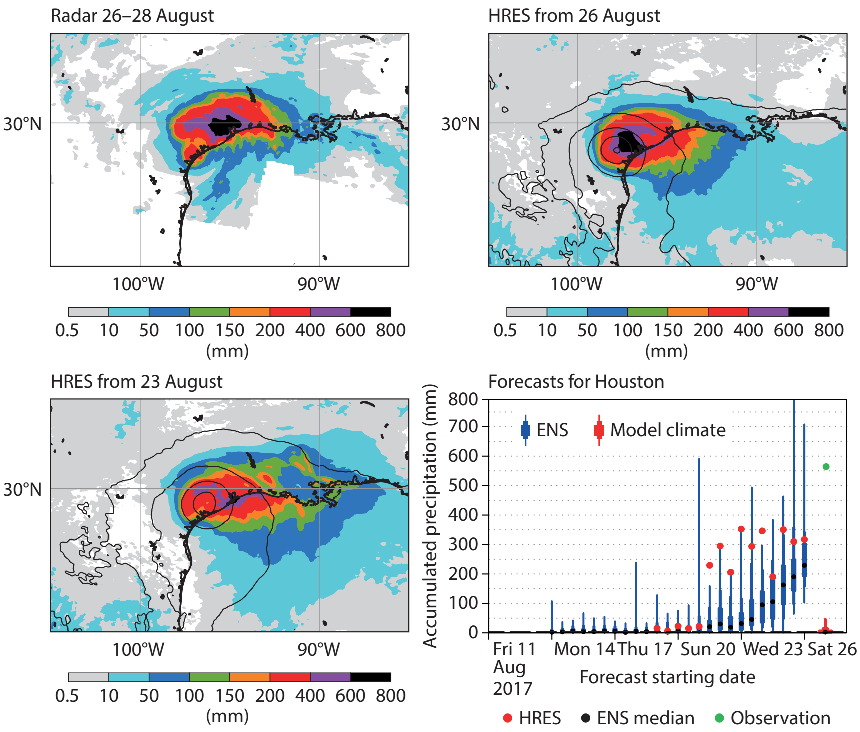

TC Harvey precipitation forecasts. The charts show accumulated precipitation in the period 26 to 28 August from the US radar network NEXRAD (top left), HRES starting from 26 August (top right) and HRES starting from 23 August (bottom left). The bottom-right plot shows the HRES (red) and ENS (blue) predicted accumulated precipitation in the period 26 to 28 August for Houston for different forecast starting dates. The box-and-whisker symbols mark the 1st, 10th, 25th, 75th, 90th and 99th percentile. Contours show mean sea level presssure as predicted for 27 August 12 UTC.

Leeward Islands. Early forecasts indicated a northward turn at some point, but the exact timing of this made a huge difference for the location where Irma would hit the US coast. Five days before the landfall, in the ensemble starting from 00 UTC on 5 September, possible landfall locations ranged from Louisiana in the west to North Carolina in the east. Later the range narrowed, but even three days before landfall, in the forecast from 00 UTC on 7 September, the tracks ranged from west of the Florida peninsula to east of the peninsula. These scenarios meant there was considerable uncertainty over where the impact of the cyclone would be strongest. On the day of the landfall, there were uncertainties in the final details concerning the strength of the cyclone and the exact landfall location. This created large uncertainty in the storm surge prediction. In the end, the cyclone hit Key West and made landfall just west of the tip of Florida. Meanwhile a storm surge caused some flooding in Miami on the east coast.

Summary

Different tropical cyclones pose different types of risks. The primary risk is the winds that damage or destroy buildings, as happened with TC Irma on the Leeward Islands. Another risk is storm surges along coasts, aggravated by waves, as for TC Irma in Florida. A slow-moving cyclone (as in the case of TC Harvey) or a cyclone that hits steep coastlines can also create extreme rainfall accumulations over land leading to potentially devastating flooding. In the case of Harvey the near-stationary nature of the cyclone was the key point to predict.

For both Harvey and Irma, ECMWF forecasts struggled to correctly predict the intensity of the cyclones. This can at least partly be explained by the relatively small scale of tropical cyclones compared to the model resolution. Other phenomena that are difficult to represent and thus limit predictability include eyewall replacements, which temporarily weaken TCs; rapid intensification; intrusion of dry air; and land interaction. These elements are the subject of intense research among tropical cyclone scientists. The processes involved are not yet fully understood and even limited-area models find it difficult to capture them.

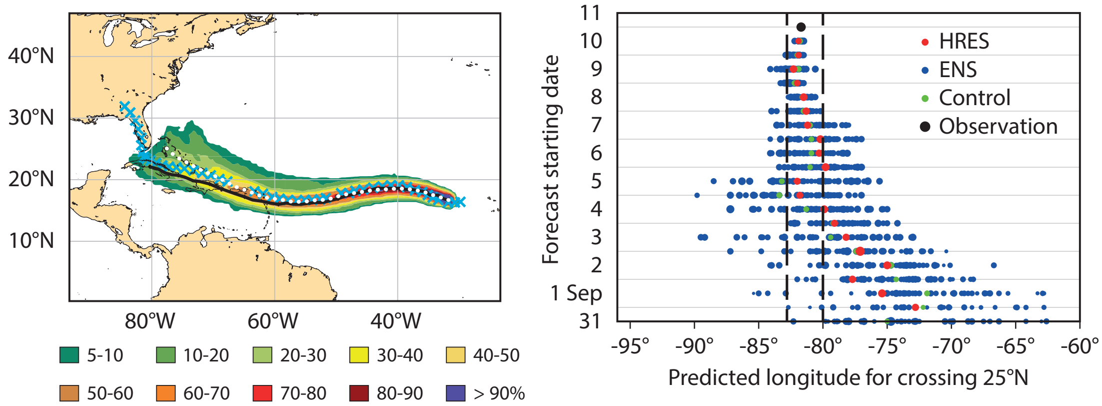

Strike probability for TC Irma. The left-hand chart shows the probability that Irma will pass within a 120 km radius during the next 240 hours, according to the forecast from 31 August. The solid line is the HRES, the dotted line the ENS mean, and the crosses show the path as subsequently observed. The right-hand chart shows the predicted longitude for crossing 25°N (latitude of Miami) in successive ENS (blue), HRES (red) and control forecasts (green). The size of each dot is scaled to the predicted strength of the cyclone at the crossing time. The black dot indicates the observed location and the dashed lines show the width of Florida.

OpenIFS users explore atmospheric predictability

GLENN CARVER

The fourth OpenIFS user workshop, held in June 2017 at the International Centre for Theoretical Physics (ICTP) in Trieste, Italy, was devoted to 'Atmospheric Variability: seasonal predictability and teleconnections'. It attracted scientists from institutes in Europe and further afield and was undoubtedly the most successful OpenIFS workshop to date.

The OpenIFS activity at ECMWF provides a supported version of the operational Integrated Forecasting System (IFS) for research and education. User workshops are an opportunity for scientists to interact, present their work with OpenIFS and learn more about ECMWF. Each meeting focuses on an active research area at ECMWF. This was the first such event to focus on seasonal forecasting.

Italian connection

The choice of location reflected increasing interest in OpenIFS from Italian research groups. ICTP was chosen for its excellent facilities and international reputation, particularly for seasonal prediction. The workshop took place over five days instead of three for previous meetings, as feedback indicated participants wanted more time for hands-on activities with the model.

Fifty participants ranging from young to senior scientists took part, twice as many as in previous OpenIFS workshops. For the first time, the workshop was also significantly oversubscribed with over ninety applicants. The meeting was opened by two keynote presentations from ECMWF. Professor Erland Källén, then Director of Research, spoke about 'Research and development at ECMWF'. He was followed by Dr Franco Molteni on 'Experimentation on extended-range prediction and multi-year variability at ECMWF'. Both talks were well received with many questions.

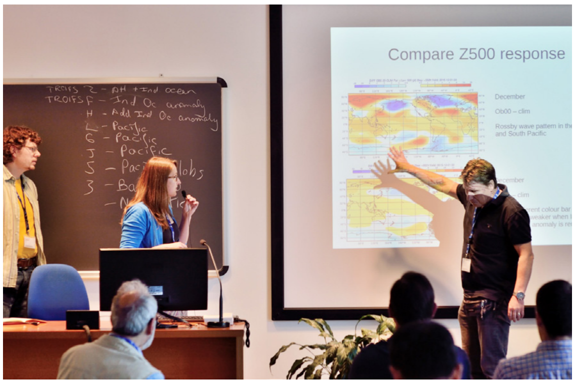

OpenIFS experiments. Participants presented and discussed their OpenIFS seasonal forecast experiments with different sea-surface temperatures. (*Photo: Filip Váňa*)

Each morning was a combination of invited and contributed presentations. With eleven invited speakers, nine contributing speakers and twenty-one poster presenters, the scientific standard was high and engaging. A range of topics on the workshop theme were presented, some based on non-ECMWF models such as the European Earth system model EC-Earth and the SPEEDY global circulation model developed at ICTP. Some speakers gave examples of using OpenIFS for research on other topics and teaching at universities. The afternoons were spent first learning about OpenIFS and then running practical experiments in teams. The last day of the workshop was devoted to team presentations and discussion of the results.

El Niño 2015/16 experiments

El Niño is a term used to describe a significant warming of the sea-surface temperature (SST) mainly in the central and eastern Pacific, which occurs irregularly every few years. This interacts strongly with the atmosphere and has consequences for global weather patterns. The strength and area of each El Niño varies, but the event of winter 2015/2016 was one of the strongest on record.

Franco Molteni designed an experiment with OpenIFS based on the 2015/16 El Niño. The OpenIFS model was run in two experiments, each with a 10-member ensemble from November 2015 to the end of April 2016; one used observed SST values, the other climatological values, a 20-year mean SST from ERA-Interim. These forecasts were provided to participants to explore the impact of El Niño by comparing a seasonal forecast with SST anomalies to one without. Ten teams of participants were then asked to design experiments in which the observed SST anomaly over an oceanic area of their choice was changed to either the ERA-Interim climatology (removing the anomaly) or to some multiple of the difference (for example doubling the anomaly).

Using Metview tools developed by ECMWF's Sándor Kertész, participants were able to alter the SST field used in these forecasts to create their own experiments and ask questions such as: what would the change in European weather be if we removed or increased the Pacific El Niño anomaly? By using ensembles of forecasts, some measure of significance could be determined for these changes.

These ensemble seasonal experiments were run from ICTP on the ECMWF high-performance computing facility (HPCF), the first time an OpenIFS workshop has made use of the HPCF remotely. Metview was used to compute

Group photo. Fifty scientists took part in the workshop. (*Photo: Filip Váňa*)

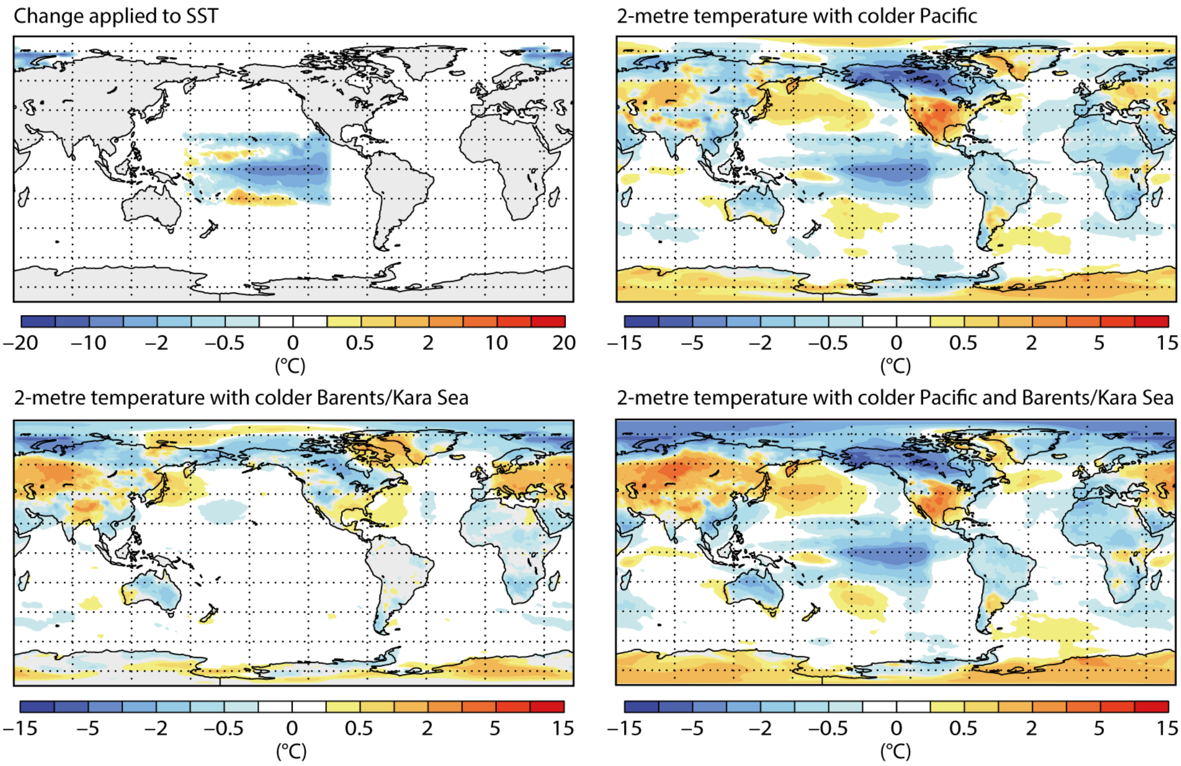

Seasonal forecast experiment with OpenIFS. The charts show the effect of combining modifications to the sea-surface temperatures (SST) in the tropical Pacific and the Barents/Kara Sea. Three separate 10-member seasonal forecasts were made in which the warm anomaly in the SST was replaced by a cold anomaly throughout the forecast. The top left panel shows the SST changes that were applied to the observed SST field. The top right panel shows the 2-metre temperature monthly mean difference between the experiment with the modified SST applied just over the tropical Pacific and the unmodified, observed SST; the bottom left panel shows the same but with the modified SST applied just over the Barents/Kara Sea region; and the bottom right panel shows the same but with the modified SST applied in both the Pacific and the Barents/Kara Sea. It can be seen that the colder SST combines to give increased warming over Eurasia. All figures are monthly means for January 2016 from forecasts from the beginning of November 2015. (*Charts: Tido Semmler and Team 3*)

monthly means from the model results, which were then transferred to the ICTP classroom for plotting using Metview. Participants were able to look at the statistical significance of the changes to the SST used in their experiments. Some teams developed their own diagnostic tools and applied them to the results, often working late into the evening!

There was a range of interesting experiments. Some teams chose to alter the tropical Pacific SST anomaly either over large areas, to look at teleconnections over Europe, or over smaller Pacific coastal areas for local impacts. Others focused on the role of the Indian Ocean, whilst two teams looked at the impact of changes to the Arctic sea-ice cover. The tools provided allowed the SST values to be altered in multiple areas, allowing for some interesting experiments looking at the interaction of different anomalies. All teams gave excellent presentations on their findings, resulting in scientifically interesting discussions on the interpretation of the results, with perhaps a few questions remaining too.

Plans for the future

The sensitivity experiments using the El Niño of 2015/16 yielded some interesting results and a follow-up publication is planned. Participants said they found the workshop informative and interesting and the opportunity to meet ECMWF scientists was much appreciated. Some participants would have liked to learn more about the model itself. However, fitting additional lectures into a very full programme would have been difficult given the large number of high-quality invited and contributed presentations. This is something to consider for future workshops.

OpenIFS workshops are important to the growing user community and will continue to be hosted at different European institutes. They also contribute to raising awareness of ECMWF and its ongoing research programme. The presentations highlighted the substantial research and development now under way using the OpenIFS model.

The workshop was organised jointly by: ICTP, ECMWF, CETEMPS University of L'Aquila, and ISAC-CNR Bologna. We gratefully acknowledge financial support from the EU ESiWACE programme, CETEMPS, ICTP and a Young Scientist Travel Award from the European Meteorological Society, which was won by Lenka Novakova from the University of Reading. The workshop would not have been such a success without the contribution and support from co-organisers Fred Kucharski (ICTP), Paolo Ruggieri (CETEMPS L'Aquila) and Susanna Corti (ISAC-CNR Bologna) as well as a number of ECMWF colleagues, particularly: Sándor Kertész for preparing the Metview tools; Filip Váňa for OpenIFS support; Franco Molteni, who suggested the practical experiment; Tim Stockdale, who provided the IFS experiment; Erland Källén; Sarah Keeley and Iain Russell. We look forward to the next user meeting in 2019.

ECMWF forecasts support Portugal wildfire response

FRANCESCA DI GIUSEPPE, CLAUDIA VITOLO, FREDRIK WETTERHALL, FLORIAN PAPPENBERGER (all ECMWF), **LOURDES BUGALHO** (IPMA)

A catastrophic forest fire in Portugal claimed more than 60 lives this summer. All the casualties were recorded in the Pedrógão Grande area, 50 km southeast of Coimbra, between 17 and 18 June. The dry thunderstorm and heatwave conditions in the region, with temperatures above 40°C, were highly unusual for the season. Moreover, relative humidity levels below 30% were conducive to the intensification of the deflagration and the spread of the wildfire, which raged out of control for several days.

The Instituto Português do Mar e da Atmosfera (IPMA) is the information provider for weather-related hazards to the Portuguese civil protection agency, which coordinates the response to emergencies. Since April 2017 IPMA has had access, on a pre-operational basis, to forecast data in near real time from the European Fire Forecast Information System (EFFIS) of the Copernicus Emergency Management Service (EMS), to which ECMWF contributes. The fire forecast products indicated the presence of extreme fire danger conditions in the area several days in advance. According to IPMA, they contributed to better planning and a fast response by crisis units.

Increasing use of forecasts

Traditionally Portugal has relied on observations to assess fire danger conditions. In summer, daily fire danger values based on the FWI (Fire Weather Index) are computed for 83 weather stations where weather measurements are routinely available and the vegetation status is recorded. The resulting fire danger rating is then extrapolated to a regular mesh with a resolution of about 1 km. This dataset provides a daily fire danger assessment for Portugal. It is clear that by using predicted conditions from advanced numerical weather prediction models, much longer-range assessments can be achieved (for example 1 to 2 weeks), enabling better planning and resource sharing within and between countries. Since 2012, IPMA has used forecast products on an experimental basis by calculating FWI values based on 72-hour forecasts of the ALADIN limited-area model. Currently it also uses ECMWF forecasts to calculate FWI values up to 72 hours ahead. In addition, in April IPMA began to receive the 10-day outlook forecast from the GEFF high-resolution forecast run. GEFF (Global ECMWF Fire Forecasting system) is the modelling component of EFFIS.

Benefits and limitations

Monitoring and forecasting are concatenated daily at Portuguese weather station locations. From IPMA's experience we know that FWI forecast skill usually degrades fast after 98 hours. For accurate information to be available at least five to seven days ahead would be a real breakthrough enabling much-improved decision management. Of course, the use of weather forecasts instead of observations means that FWI values might be affected by model biases, which may be amplified or

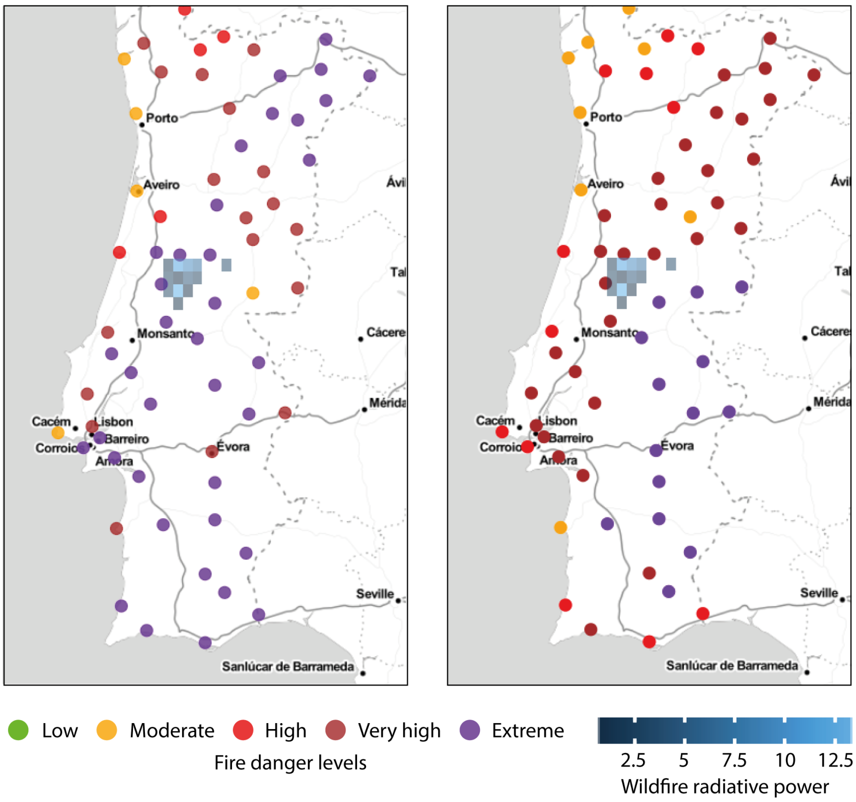

Fire danger observations and 10-day forecast. Fire danger level classification for 18 June 2017 at the weather station locations used by IPMA to monitor fire danger in Portugal (left) and predicted ten days before by the GEFF modelling component of EFFIS (right). The extent and intensity of the observed fire is shown in terms of the fire radiative budget available from MODIS through the Global Fire Assimilation system.

CaliVer

CaliVer is a library for the R statistical language that contains reproducible algorithms for the calibration, verification and visualisation of the modelling component (GEFF) of the EFFIS service. CaliVer uses the available historical dataset of FWI from ERA-Interim to provide a consistent classification of fire danger classes across countries. For more information, visit: https://github.com/ecmwf/caliver damped by nonlinear transformations in the fire model. For example, a dry bias in the model in a certain region will lead to the persistent prediction of relatively high fire danger values. If warning levels are defined on the basis of local observations (as they are in the Portuguese system), this may result in a high false alarm rate. From a computational point of view, tailoring fire danger levels to a given area and validating the performance of an early warning system based on weather forecasts is a demanding task as it requires handling large historical datasets. To support the use of fire forecast data, ECMWF has developed a freely available post-processing tool called CaliVer (Calibration and Verification) to define warning levels from model outputs. Comparing the warning levels adopted by IPMA calculated from the historical record of the observed FWI with the danger classes from the reanalysis dataset, it turns out that the two classifications have different distributions: IPMA's is more conservative for higher danger levels. As a result, the use of IPMA's levels would generate a larger number of false alarms if applied to the FWI calculated by GEFF for the EFFIS platform. In practice, observed and predicted FWI values are not directly comparable if there are large biases in the meteorological inputs. However, as the definition of specific warning levels can be regarded as a bias correction procedure, it is reasonable to assume that the fire danger classes are broadly comparable between different systems. Results for the June wildfires in Portugal show that predictions from the GEFF model signalled the very high to extreme fire danger that was observed on 18 June 2017 ten days ahead. The extended predictability in this particular case may have been boosted by comparing fire danger classes rather than FWI values.

Another advantage of the modelling approach chosen by EFFIS is that the forecasts it produces extend to regions not covered by in-situ observations. Gridded data makes it possible to perform simple impact analyses. For example, by layering maps of population or of land use, it is possible to derive the potential number of people affected by an event.

The use of weather forecasts in the management of fire events thus opens up new possibilities. This has been recognised at a wider European level, and in fact EFFIS only relies on the use of weather forecasts to monitor fire danger at the European scale. A wider socio-economic benefit study under way at ECMWF will attempt to assess the benefits of ECMWF fire forecasting in Europe. Watch this space.

Combined risk classification

The table shows how risk classes can be based on a combination of exposure (population density in this case) and hazard (fire). Population density classes range from 'Insignificant' (number of residents per km^2 < 300) to 'Extreme' (number of residents per km^2 >= 10,000). The hazard classes are taken from GEFF warning levels.

		Exposure				
		Insignificant	Minor	Moderate	Major	Extreme
Hazard	Very low	Low	Low	Low	Low	Medium
Hazard	Low	Low	Low	Low	Medium	Medium
Hazard	Moderate	Low	Low	Medium	Medium	High
Hazard	High	Low	Medium	Medium	High	High
Hazard	Very high	Medium	Medium	High	High	High
Hazard	Extreme	Medium	High	High	High	High

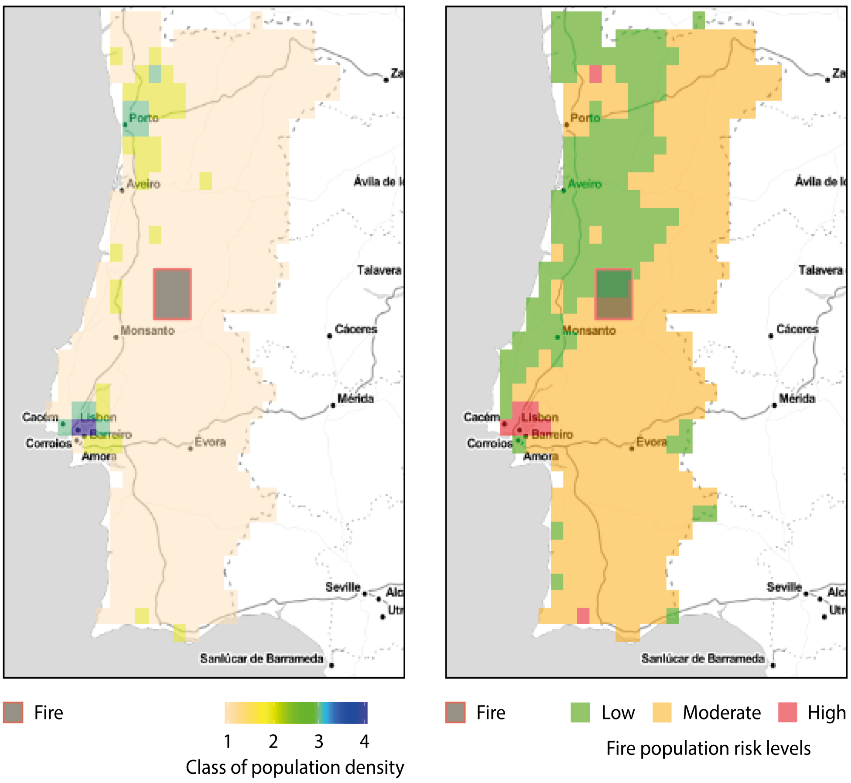

Predicted impact on the population during the event. Population density figures provided by NASA's SEDAC socioeconomic database (left) can be used to derive a risk map based on a combination of fire danger and population density ten days ahead (right).

The August 2017 heat wave in southern Europe

FRANCESCA DI GIUSEPPE, LAURA FERRANTI (both ECMWF), **CLAUDIA DI NAPOLI** (University of Reading)

At the beginning of August, unusually high temperatures, in some cases unprecedented, were recorded in a large area spanning much of the Iberian Peninsula, southern France, Italy, the Balkans and Hungary. For some days top temperatures rose above 40°C across southern Europe, exacerbating the impact of an extended drought and the lingering impact of a June heat wave.

Summer heat waves in Europe develop when there is a blocking of westerly winds and high-pressure weather conditions prevail. In early August 2017, the stability of the atmosphere caused high-pressure conditions to last for several days, resulting in extreme warming of the air and leading to unusually high temperatures and no wind.

ECMWF forecasts pointed to the possibility of unusually warm conditions well in advance. The monthly ensemble-based forecast showed a warm anomaly for 31 July to 6 August in many of the affected areas as early as four weeks ahead. The signal persisted at shorter lead times and the forecast starting on 27 July is in very good agreement with the analysis. The World Meteorological Organization (WMO) criterion for a heat wave is that "the daily maximum temperature of more than five consecutive days exceeds the average maximum temperature by 5 °C, the normal period being 1961–1990". Using this definition, the operational analysis shows an extensive heat wave over parts of southern and southeast Europe between 1 and 10 August. In places, temperatures rose well above 5°C when compared to the ERA-Interim mean climate (1979–2016). The extent of the heat-wave conditions was very well captured by ECMWF's high-resolution 10-day forecast (HRES) even though there was a slight overestimation of their duration in some areas. The heat-wave duration is shown for the forecast starting on 1 August as an example. However, the signal was present for a few days before. This is not surprising as temperature is one of the most predictable meteorological surface variables. For example, the

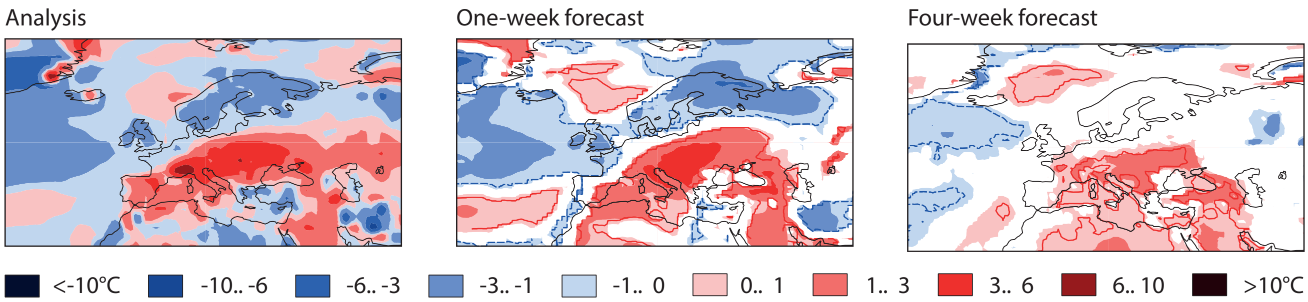

Temperature anomaly forecasts. Weekly mean forecasts of two-metre temperature anomalies for the period 31 July to 6 August 2017 starting on 6 July (right) and on 27 July (middle). The analysis is shown in the left-hand panel. The anomalies are calculated relative to a 20-year model climatology. Areas where the ensemble forecast is not significantly different from the climatology are left blank.

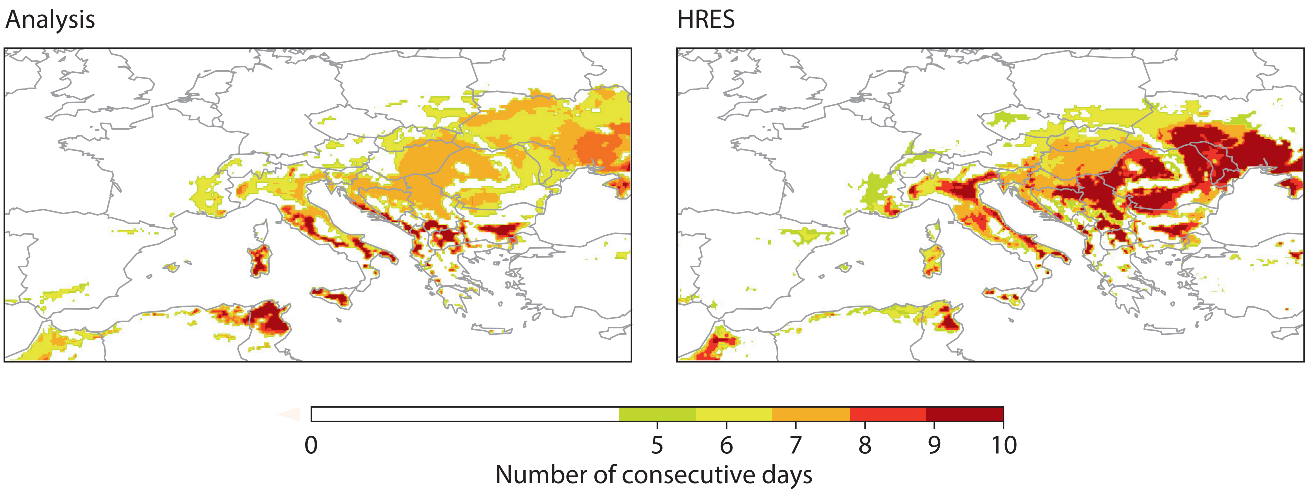

Heat-wave conditions. The maps show locations where heat wave conditions occurred and the number of consecutive days for which they occurred according to the operational analysis for 1 to 10 August 2017 (left) and according to ECMWF's high-resolution forecast issued on 1 August for the period up to 10 August.

high-resolution forecast for maximum two-metre temperature during the first ten days in August was in almost perfect agreement with observations from the high-resolution network of European in-situ data collected through EFAS (not shown).

Extreme heat can lead to a variety of health risks, including dehydration, hyperthermia, and even death, especially during sustained periods of high temperatures. The thermal health hazard experienced during the August 2017 heat wave can be assessed using the Universal Thermal Climate Index (UTCI). UTCI is an index representing the human body's discomfort to thermal stress. It is elaborated as an equivalent temperature via an advanced model of human thermo-regulation that, coupled with a clothing insulation model, estimates the effect of wind speed, water vapour pressure and short- and long-wave radiant fluxes on human physiology. Maps of predicted UTCI show high values in many of the areas affected by the heat wave (not shown). Such maps are computed at the European scale using ECMWF ensemble forecasts as part of the activities of the EU-funded ANYWHERE project.

In summary, forecasts of two-metre temperature can provide valuable information ahead of and during heat-wave episodes. Impact-oriented products, such as heat-wave duration, maximum temperatures reached over a given period of time and UTCI maps, can provide supplementary information to meet the requirements of different users.

ECMWF supports field campaign in the Azores

MAIKE AHLGRIMM (ECMWF), **SCOTT M COLLIS** (Argonne National Laboratory), **MICHAEL JENSEN, JIAN WANG** (both Brookhaven National Laboratory)

Real-time ECMWF forecast products, including from the Copernicus Atmosphere Monitoring Service (CAMS), have aided aircraft flight planning in a field campaign to gather aerosol and cloud observations. The campaign was hosted at the US Department of Energy's Atmospheric Radiation Measurement (ARM) site on Graciosa Island in the Azores. ECMWF data proved a key decision aid in this successful mission, which collected a rich dataset for improving the representation of shallow, radiatively important clouds in weather and climate models.

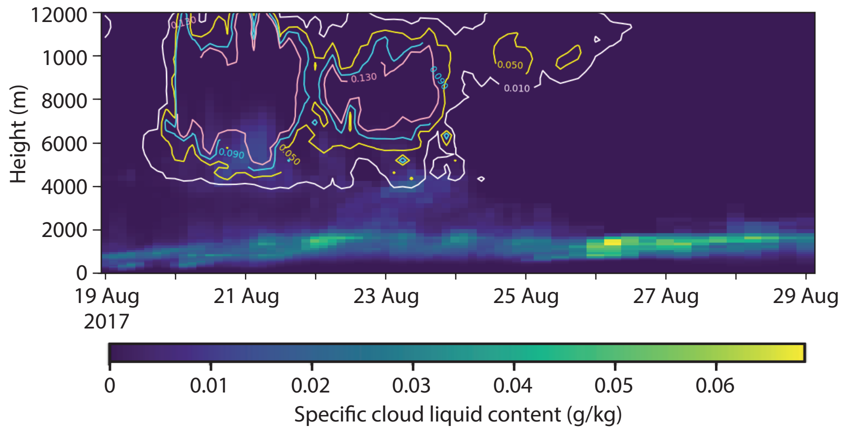

Example forecast guidance chart. This time–height plot of cloud liquid water content (shading) and ice liquid water content (contours, in g/kg) was produced with ECMWF forecast data over the Graciosa domain.

Aerosol and cloud experiments

The ARM facility maintains instrumented research sites in various climatic regions around the globe. Ground-based remote sensing equipment, such as zenith-pointing and scanning cloud and precipitation radars and lidars, together with more traditional meteorological and radiation instruments, are used to observe the atmosphere, clouds and aerosol in detail to improve our understanding of radiative processes in the atmosphere. ARM also has a Gulfstream 159 aircraft to enable the in-situ sampling of clouds and aerosols to complement the remotely sensed observations.

ECMWF and ARM's mutually beneficial relationship goes back two decades. ARM's observations and related research contribute to improved representations of radiation and microphysical processes in the Integrated Forecasting System (IFS). ECMWF, in turn, supports the permanent ARM sites and field campaigns with analysis and forecast products.

In July 2017, the ARM site on Graciosa Island hosted the first part of the Aerosol and Cloud Experiments in the Eastern North Atlantic (ACE-ENA) field campaign led by principal investigator Jian Wang. The campaign gathered comprehensive observations of marine boundary layer (MBL) clouds and aerosol at the ENA site. MBL clouds cover vast areas of the oceans and have a strong impact on the Earth's radiation budget. There are still many uncertainties in how to represent the characteristics of MBL clouds well in atmospheric models, and they remain a leading cause of uncertainty in the prediction of future climate. It is thought that drizzle and the number of cloud condensation nuclei play an important role in determining the cloud structure and cover.

While the ground-based remote sensing instruments can characterise some of the cloud and drizzle properties, they cannot measure aerosol and cloud particles directly. During ACE-ENA, the Gulfstream aircraft, equipped with a selection of aerosol and cloud probes, obtained crucial in-situ observations around and within the clouds to characterise typical aerosol and cloud properties. These can be used to verify remote sensing retrievals and to explore how near-surface aerosol observations relate to those aloft.

Optimising flight times

Real-time ECMWF forecast products, including from CAMS, supported aircraft flight planning by helping to identify cloud and aerosol scenarios of interest. Forecast products were downloaded and ingested in standard formats using a Python-based framework. A number of additional diagnostics were derived and a series of visualisations produced, including time series, skew-T thermodynamic diagrams and two-dimensional colour meshes. The charts were then disseminated to the campaign team.

These products were critical for making optimal use of the campaign resources. The Gulfstream aircraft could not fly on every consecutive day due to prescribed pilot rest periods. Understanding which days were suitable for the science objectives was essential to take full advantage of the available flight time. Given the forecast skill and resolution over the region, data from ECMWF provided the necessary information to anticipate good (or poor) conditions for sampling single-layer marine stratocumulus, making it a key decision aid.

ECMWF data also helped to determine the direction in which the clouds were moving. Scanning cross-wind allows the radar to systematically sample the clouds as they are advected overhead. The aircraft was flown along the scanning direction of the radar in order to sample the same cloud volume. Validation of remote sensing retrievals with the in-situ observations from the short field campaign will provide confidence in the retrieval products and thus extend the utility of the long-term record of ground-based observations. The campaign was a great success, with 20 missions flown and a rich dataset collected for improving the representation of shallow, radiatively important clouds in weather and climate models.

A second part of the field campaign will take place in early 2018 to sample cloud and aerosol conditions more typical of the winter season. Past observations from the ENA site have already contributed to an improved treatment of boundary layer cloud processes in the IFS and those obtained during ACE-ENA promise to address some of the questions yet unanswered.

Scientific exchange boosts calibration effort

JAN BECKER (DWD)

From 18 June to 14 July 2017, Jan Becker from Germany's national meteorological service (DWD) visited ECMWF as part of a scientific exchange focused on the calibration of forecasts. Here he explains what the visit achieved.

At DWD we need a large amount of historical forecast data stored in the Meteorological Archival and Retrieval System (MARS) at ECMWF in order to calibrate ECMWF ensemble forecasts (ENS) by means of Model Output Statistics (MOS). MARS experts at ECMWF gave valuable advice on the optimisation of DWD's management of MARS requests to retrieve this data. As a result, the data transfer to DWD has been accelerated significantly. An operational version of MOS-calibrated ECMWF ensemble forecasts (00 UTC run) will soon be available at DWD and the production of a corresponding version for the 12 UTC run is envisaged.

The visit also provided an opportunity to examine the MARS hardware. This enabled me to see first-hand how the conceptual redesign of the request management leads to significantly faster access to the data.

Jan Becker. His visit to ECMWF has helped to optimise DWD's retrieval of forecast data for calibration purposes.

Since ECMWF's Integrated Forecasting System (IFS) is enhanced continuously, archived historical ensemble forecasts are based on different IFS cycles. However, changes in the models used to produce the historical ensemble forecast dataset typically cause a deterioration of forecast quality in the resulting MOS forecasts. This loss in forecast quality could probably be prevented or at least attenuated if suitable re-forecasts based on recent IFS cycles (ideally the current operational one) were used for MOS development instead of the historical forecasts mentioned above. I investigated locally with ECMWF experts how ensemble re-forecast data available at ECMWF could be used to set up an operational MOS system at DWD for ECMWF ensemble forecasts. A gap analysis raised a number of issues:

- the need for a scheduling concept to harmonise the temporal calculation scheme of (re-)forecasts at ECMWF with the schedule for the development and operational use of corresponding MOS equations at DWD
- further requirements concerning ECMWF ensemble re-forecasts, in particular the provision of 3-hourly re-forecast data
- further open questions, such as how the number of ensemble members (used for MOS-development in the first instance) affects the quality of resulting ENS-MOS forecasts, and whether replacing the data from 50 ensemble members with data from 10 ensemble members in MOS development is reasonable from a statistical point of view.

My current work focuses on the conceptual development and technical implementation of an optimised statistical combination of probabilistic forecasts emerging from individual MOS-calibrated NWP forecasts. The MOS-calibrated ECMWF ensemble forecasts will be used as important input for this optimised statistical combination of forecasts. In a talk at ECMWF, I was able to report on progress made in this DWD project. Overall the visit was very successful and took place in an exciting and stimulating environment conducive to new ideas.

Progress with running IFS 4D-Var under OOPS

STEPHEN ENGLISH, DEBORAH SALMOND, MARCIN CHRUST, OLIVIER MARSDEN, ALAN GEER, ELIAS HOLM, SÉBASTIEN MASSART, MATS HAMRUD (all ECMWF), ROEL STAPPERS (Met Norway), RYAD EL KHATIB (Météo-France)

Work to implement the Object-Oriented Prediction System (OOPS) as a new 4D-Var data assimilation framework for ECMWF's Integrated Forecasting System (IFS) is making good progress, with first OOPS-IFS test results showing good performance. OOPS was conceived at ECMWF by Yannick Trémolet and Mike Fisher as a unified, easy-to-use framework for running different variational data assimilation formulations with a variety of forecast models. It will replace the control layer of data assimilation code that has supported ECMWF for the past 20 years. OOPS is an international effort, involving major input from ECMWF, Météo-France and the HIRLAM-ALADIN community.

The OOPS code is available under an Apache-2 licence. This enables wider collaboration with other centres. The collaboration with the European numerical research centre CERFACS has already been valuable, and there are prospects for working closely with the US Joint Center for Satellite Data Assimilation (JCSDA). Academic partners also benefit from the simplified models in OOPS to test new data assimilation approaches. It is then relatively straightforward to test ideas demonstrated in this way with the full system, in collaboration with ECMWF.

Benefits of OOPS

OOPS is an abstract control layer that can manipulate elements of the data assimilation system without needing to know their model-specific implementation details. OOPS includes simplified models (Quasi-Geostrophic and Lorenz-95) to enable early testing of data assimilation algorithms. This enabled the early demonstration of new concepts such as the saddle-point formulation of weak-constraint 4D-Var. OOPS-IFS will bring some significant benefits:

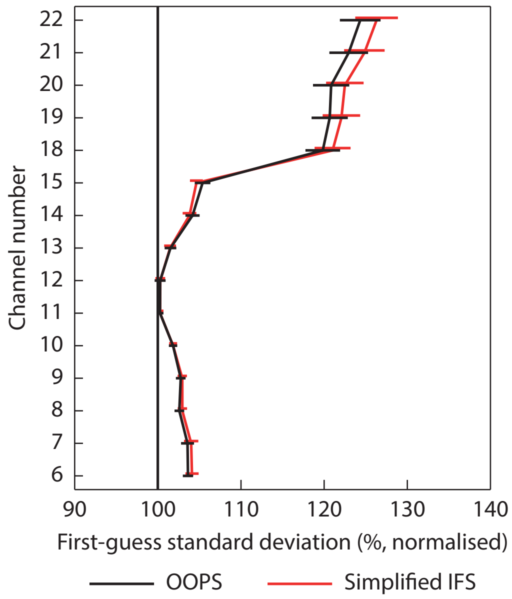

OOPS performance experiment. The chart shows the normalised change in fit of ATMS observations to the first-guess equivalents using T255/T159 OOPS and a simplified IFS, compared to the full IFS, which is represented by the 100% line. The simplified IFS uses the same switches as currently available in OOPS (No VarQC, No Wave model, No sink variable for surface temperature, No JC, No observation error correlations) and uses reduced observations: only conventional (no AIREP), ATOVS and ATMS. Channels 6–15 are sensitive to temperature from the mid-troposphere to the upper stratosphere; channels 18–22 to humidity from the lower to the upper troposphere.

- it will be easier to develop and test alternative minimisation algorithms
- it will be possible to test approaches such as the saddle-point formulation with a full system
- OOPS will provide a common framework for the development of coupled data assimilation
- OOPS reduces interdependencies in the code that make it hard to change one area without causing unexpected impacts elsewhere
- the multi-incremental assimilation will be run as a single executable, reducing I/O costs.

Towards implementation

Interfacing OOPS with the IFS, including atmospheric 4D-Var and NEMOVAR ocean 3D-Var, has necessitated significant refactoring of the Fortran code. The major goal in 2016 and 2017 has been to enable the IFS 4D-Var system to be run from OOPS with the same performance and capabilities as in operations, so that we can replace the current IFS control layer with OOPS. This is a major undertaking due to the complexity of the IFS: in addition to the core 4D-Var algorithm, the IFS has many essential components which must work consistently under OOPS. Late in 2016, a highly simplified 4D-Var system was working under OOPS.

Since then progress has been very encouraging. In 2017 many more elements of the full 4D-Var system have been added. During this work some issues in the IFS itself have been discovered, notably an error in the virtual temperature conversion in the thermodynamic balance operator of the background error covariances. A correction for this has been applied in the IFS with a notable positive impact. In this way the OOPS development is already benefiting operational scores.

However, some elements remain to be implemented, including the variational quality control (VarQC), second level preconditioning, the 'sink' skin temperature control variable needed for radiance assimilation, and correlated observation errors. These features are needed to assimilate the wide variety of observations required to produce high-quality initial conditions. Without them OOPS can only run with a limited set of observations.

Initial tests encouraging

In order to test OOPS-IFS, it has been compared to both the full IFS and a simplified IFS, where data assimilation components and observations not yet available under OOPS are switched off. Here, only conventional observations (excluding aircraft) and the important satellite radiances from AMSU-A and ATMS (mostly giving temperature information) are assimilated. Comparing OOPS to the simplified IFS allows us to assess whether its current capability is working well and how close we are to an operational level of skill.

The chart shows that OOPS produces a model short-range forecast that does not fit observations from the ATMS satellite instrument as well as the full IFS. This is to be expected as not all components of the full system are yet in place. However, it fits ATMS slightly better than the simplified IFS. Other observation fits and forecast impacts have also been examined, albeit for short periods. They lead to the same conclusion. Other data assimilation diagnostics are giving good results with OOPS, including the adjoint test and the minimisation and convergence of 4D-Var. The conclusion is that technically OOPS-IFS is performing similarly to the IFS when run in the same configuration, but that the remaining components that are still lacking in OOPS are critical to its pre-operational readiness. It is planned to implement these remaining components into a test branch of the IFS in 2017 with the intention of inclusion in IFS Cycle 46r1.

The target is that OOPS-IFS will be operational in 2019.

How to deal with model error in data assimilation

JACKY GODDARD, PATRICK LALOYAUX, SIMON LANG, MARTIN LEUTBECHER

ECMWF is developing an updated 4D-Var data assimilation system for its Integrated Forecasting System (IFS) which takes into account model error in estimating the initial conditions at the start of a forecast. This 'weak-constraint 4D-Var' has recently been implemented for the stratosphere but further work is needed before it can be made operational for the troposphere.

Strong-constraint 4D-Var

Strong-constraint 4D-Var was implemented at ECMWF in 1997 to produce a more accurate and physically consistent estimate of the state of the atmosphere at the start of a forecast. This method was designed to optimally blend information from the observations and the model in the presence of random (zero-mean) errors.

In reality many conventional and satellite observations contain systematic errors due to instrument configuration or approximations in radiative transfer calculations. To take into account these biases in the observations, a variational bias correction scheme (VarBC) was embedded inside the 4D-Var system. The scheme estimates the instrument biases by finding corrections that minimise the systematic observation departures.

Strong-constraint 4D-Var relies on the assumption that the numerical model's representation of the evolution of atmospheric flow is perfect. The error in the model trajectory thus depends only on the background (short-range forecast) error. As data assimilation processes have advanced, it is no longer possible to ignore the model error which accumulates during the 12-hour assimilation window due to inaccurate surface forcing, simplified representations of moist physics and clouds, and various other imperfections. To address this issue, ECMWF has been developing a weak-constraint 4D-Var system where the model error is explicitly taken into account in the data assimilation.

Weak-constraint 4D-Var

The perfect model assumption is relaxed in weak-constraint 4D-Var by adding a correction term in the model integration to account for the different sources of model error. The 4D-Var control variable is augmented by this correction term and a corresponding term is added to the cost function which constrains model error in accordance with its covariance statistics (see box). In order to account for systematic model error and to make the implementation affordable on today's supercomputers, the model error is assumed to be constant for the 12-hour assimilation windows. With this assumption the size of the control variable almost doubles compared to strong-constraint 4D-Var.

Weak-constraint 4D-Var has been evaluated against strong-constraint 4D-Var. Experiments have shown that in the current setup the model error forcing needs to be restricted to the stratosphere above 40 hPa to avoid the erroneous interpretation of aircraft observation error as model error.

GPS Radio Occultation measurements are extremely valuable to assess improvements in analyses and short-range forecasts for the stratosphere as they are considered to be bias-free. They are based on analysing the bending caused by the atmosphere along paths between a GPS satellite and a receiver placed on a low-Earth-orbiting satellite. As illustrated in the first figure, such measurements show that weak-constraint 4D-Var produces an analysis and a background with a smaller systematic error compared to strong-constraint 4D-Var. Similar improvements have been found with other instruments, including AMSU-A and radiosondes.

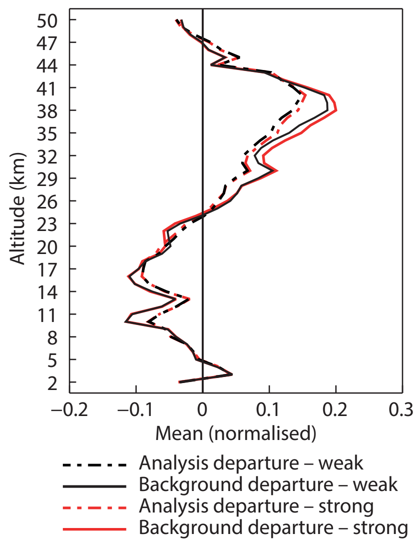

Weak- and strong-constraint 4D-Var compared. The plot shows analysis and background mean departures with respect to GPS-RO measurements for weak-constraint 4D-Var and strong-constraint 4D-Var over the period 1 January 2016 to 30 April 2016. Weak-constraint 4D-Var produces a better analysis with a smaller bias in the stratosphere.

Model error covariance matrix

Each term in the 4D-Var formulation requires the specification of an error covariance matrix to describe the error statistics of the different sources of information. Therefore, weak-constraint 4D-Var requires a model error covariance matrix in addition to the classic background and observation error covariance matrices. This model error covariance matrix describes how fast the model error can change between assimilation cycles and how the model error between different levels is correlated.

A large set of samples is required to generate covariance statistics for model errors. It is not possible to explicitly generate a sample of model errors, so instead a proxy has to be used. To create this proxy set of statistics we have run ECMWF ensemble forecasts without initial perturbations. In these runs, members diverge from each other solely due to the stochastic representation of model uncertainties, which is used operationally in the ensemble forecasts. The differences between members after 12 hours of model integration are used to construct the model error covariance matrix as they provide an estimate of the integrated effect of model error over 12 hours.

Results from operations

Weak-constraint 4D-Var was made operational for the stratosphere in IFS Cycle 43r1 launched in November 2016. Very few diagnostics exist at the moment to quantify the actual model error and to monitor the performance of weak-constraint 4D-Var. We present here a first attempt where the model error correction term estimated by the weak-constraint 4D-Var system is compared to the forecast error computed as the difference between the analysis and the forecast after 12 hours. Although the forecast error is not only due to model error but also due to the analysis error, we expect weak-constraint 4D-Var to estimate some of the systematic features of the forecast error. The model error correction term is represented in the top panel of the second figure as it has evolved since its introduction into operations. As expected, its evolution is slow, focusing on the representation of systematic errors in the model, and there is no model error estimation below 40 hPa as the weak-constraint 4D-Var is active only in the stratosphere. The forecast error in the bottom panel of the figure shows more daily variability. However, the systematic signals present, mainly between 5 hPa and 1 hPa, are captured by weak-constraint 4D-Var, which estimates a model error correction term to correct them.

Outlook

The weak-constraint formulation is an ongoing development at the cutting edge of 4D-Var. It reduces analysis and background biases in the stratosphere and improves the fit to satellite and radiosonde observations. Characterising the statistical properties of model error is one of the main current challenges. Another important aspect is also to disentangle correctly the different sources of error between observations and the model. At the moment this prevents us from activating the weak-constraint formulation in the troposphere, where aircraft observation error can be erroneously interpreted as model error.

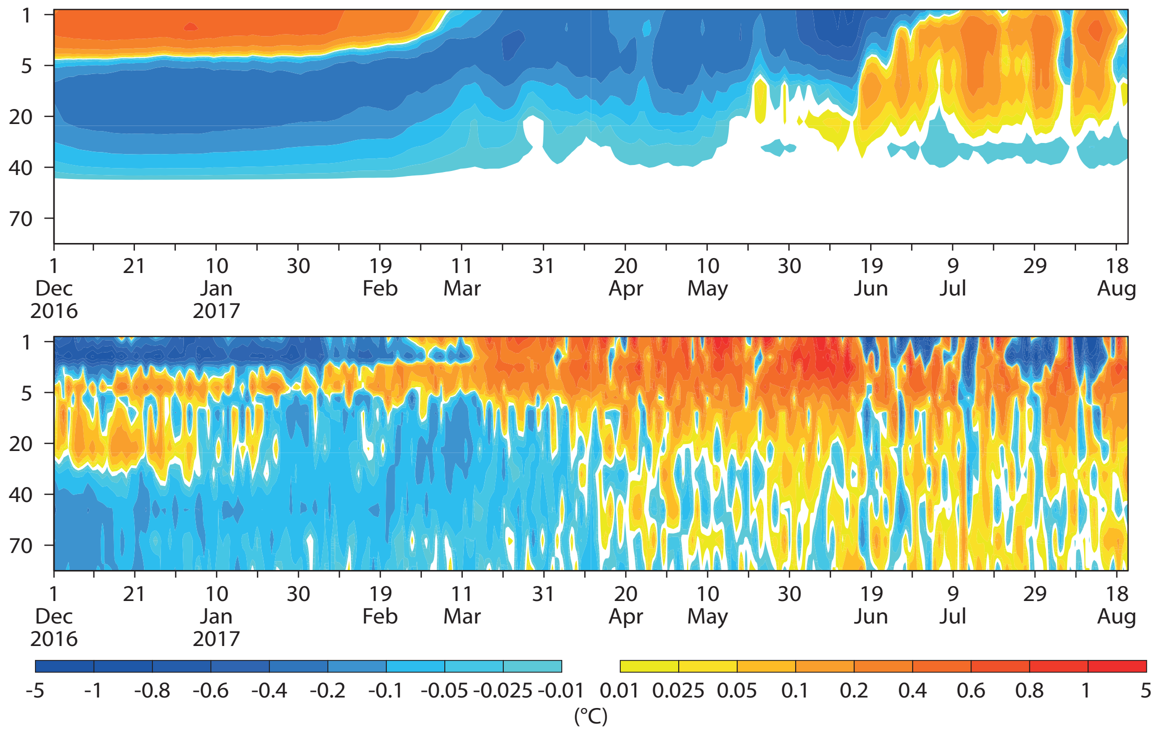

Evolution of model error correction term and forecast error. The charts show time series of the temperature model error correction term estimated by the weak-constraint 4D-Var system (top) and the temperature forecast error (bottom) after 12 hours over the south polar region. Positive (negative) model error correction terms estimated by the weak-constraint 4D-Var system correct for the systematic negative (positive) forecast errors, especially in the upper stratosphere between 5 hPa and 1 hPa.

Copernicus users rate services highly

KARL HENNERMANN, ANABELLE GUILLORY, XIAOBO YANG

A survey has found high levels of satisfaction among users of the EU-funded Copernicus Atmosphere Monitoring Service (CAMS) and the Copernicus Climate Change Service (C3S) run by ECMWF. It has also helped to identify a number of areas where improvements can be made. The survey was carried out by the Copernicus User Support Team in collaboration with the European Commission. Between 13 June and 8 July 2017, 254 CAMS users and 1,298 C3S users took part in the survey. The overall satisfaction rating is 87.5% for CAMS (an average rating of 3.5 out of 4) and 85% for C3S (3.4 out of 4).

CAMS

About half of users are academics and researchers, the others are divided equally between the private and the public sectors. Two thirds are based in Europe. User satisfaction is consistent across sectors and locations.

Most users find CAMS products and services useful. However, user uptake and satisfaction vary significantly across products and services. The most popular services are the provision of global atmospheric composition data, European air quality data, and solar radiation data. These services have the highest numbers of users, are rated as the most useful, and receive the highest satisfaction ratings. The user demographics, usage patterns and satisfaction ratings reflect the fact that CAMS started as a science programme, while also demonstrating that CAMS has established a firm foothold in the private and public sectors. The survey identified four main areas for improvement:

- the footprint of CAMS in Eastern Europe
- the attractiveness of CAMS to public-sector users
- data access tailored to users' specifications
- more detailed scientific documentation.

What users say about CAMS

"CAMS is amazing and enables researchers in atmospheric and earth sciences the world over."

"The data has really helped to conduct some analysis that would not have happened without it."

"For me it was actually rather troublesome to obtain data."

"The parameters I have used (aerosol) perform very well in comparison with other models. It is extremely useful to have these tools available."

"The ability to request more than two download streams would be very helpful."

"The people at the helpdesk are very helpful!"

What users say about C3S

"We are waiting for ERA5. Pre-release support has been helpful."

"It would be important to have access not only to charts but also to numerical values."

"I was only aware of ERA-Interim products, not any others. However, I am very satisfied with ERA-Interim."

"It would be great if the seasonal forecast data were freely downloadable."

"C3S is still developing its services, we'll have to wait and see. WebAPI is considered very useful, especially for business but it's not user-friendly."

"Allow data processing online so we don't have to download large amounts of data and then analyse them locally."

"More training, like summer schools and workshops."

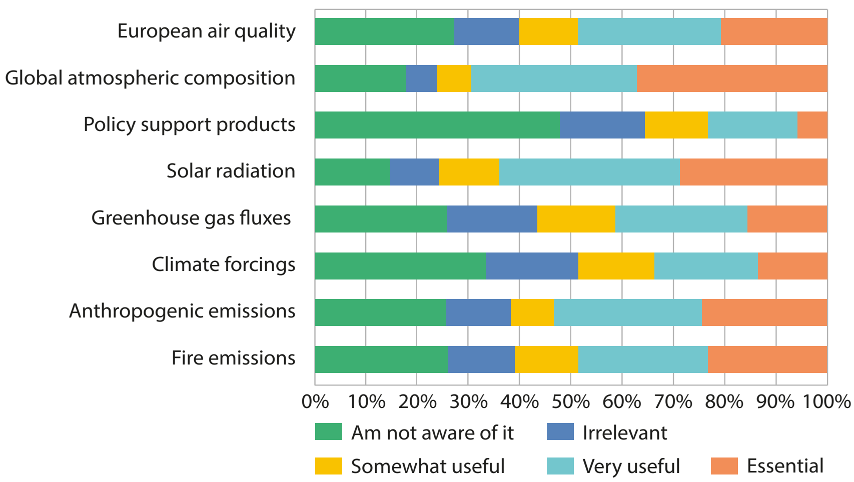

Usefulness of CAMS products and services. Many users find CAMS very useful or essential but there is considerable variability across products and services.

C3S

C3S serves a global user community: 60% of survey participants are based outside Europe. Globally three quarters of respondents are academics and researchers, and within Europe two thirds work in this sector. Other users are equally divided between the private and the public sectors.

The perceived usefulness of C3S products and services varies significantly: the ERA-Interim climate reanalysis (an ECMWF dataset for which the Copernicus User Support Team provides a technical support service) remains the most popular dataset and the support service provided receives the highest satisfaction ratings. Other products included in the survey are primarily monthly charts, plus some pre-release data and services. Many features suggested by respondents are already planned for implementation in upcoming products and services, including the ERA5 climate reanalysis, seasonal forecasts and the Climate Data Store (CDS), indicating that C3S is well aligned with user requirements.

The survey identified three main areas for improvement:

- the footprint of C3S in Eastern Europe
- the attractiveness of C3S to private- and public-sector users
- access to future products tailored to users' specifications.

The survey highlighted that some users were having difficulties downloading large amounts of data with the current infrastructure. The upcoming Climate Data Store (CDS) and the toolbox associated with it will play a vital role in attracting more users to the service. For details on the CDS, see ECMWF Newsletter No. 151, pp 22–27.

Users of CAMS as well as C3S are mostly satisfied with supporting services, in particular with the available data access mechanisms and with the Copernicus service desk.

The full survey reports are available on the CAMS (http://atmosphere.copernicus.eu) and C3S (http://climate.copernicus.eu) websites under 'Help & Support'.

The Hermes service for scalable post-processing

FLORIAN RATHGEBER,
TIAGO QUINTINO,
BAUDOUIN RAOULT

ECMWF has developed a distributed computation service for scalable post-processing, called Hermes. The new service will support ECMWF's Scalability Programme by bringing computations closer to the data, thus saving bandwidth and reducing the need for client-side processing power.

Motivation

Users have been requesting data from ECMWF's Meteorological Archiving and Retrieval System (MARS) for over 30 years. Internal users can access MARS via a command-line client, which pulls data from the archive to the user's workstation and performs any requested post-processing locally. External users, on the other hand, frequently use the MARS Web API, where the post-processing is carried out on one of the machines serving ECMWF's web infrastructure and the processed data is sent to the user via HTTP.

With increases in resolution and in light of the strategic importance placed on ensembles, transferring entire fields to the client for post-processing may involve large volumes and wastes bandwidth, especially when extracting only a small sub-area or a single point from a long time series. More control over where the post-processing happens is required since clients may not have the necessary processing power or the post-processing may take too long.

ECMWF's Scalability Programme launched in 2014 recognises the need to bring computations closer to the data. As part of this programme, the Hermes project has developed a distributed computation service to complement MARS while using a similar architecture. Hermes has focused on post-processing fields using the new interpolation package MIR. The successful conclusion of the project means that the Hermes package can be part of future data processing system developments at ECMWF.

Hermes architecture

Hermes employs a broker–worker model of distributed computing, in which a client sends a request to a broker, which queues all incoming requests for dispatch to a worker. All communication is done via the Transmission Control Protocol (TCP) using specially encoded messages. The system is transactional and automatically recovers from failures by resubmitting any request that failed to be processed, e.g. because a worker node crashed or a network connection was lost. Workers are 'stateless', which means that they process any work item independently. They can connect or disconnect from the broker at any time, allowing the system to be scaled up and down dynamically based on load and available computational resources. Hermes is not tied to any particular architecture, network topology, interconnect or even a shared file system and has been successfully deployed on desktops, ECMWF's high-performance computing facility, a dedicated cluster and in a cloud service.

Hermes leverages much of the software stack developed and maintained at ECMWF. By building on established MARS technology, Hermes

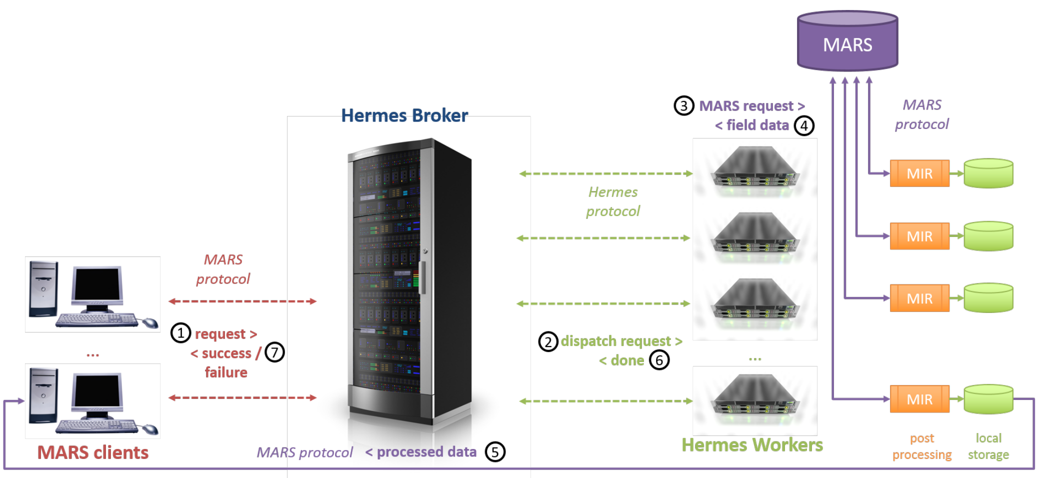

Hermes workflow. The diagram shows the workflow of processing a MARS request sent by a client (1) to the Hermes broker and dispatched by the latter to a worker (2), which retrieves the input field data from MARS (3/4), calls MIR for post-processing, sends the processed data back to the client (5) and finally notifies the broker (6) and the client (7).

is able to provide the same Quality of Service (QoS) as MARS, and similar resilience characteristics.

Hermes as an interpolation service

To address the need for scalable post-processing outlined above, initially a prototype was developed where Hermes performs an interpolation service. Support for the MARS protocol (client and server) has been added, allowing the Hermes broker to fully integrate with the MARS server and to respond to requests from the widely deployed MARS client that users are familiar with. As illustrated in the diagram, the crucial difference to the previous setup is that any post-processing is performed server-side by the Hermes worker and only the processed data is sent back to the user. This cuts bandwidth requirements, brings the processing closer to the data and enables data delivery to thin clients (such as Python).

Hermes's first deployment will be together with MARS, to serve interpolated ERA5 climate reanalysis data for the Climate Data Store, where Hermes workers will use the new interpolation package MIR (ECMWF Newsletter No. 152, pp 36–39) to process fields. The workers will be collocated with MARS data movers to minimise unnecessary data movement and optimise resource usage.

WGNE project compares tropical cyclone forecasts

JUNICHI ISHIDA (JMA), **NILS P WEDI**

The accuracy of tropical cyclone (TC) track forecasts produced by different centres is compared regularly as part of a project under the Working Group on Numerical Experimentation (WGNE). The Japan Meteorological Agency (JMA) has carried out the comparison and verification work since 1991. This intercomparison has helped to validate global models in the tropics and subtropics. The results presented here demonstrate a steady increase in the ability of global models to predict TC positions. Improvements are less pronounced for TC intensity.

Scope and method

The first forecasts considered in the project were for TCs in the western North Pacific basin produced by three global models from 1991. For later years the verification extends to all ocean basins where TCs regularly occur. Today more than ten global models regularly participate in the project. In recent years, the project has been extended to include the verification of intensity forecasts and forecasts by regional models.

The verification carried out by JMA includes annual average TC position forecast errors, systematic biases including in the along- and cross-track directions, and systematic errors common to all or most global models (e.g. forecast busts). Each participating numerical weather prediction (NWP) centre provides a gridded dataset of the 6-hourly mean sea level pressure field. The horizontal resolution of the gridded dataset differs from one NWP model to another. A minimum pressure location in the mean sea level pressure field is defined as the central position of TCs. A surface fitting technique is used so that the central position is not necessarily on a grid point of the mean sea level pressure fields provided. The initial TC position for each participating model is found by searching in a 500 km radius compared to best track data derived from a post-event reanalysis of all available data. Subsequent positions in the forecast data are searched in a 500 km radius around the (extrapolated) model track. The TC tracking ends when no more appropriate minimum pressure locations exist. Only TCs that have a maximum sustained wind of 34 knots or stronger during their lifetime are verified in this project.

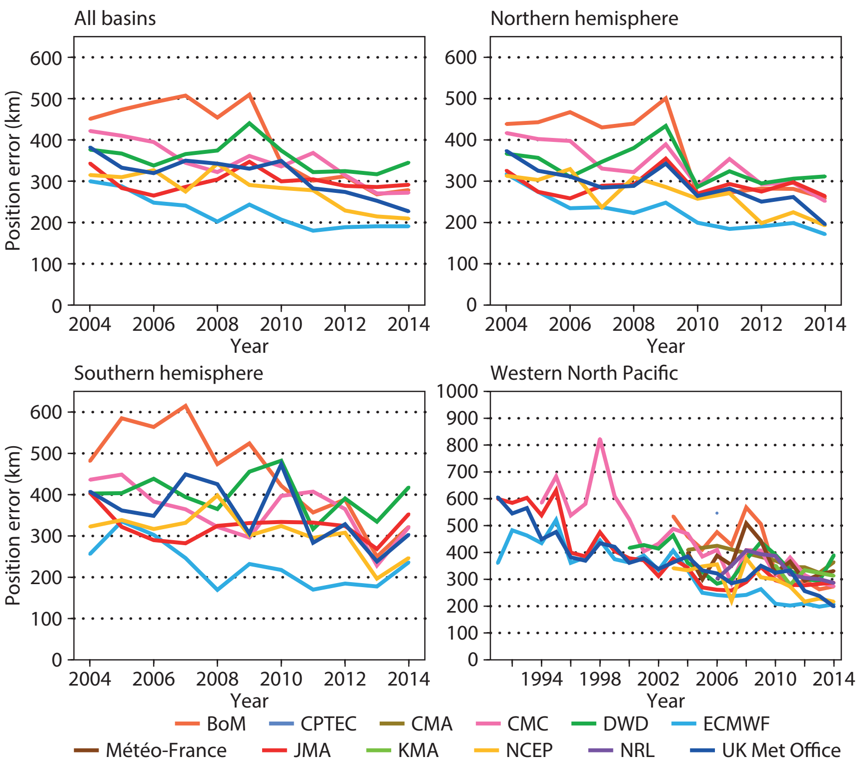

Timeseries of TC position forecast errors. The charts show the evolution of annual average TC position errors from 3-day forecasts for the models participating in the intercomparison project for all TC basins (top left), the northern hemisphere (top right), the southern hemisphere (bottom left) and the western North Pacific (bottom right).

Results and outlook

Results show that TC track forecasts have improved significantly both globally and in each TC basin. In the

NWP Centre	Year
Japan Meteorological Agency (JMA)	1991
ECMWF	1991
UK Met Office	1991
Canadian Meteorological Centre (CMC)	1994
Germany's National Meteorological Service (DWD)	2000
US National Centers for Environmental Prediction (NCEP)	2003
Australia's Bureau of Meteorology (BoM)	2003
China Meteorological Administration (CMA)	2004
Météo-France	2004
US Naval Research Laboratory (NRL)	2006
Brazil's Centre for Weather Forecasts and Climate Studies (CPTEC)	2006
Korea Meteorological Administration (KMA)	2011

Participating centres and the year when they joined the project. The first forecasts included in the comparison date back to 1991. Today 12 centres participate in the WGNE project.

WGNE

The Working Group on Numerical Experimentation (WGNE) has been jointly established by the World Climate Research Programme (WCRP) Joint Scientific Committee (JSC) and the World Meteorological Organization's Commission for Atmospheric Sciences (CAS), which is responsible for the World Weather Research Programme (WWRP) and the Global Atmosphere Watch (GAW) programme. Among other things, WGNE aims to foster the development of atmospheric circulation models for use in weather, climate, water and environmental prediction on all timescales and to diagnose and resolve shortcomings. One of its tasks is to promote coordinated numerical experimentation to validate model results.

western North Pacific, for example, by 2014 forecasts could be made 2.5 days further ahead than 21 years earlier, at the same level of accuracy. Time series of errors in different TC verification basins and of global and hemispheric errors show a steady improvement in TC track forecasting over the years for all models, although for some models there is substantially more interannual variability than for ECMWF's Integrated Forecasting System (IFS).

The performance of ECMWF forecasts is relatively good, with an average 200 km position error in all regions in recent years. But clearly further improvements are required to better support severe weather warnings in affected regions. ECMWF shares biases with other modelling centres, such as systematically underestimating the intensity of deep tropical cyclones. High-resolution forecasts, such as those produced at ECMWF and JMA, have the largest number of opposite-sign errors, where the predicted minimum pressure is lower than the analysed one.

ECMWF is well placed to address some of these issues. The accuracy of track and intensity forecasts should improve as the resolution of ECMWF's ensemble forecasts is increased; the data assimilation of observations is improved to reduce initial position errors; and advances are made in Earth system modelling by improving the model dynamics and physics, and by increasing complexity through coupling the ocean to the high-resolution TCo1279 forecast.

Further details on the method and the results presented here can be found in an article by M. Yamaguchi et al. accepted for publication in the *Bulletin of the American Meteorological Society* (doi:10.1175/BAMS-D-16-0133.1).

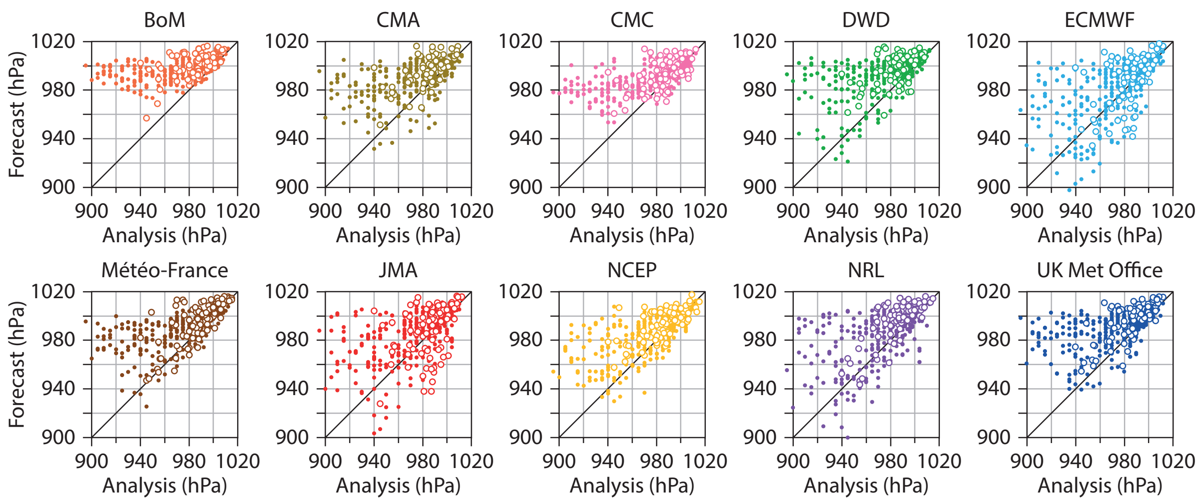

Intensity biases in TC forecasts. Scatter plots of analysed (best track) versus 3-day forecast minimum sea level pressure, indicating how the different models share similar TC intensity biases especially for TCs analysed to have very low minimum pressure. The verification period is 2012–2014.

doi:10.21957/bv418o

25 years of ensemble forecasting at ECMWF

ROBERTO BUIZZA, DAVID RICHARDSON

Twenty-five years ago ECMWF was one of the first forecasting centres which started to issue operational ensemble forecasts. The availability of such forecasts marked a paradigm shift in weather prediction: for the first time, forecasters and users were provided with reliable and accurate estimates of the range of possible future scenarios, and not just with a single realisation of the future. Today ensembles are used not only in forecast mode, to provide forecasts for the short- and medium-range, the monthly and the seasonal timescale, but also in analysis mode, to provide estimates of the initial state of the Earth system. These ensemble-based forecasts and analyses provide more complete information than single, deterministic forecasts, for example through indices of the risk of severe events; probabilities of the occurrence of weather events; the range of possible values at specific locations; alternative weather scenarios; and weekly-mean anomalies.

ECMWF ensembles have been developed, implemented and maintained thanks to the work of very many people at ECMWF and in its Member States, and of visitors who, over the years, have spent time working with us to understand their performance and to improve them further. This started well before 1992, with trials that helped us to identify the strategy to be followed, and it is continuing today, thanks to the interactions with scientists within global projects such as the World Meteorological Organization's TIGGE and S2S (sub-seasonal to seasonal prediction) projects. As we explain in this article, ensemble forecast performance is strongly linked to the quality of the model and the assimilation system used to generate the initial conditions; the assimilation of an increasing number of observations; the strategy that we have followed to simulate initial and model uncertainties; and the ensemble forecast configuration. Neither the performance of ensemble forecasts nor the range of ensemble-based products that we currently offer would be what they are today without those many contributions.

This article presents some examples of ensemble-based forecasts and explains their added value; it briefly reviews how we got where we are today, starting from ECMWF's first ensemble forecasts in 1992; it discusses the key characteristics of an ensemble system and the design of the ensembles operational today at ECMWF; it charts the evolution of ensemble forecast quality; and it describes the development of ensemble-based products to meet different user requirements. Finally, it looks to the future to highlight areas where further improvements can be made.

Example forecasts

Figures 1 to 5 give an impression of the breadth of information that ensemble-based products can provide.

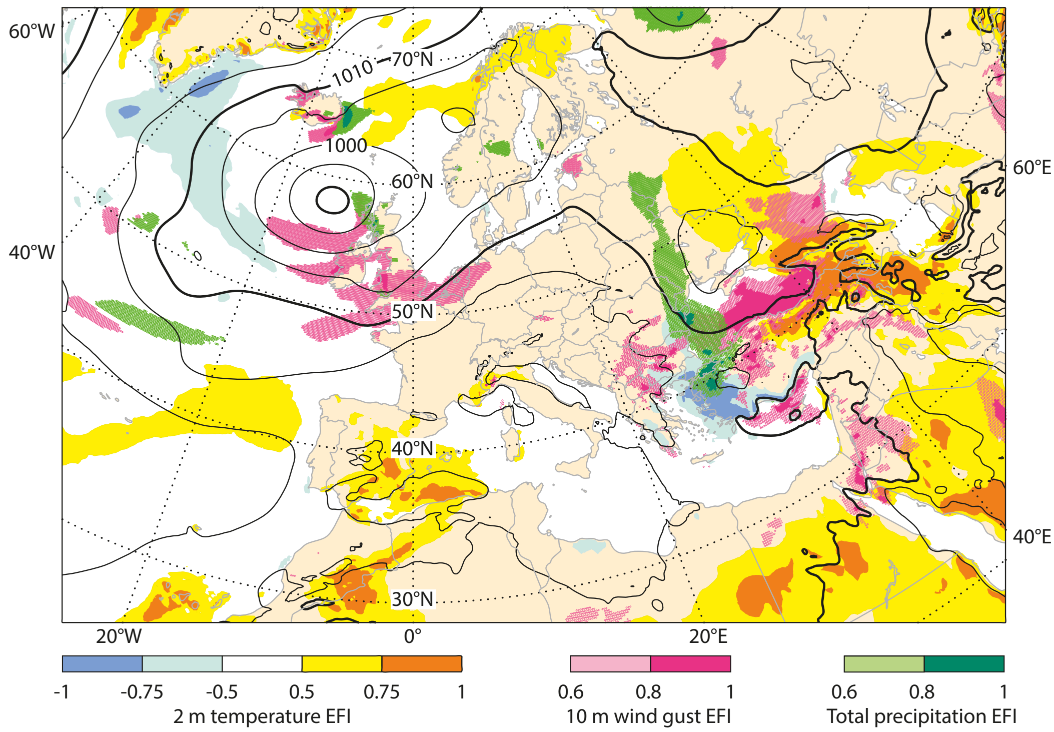

Figure 1 ENS-based Extreme Forecast Index (EFI) forecast issued on 27 July 2017 at 00 UTC and valid for 29 July. The map shows areas where the ensemble forecast distribution differs substantially from the model climatological distribution for three variables: 2-metre temperature, 10-metre wind gusts and precipitation. The black contours show the ensemble-mean forecast for mean sea level pressure. For example, the EFI map shows that southeastern Europe is predicted to experience extreme wind gusts and precipitation.

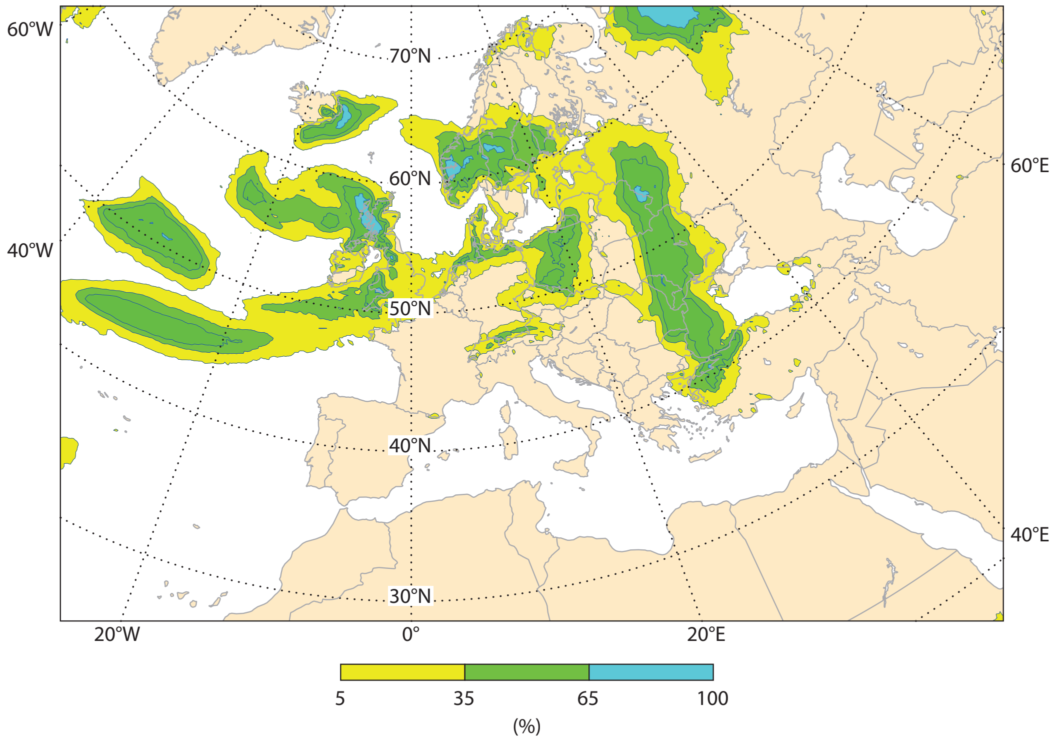

Figure 2 ENS-based probability of 24-hour precipitation in excess of 5 mm, issued on 27 July 2017 at 00 UTC and valid for 29 July.

The Extreme-Forecast-Index (EFI) forecast issued on 27 July 2017 (Figure 1) identified southeastern Europe as an area that could be affected by anomalous precipitation and wind anomalies. In terms of precipitation, products such as the probabilistic forecast of rainfall in excess of 5 mm/day confirm this (Figure 2). The forecast identifies northwest Turkey as a region that could be affected by rainfall events. A forecaster interested in more local weather, say for Istanbul, could then click on the EFI map and generate a 15-day ENS meteogram for this city (Figure 3). This product shows the whole range of possible values that surface variables such as cloud cover, precipitation, wind and temperature can reach, and it contrasts them with average, climatological values. Figure 3 shows that, indeed, Istanbul is expected to experience anomalous precipitation on 28 and 29 July. It also indicates that the two-metre maximum temperature for 28 July will be very low, close to the climatological minimum for this time of the year.

The meteogram for Istanbul also indicates that, after day four, the city will experience anomalous winds. To understand the synoptic-scale pattern associated with this event, we can look at the clusters for the 500 hPa geopotential height over Europe (Figure 4). They indicate that there is very little uncertainty over southeastern Europe (the three clusters have a very similar circulation in that region), with a low-pressure anomaly. Indeed, the small difference between the weather scenarios over Turkey for 31 July (Figure 4, right-hand column) is also reflected in the small spread in the wind forecast for Istanbul for that day (Figure 3, second diagram from the bottom).

Ensembles also provide very valuable information for longer time ranges. An example is given in Figure 5, which shows a series of monthly ensemble-mean forecasts of the weekly-average two-metre temperature anomaly over Europe, predicted for the week of 19 to 26 of June 2016. The plots show that the ensemble was able to predict the heat wave that affected Europe up to four weeks ahead.

1992: the start of a paradigm shift

From the early days of numerical weather prediction (NWP), it was clear that there are some cases when forecast errors remain small even for long forecast ranges, and others when even a 1-day forecast is wrong. This operational experience was supported by scientific studies that pointed out that, due to the chaotic nature of the atmosphere, even small initial errors can grow very rapidly and affect forecast quality at a very short range.

In the seventies and the eighties, we started investigating whether we could determine in advance, say when a forecast is issued, whether the future weather was easier or more difficult to predict than on average. In other words, we were looking for an objective method that could provide us with a level of forecast confidence. At that time different approaches were tested at the major NWP centres. It quickly became clear that the only feasible way to address this problem was to use ensembles. The main idea behind an ensemble approach is very simple: generate *N* forecasts, each of them designed to take into account possible uncertainties, and use the *N* forecasts to estimate the range of possible outcomes, and/or the most probable set of values, and/or the probability that temperature (or other variables) will be higher or lower than a certain value.

In the 1980s, different techniques were tried to develop reliable and accurate ensembles. These two adjectives,

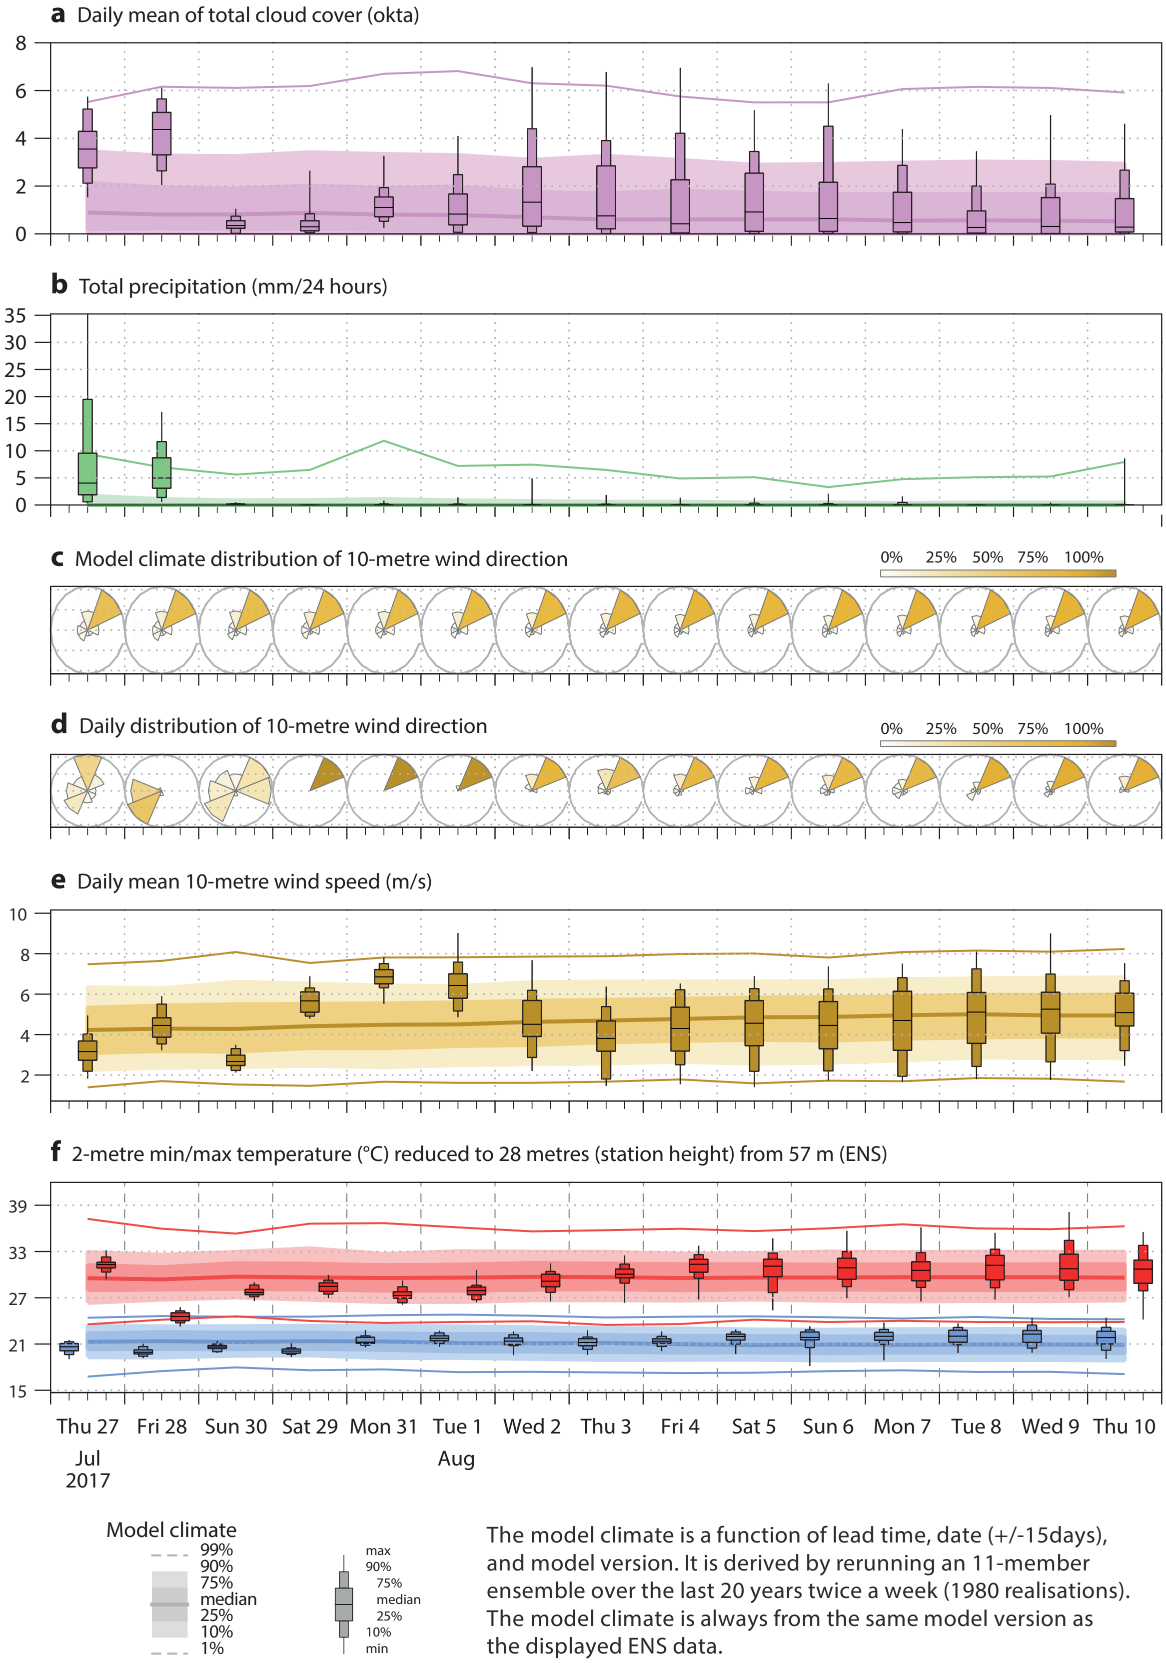

Figure 3 15-day ENS meteogram for Istanbul, issued on 27 July 2017 at 00 UTC, showing (a) daily mean total cloud cover, (b) total daily precipitation, (c) model climate 10-metre wind direction distribution, (d) daily 10-metre wind direction distribution, (e) daily mean 10-metre wind speed, and (f) 2-metre minimum and maximum temperature. For each variable, the plot shows the ensemble distribution (box-and-whiskers in most cases) and the model climate so that users can assess the range of possible future weather scenarios and how this compares with the model climate.

'reliable' and 'accurate', are key, since they define whether an ensemble is capable of providing valuable information. An ensemble is reliable when there is, on average over many cases (say a season), a good correspondence between a forecast probability and the probability of occurrence. More precisely, in a reliable ensemble, if an event is predicted with an 80% probability, it occurs 80% of the time when such a prediction is made. An ensemble is accurate when the average error of the ensemble mean is small. In a reliable ensemble, the average spread is equal to the average error of the ensemble mean. An ensemble is sharp when the spread of the ensemble members is small (so event probabilities tend towards 0 or 100%). A good ensemble forecast is as sharp as possible while still being reliable.

In the 1980s in the US initial tests used lagged ensembles, which mixed forecasts started at different times and on different days, e.g. the nine forecasts issued every six hours over the past two days. ECMWF tried to generate ensembles starting all at the same time, but with initial conditions perturbed in a random way. Results indicated that the US method delivered forecasts with a reasonable quality for the medium forecast range, beyond about a week, but not for the shorter forecast range, since the 'oldest' forecasts were too old to be accurate. The ECMWF methods did not deliver good results since the random perturbations did not lead to very different forecasts: the forecasts remained too similar to provide valuable information on possible future scenarios.

The beginning of the 1990s saw the development and testing of more promising methods both at ECMWF and at the US National Centers for Environmental Predictions (NCEP). 1992 saw the implementation of the first two operational ensemble systems in those two places. In 1995 the Meteorological Service of Canada (MSC) followed suit and others a few years later, both at the global scale and for specific regions.

These implementations generated a paradigm shift in operational NWP from a deterministic approach, based on a single forecast, to a probabilistic one, in which ensembles are used to estimate the probability density function of initial

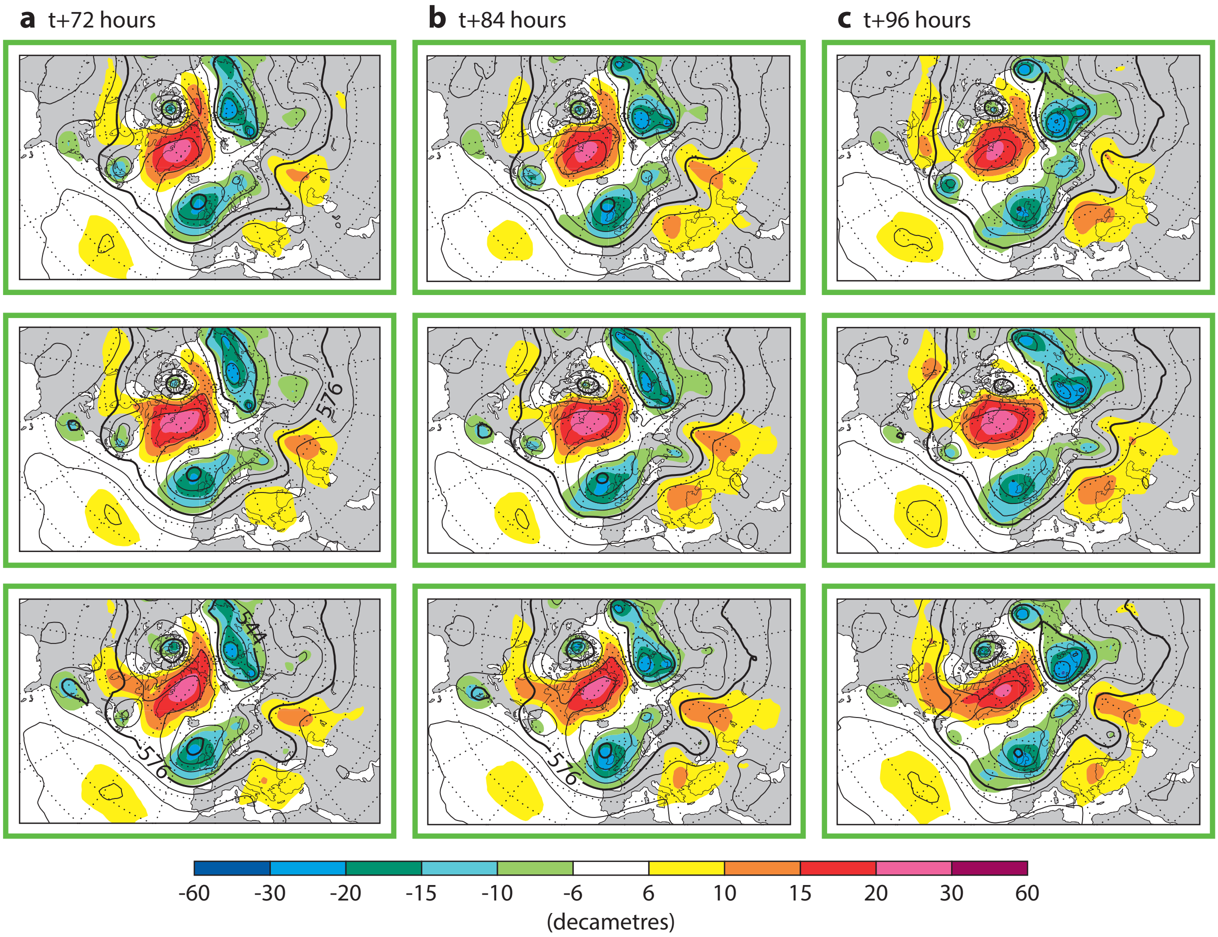

Figure 4 ENS-based 500 hPa geopotential height forecast clusters starting from 27 July 00 UTC for the European region, valid for (a) t+72 hours, (b) t+84 hours, and (c) t+96 hours. Cluster 1 (top row) includes 31 of the 51 ENS members, cluster 2 (middle row) 11 members and cluster 3 (bottom row) 9 members. Each panel shows the 500 hPa geopotential height (black contours) predicted by a representative member (RM) of the cluster and the anomaly (shading, computed with respect to the model climate). Each day, up to a maximum of five clusters are generated by an algorithm that defines in an objective way how many clusters are needed to represent the ensemble forecast distribution and identifies the RMs. The frame colour of each plot represents different climatological regimes. The green colour shown here stands for a negative North Atlantic Oscillation (NAO).

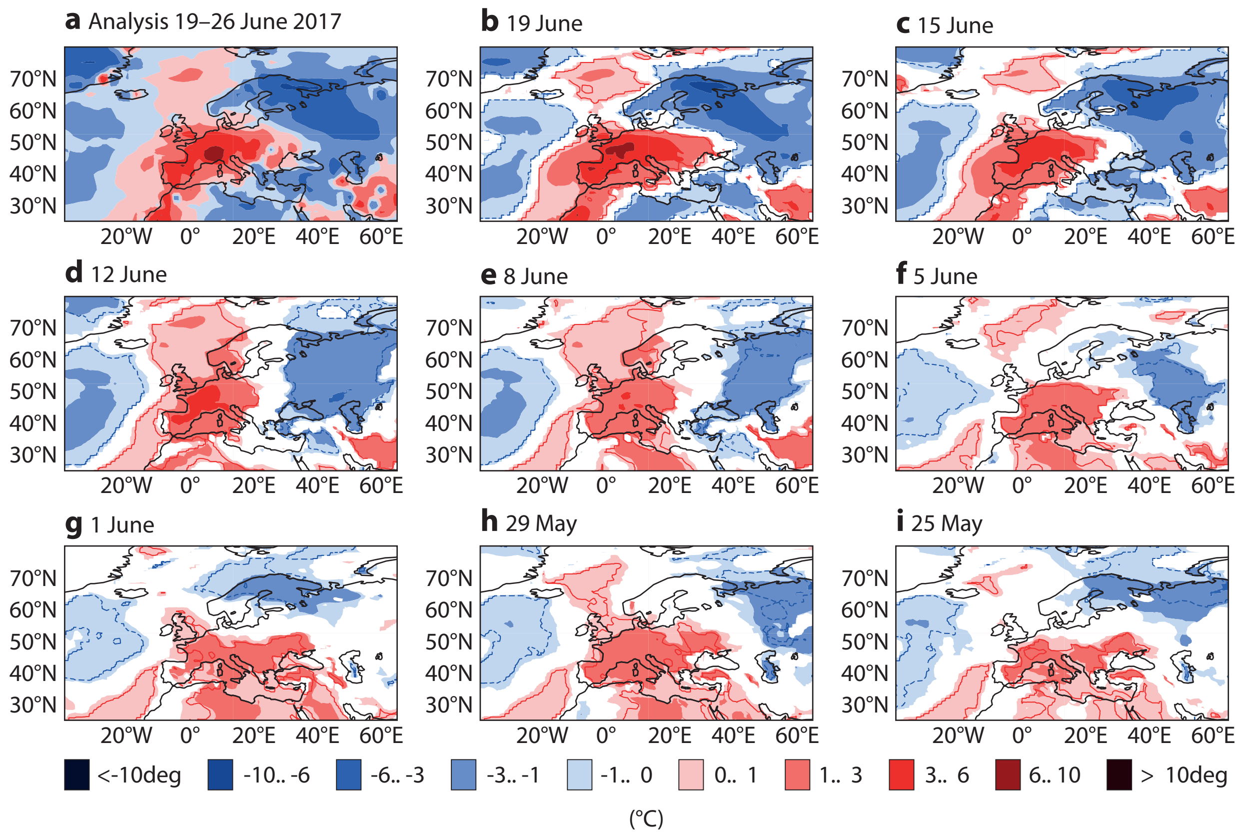

Figure 5 Analysis and ENS-based monthly forecasts for weekly-average anomalies of 2-metre temperature between 19 and 26 June 2016, showing (a) the observed anomaly, defined by the average of ECMWF analyses over that period, and showing the forecasts starting on (b) 19 June, (c) 15 June, (d) 12 June, (e) 8 June, (f) 5 June, (g) 1 June, (h) 29 May and (i) 25 May. Each forecast panel shows the ensemble-mean 2-metre anomaly.

and forecast states. Products such as the ones shown in Figures 1–5 would not exist if it was not for the development and operational implementation of these ensembles.

Added value of ensemble forecasts

Today it is widely accepted that forecasts have to include uncertainty estimations, confidence indicators that allow forecasters to estimate how 'predictable' the future is. Short- and medium-range forecasts, monthly and seasonal forecasts, and even decadal forecasts and climate projections are today based on ensembles, so that the uncertainty associated with the forecast can be estimated. Furthermore, ensembles are widely used to provide an estimate of the initial state uncertainty, to estimate the analysis error more accurately.

Ensemble-based, probabilistic forecasts are more valuable than single forecasts. This is mainly due to the fact that they provide probabilities for different events to occur. In other words, ensembles provide users with more complete information about future weather scenarios. One way to measure such a difference is to apply simple cost–loss models using a measure called the Potential Economic Value (PEV) of a forecasting system (*Richardson* 2000). Another reason why ensemble-based, probabilistic forecasts are more valuable than single forecasts is that they provide more consistent (i.e. less changeable) successive forecasts. For example, consecutive ensemble-mean forecasts issued 24-hour apart and valid at the same time are generally found to jump less than corresponding single forecasts such as the high-resolution forecast or the ensemble control forecast (the ensemble member that starts from the 'most likely' initial state, defined by the unperturbed analysis). By using the whole ensemble, the unpredictable features are averaged out, and the predictable features (the signals) can be extracted.

Design of medium-range global ensembles

Ensembles are designed to simulate the sources of forecast errors linked to initial condition and model uncertainties. Model uncertainties arise because the models that we use to generate weather forecasts are imperfect, simulate only certain physical processes on a finite mesh, and do not resolve all the scales and phenomena that occur in the real world. Initial condition uncertainties arise because observations are affected by observation errors and do not cover the whole globe with a uniform density and frequency. Furthermore, the process of estimating the initial state of the system, from which a forecast is computed, is based on some statistical assumptions and approximations.

In the first version of the ECMWF global ensemble (*Molteni et al.*, 1996), initial uncertainties were simulated using singular vectors (SVs), which are the perturbations with the fastest growth over a finite time interval (*Buizza & Palmer*, 1995). SVs provided a very good basis to define the initial perturbations of the ECMWF ensemble: compared to the random initial perturbations tried in the 1980s, they led to a very good growth rate in the spread of the ensemble, similar to the forecast error growth rate. SVs remained the only type of initial perturbations used in the ECMWF ensemble until 2008, when the Ensemble of Data Assimilations (EDA) started being used, together with singular vectors (*Buizza et al.*, 2008). EDA-based perturbations were added to improve the simulation of initial errors linked to the characteristics of the observing system (observation errors, coverage, scalability...). Today SVs remain an essential component of the ECMWF ensemble, and they keep providing dynamically relevant information about initial uncertainties that could have a strong, negative impact on forecast errors.

There are different ways to simulate initial and model uncertainties. For example, in the first version of NCEP's global ensemble, bred vectors (BVs) were used to simulate initial uncertainties instead of SVs. The BV cycle aims to emulate the data assimilation cycle. It is based on the notion that analyses generated by data assimilation will accumulate errors that have a tendency to grow by virtue of perturbation dynamics (*Toth & Kalnay*, 1997). On the one hand, errors that have a tendency to stay constant or to decay will be reduced when detected by an assimilation scheme in the early part of the assimilation window. What remains of them will decay by the end of the assimilation window due to the dynamics of such perturbations. On the other hand, even if errors that have a tendency to grow are reduced by the assimilation system, what remains of them will have amplified by the end of the assimilation window.

In 1995 the ECMWF and the NCEP ensembles were followed by the Canadian ensemble. The Canadians adopted a Monte Carlo approach, designed to simulate both initial uncertainties due to observation errors and data assimilation assumptions, and model uncertainties (*Houtekamer et al.*, 1996). The Canadian ensemble was the first to include a simulation of model uncertainties. It tried to include as many sources of error as possible.

Following the Canadian example, a stochastic scheme designed to simulate model uncertainties was introduced in the ECMWF ensemble in 1999 (*Buizza et al.*, 1999). Since then, many other operational ensembles have also included such schemes to simulate model uncertainties. *Buizza* (2014) provides a review of the main characteristics of the operational global ensembles available in the TIGGE database.

At present, as detailed by *Palmer et al.* (2009), four main approaches are followed in ensemble prediction to represent model uncertainties:

- A multi-model approach, where different models are used to construct ensembles; models can differ entirely or only in some components (e.g. in the convection scheme);
- A perturbed parameter approach, where all ensemble integrations are performed with the same model but with different parameters defining the settings of the model components; one example is the Canadian ensemble (*Houtekamer et al.*, 1996);
- A perturbed-tendency approach, where stochastic schemes designed to simulate the random model error component are used to simulate the fact that tendencies are known only approximately: one example is the ECMWF Stochastically Perturbed Parametrization Tendency scheme (SPPT) (*Buizza et al.*, 1999);
- A stochastic back-scatter approach, where a Stochastic Kinetic Energy Backscatter scheme (SKEB) is used to

	Forecast length	Resolution		Number of members	Re-forecasts (number of members; number of years)
		Horizontal	Vertical (number of layers)		
ORAS5 (Ocean reanalysis System-5)	--	100 km	75 layers	5 (once a day)	–
EDA (Ensemble of Data Assimilations)	--	18 km	137 (to 0.01 hPa)	25 (00, 12 UTC)	–
ENS for boundary-condition generation	6.5 days	18 km	91 (to 0.01 hPa)	51 (06, 18 UTC)	–
ENS for the medium range	15 days	18 km	91 (to 0.01 hPa)	51 (00, 12 UTC)	Yes (22/week; 20 years)
ENS for the monthly range	46 days	36 km	91 (to 0.01 hPa)	51 (00 UTC on Mon and Thu)	Yes (22/week; 20 years)
SEAS4 (seasonal-range ensemble)	7 and 13 months	80 km	91 (to 0.01 hPa)	51 (1st of each month)	Yes (15/month; 30 years)

Table 1 Key characteristics of the ECMWF ensembles. All ensembles simulate initial and model uncertainties.

simulate processes that the model cannot resolve, e.g. the upscale energy transfer from the scales below the model resolution to the resolved scales; an example is the SKEB scheme currently used in the ECMWF ensemble, which is due to be switched off in 2018 since, in its current formulation, it does not deliver any significant benefits.

Ensemble configurations

Two key aspects that define the characteristics of an ensemble are the methodology used to simulate initial uncertainties and the approach adopted to simulate model approximations. A third key characteristic of an ensemble is its resolution, both horizontal and vertical. A fourth aspect is the forecast length, and the fifth key aspect of an ensemble configuration is the number of ensemble members.

Theoretical work done in the 1970s and 1980s suggested that one needs at least about 10 members for a good ensemble-mean forecast, i.e. one which has enough members to average out the unpredictable scales or features. But are 10 members enough to get a good probability distribution forecast, and not just a good ensemble-mean? Results obtained in the 1990s and 2000s based on the comparison of ensemble sizes of up to a few hundred indicated that reliability and accuracy are very sensitive to ensemble size. On synoptic scales (scales of a few hundred kilometers), increasing the ensemble size from say 10 to about 50 was found to have a clear and detectable impact on ensemble forecast performance. Further increases beyond 50 have a smaller effect, on average, but can still have a detectable impact if one wants to predict rare events. When ensembles are used to estimate the analysis uncertainties, increasing the ensemble size to a few hundred members was found to bring clear benefits. Today most operational ensemble forecasts have between 20 and 50 members, while ensembles of analyses have up to about 300 members (although this number can be substantially lower depending on the computational cost of the analysis method).

Resolution, forecast length and the number of ensemble members are key cost drivers of ensemble production. Given that we need to generate forecasts in a reasonable amount of time if we want them to be valuable (say about 1 hour), and that we have a finite amount of computing resources, compromises have to be made when an ensemble configuration is defined. Ideally, we would like to use as many members as possible and the highest resolution possible to be able to simulate also the finest scales so that we can provide detailed forecasts, including of severe weather. We would also like to extend the forecast length as much as possible to provide a bigger set of users with ensemble-based, probabilistic forecasts.

The compromise struck in ECMWF's Integrated Forecasting System (IFS) is to use a number of different configurations for the ocean, the Ensemble of Data Assimilations, boundary condition generation, medium-range and monthly ensemble forecasts (ENS), and seasonal forecasts (see Table 1 for details).

Figure 6 illustrates how the three ECMWF ensembles and the high-resolution analysis and forecast are linked together:

- The 25-member EDA and the high-resolution analysis (both using 4D-Var) are used to generate the initial conditions of the two coupled forecast ensembles, ENS and SEAS4;
- The 5-member ORAS5 (Ocean Reanalysis and Analysis, version S5) is used to initialise the dynamic ocean and sea-ice components of the two coupled ensembles, ENS and SEAS4;

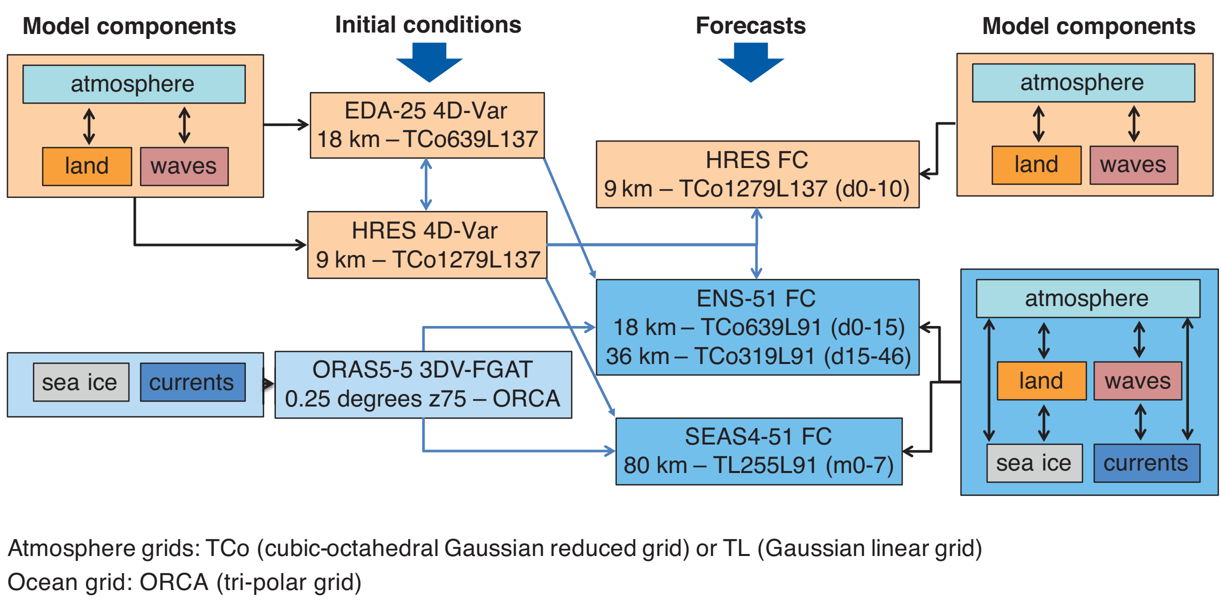

Figure 6 Diagram showing the links between the three ECMWF ensembles (EDA, ENS and SEAS4) and Earth system components used to generate the initial conditions and the forecasts.

- The high-resolution analysis is used to generate the initial conditions of the single, high-resolution forecast (HRES).

Figure 6 also illustrates the Earth-system components included in the initial condition and forecast models:

- The EDA, the high-resolution analysis and the HRES use the ECMWF land and atmosphere model and the ECMWF wave model (ECWAM);
- The ocean analysis ORAS5 uses the NEMO (Nucleus of European Modelling of the Ocean) ocean model (see *https://www.nemo-ocean.eu/*) and the LIM2 (Louvain-la-Neuve) sea-ice model (see *http://www.elic.ucl.ac.be/repomodx/lim/*);
- The two coupled ensembles ENS and SEAS4 use all model components: IFS+ECWAM+NEMO+LIM2.

The high-resolution forecast is due to be coupled with the ocean and sea-ice models (NEMO+LIM2) from the beginning of 2018, after the implementation of IFS Cycle 45r1.

Ensembles for sub-seasonal and seasonal timescales

Since the beginning of the 2000s, ensembles have also been used to generate monthly and seasonal forecasts. These extended-range ensembles are global and have a coarser resolution than the medium-range ensembles to limit production costs (see Table 1). Since extracting predictable signals for the extended range is very difficult, these ensembles have been complemented by re-forecast suites, which are smaller-size ensembles with the same configuration as the operational ensembles (apart from their size) generated for the last few decades. After the medium-range and the monthly ensembles were joined in 2008, with the implementation of the VAREPS approach, we have been able to exploit the re-forecasts to design new and better products for the medium range too.

Today the re-forecasts are used to estimate the ensemble characteristics (reliability and accuracy, and model biases), and they help to generate forecast products across the whole forecast range, for example products such as the EFI (shown in Figure 1), the 15-day meteograms (Figure 3) and some of the extended-range probabilistic forecasts (e.g. the ones shown in Figure 5) by providing a model climatology.

ECMWF is upgrading the seasonal ensemble to SEAS5, which will have the same resolution as the ENS monthly extension. SEAS5 is due to become operational in November 2017.

Ensembles of analyses and reanalyses

Since its inception in 1995, the Canadian ensemble has included an ensemble of analyses, generated using an ensemble Kalman filter (EnKF). The initial conditions of each of the ensemble members are defined by one of the members of the EnKF. The EnKF has been providing the Meteorological Service of Canada with information about uncertainties in the analysis.

At ECMWF and Météo-France, we started producing an Ensemble of Data Assimilations in 2008. We run an ensemble of *N* separate data assimilation procedures, each using perturbed observations and a model uncertainty scheme. Observations are perturbed to simulate the fact that observations are not perfect due to observation errors, and to take into account observation representativeness errors. Model uncertainties are simulated to take into account the fact that the models used to define the analysis are not perfect.

Since 2008, the ECMWF EDA has been used in combination with SVs to define the initial conditions of the medium-range/monthly ensemble (*Buizza et al.*, 2008). The addition of EDA-based perturbations has had a major impact on ensemble reliability and accuracy in the short forecast range over the extratropics, and for the whole forecast range over the tropics.

Since 2002, ECMWF has been producing a five-member ensemble of ocean analyses and reanalyses to initialise the ocean component of coupled ensembles. The medium-range/monthly ENS started using the ocean ensemble in 2008 from day 10, when it was merged with the monthly ensemble (see Box A). Since 2013, ENS has been using the ocean ensemble ORAS5 from initial time.

Changes in configuration over 25 years

Since its inception in 1992, the medium-range ensemble has changed configuration several times. A chronology of the main configuration changes is provided in Box A. It is interesting to compare the ENS configuration implemented in operations in December 1992 with that of 2016.

In terms of the key cost drivers of ensemble production, the main changes are:

- the horizontal resolution has increased by a factor of 20, from about 320 km to about 16 km;
- the vertical resolution has increased by a factor of almost 5, from 19 to 91 vertical levels;
- the forecast length has been extended from 10 to 46 days;
- the number of ensemble members has increased from 33 to 51;
- the frequency of ENS forecast production has increased: in 1992 we produced 99 ensemble forecasts each week (3x33), while today we produce 1428 ensemble forecasts each week up to 6.5 days (4x51x7); of these 1428 forecasts, 714 are extended up to 15 days (2x51x7); of these 714 forecasts, 102 are extended up to 46 days (2x51);
- today we also produce ensemble re-forecasts: every week, we generate 440 ensemble forecasts up to 46 days (2x11x20).

Evolution of ensemble forecast quality

Thanks to model upgrades, improvements in the data assimilation system, the use of more observations, and the ENS configuration changes discussed above, the ENS performance has increased substantially during the past 25 years.

ENS configuration changes

A

This list includes the main changes made to the medium-range/monthly ensemble (ENS) since its first day of real-time production and dissemination on 19 December 1992:

- **Dec 1992:** the ensemble starts running three days a week (Fri-Sat-Sun, at 00 UTC); initial uncertainties are simulated using only the initial-time SVs with a T21L19 resolution, computed with a 36-hour optimisation time interval, over the whole globe; only the initial-time SVs are used, and the initial perturbations are symmetric; the forecast resolution is T63L19 (~320 km); forecasts are run up to 10 days; ENS includes 33 members; there is no simulation of model uncertainties, no coupling to an ocean/sea-ice model, and no re-forecast suite;
- **Feb 1993:** to address the fact that SVs were concentrating mainly in the southern hemisphere (SH), the Local Projector Operator (LPO) was introduced, to allow SVs to be located in the northern hemisphere (NH); this had a major impact on ensemble reliability over the NH;
- **Aug 1994:** to improve the ensemble spread, the optimisation time interval (OTI) of the SV computation was increased from 36 to 48 hours; this improved the perturbation growth also beyond the OTI; from 1 May 1994, ENS forecasts were generated every day, once a day, at 00 UTC;
- **Mar 1995:** the horizontal resolution of the SVs was increased to T42; this improved perturbation growth and thus ensemble reliability;
- **Mar 1996:** a second set of SVs was introduced, targeted to grow over the SH; this had a major impact on ensemble reliability over the SH;
- **Dec 1996:** the resolution of ENS was increased to TL159L31 (~120 km), and the number of members was increased from 33 to 51;
- **Mar 1998:** a second set of SVs, called evolved SVs, that grow during the two days before the initial date, were added to the initial-time SVs; the evolved SVs simulated the effect of errors growing during the data-assimilation period; their addition improved ensemble reliability (spread) especially in the short range;
- **Oct 1998:** the stochastic model error scheme (SPPT) was introduced to simulate the effect of model uncertainties linked to physical parameterization; this had a large impact on ensemble reliability (spread) over the whole forecast range, and especially over the tropical region;
- **Oct 1999:** vertical resolution was increased from 31 to 40 levels;
- **Nov 2000:** ENS horizontal resolution was increased from TL159 to TL255 (~80 km);
- **Jan 2002:** SVs targeted to grow over the tropical region, in areas where tropical depressions were identified, were added; they led to improved spread over the tropical region, and especially in cases of tropical storms;
- **Sep 2004:** the sampling strategy applied during the generation of ENS initial perturbations was changed to Gaussian sampling;
- **Jun 2005:** the Gaussian sampling method was revised;
- **Feb 2006:** ENS horizontal resolution was increased from TL255L40 to TL399L62 (~60 km);
- **Sep 2006:** ENS was extended to 15 days, with the use of variable resolution (VAREPS), whereby the forecast resolution was truncated at day 10 from TL399 to TL255;
- **Mar 2008:** the medium-range ensemble (ENS) and the monthly ensemble were merged, using the VAREPS technique; ENS was run to 32 days once a week (Mon at 00 UTC); ENS was then coupled to the dynamical ocean model HOPE from forecast day 10; the ENS re-forecast suite with a 5-member ensemble run once a week for the past 18 years was introduced;
- **Sep 2009:** the stochastic model error scheme was revised;
- **Jan 2010:** the horizontal resolution was increased from TL319 to TL639 (~35 km) in the first 10 days, and from TL255 to TL319 (~70 km) from day 10 to day 32;
- **Jun 2010:** a new set of initial perturbations, generated using the 10-member Ensemble of Data Assimilations (EDA), was introduced in ENS; the EDA-based perturbations improved the simulation of the perturbations linked to the data-assimilation cycle, and replaced the evolved SVs; this led to improvements in ensemble reliability (spread), especially in the short forecast range and over the tropical region;
- **Nov 2010:** a second scheme, the stochastic back-scatter (SB) scheme, was introduced to simulate model error;
- **Nov 2011:** a new ocean model was introduced: NEMO with a 1-degree resolution (~100 km) replaced HOPE; the ENS extension to 32 days started being run twice a week (Mon and Thu at 00 UTC);
- **Jun 2012:** the EDA-based perturbations were revised to include perturbations in the surface fields, and the re-forecast suite was enlarged to cover the past 20 years;
- **Nov 2013:** the vertical resolution was increased from 62 to 91 vertical levels, and the coupling to the ocean model was moved from day 10 to day 0; this led to major improvements in the prediction of phenomena over the tropics, such as the Madden–Julian Oscillation (MJO);
- **May 2015:** forecast length was extended from 32 to 46 days, and the re-forecast suite was enlarged to include two 11-member ensembles run every week (on Mon and Thu), covering the past 20 years;
- **Mar 2016:** the horizontal resolution was increased to TCo639L91 (cubic-octahedral grid; ~18 km) up to day 15, and to TCo319L91 (~36 km) from day 15 to day 46;
- **Nov 2016:** the ocean model resolution was increased from 1 degree to 0.25 degrees (~25 km), and the number of vertical layers from 42 to 75; the interactive sea-ice model LIM2 was introduced.

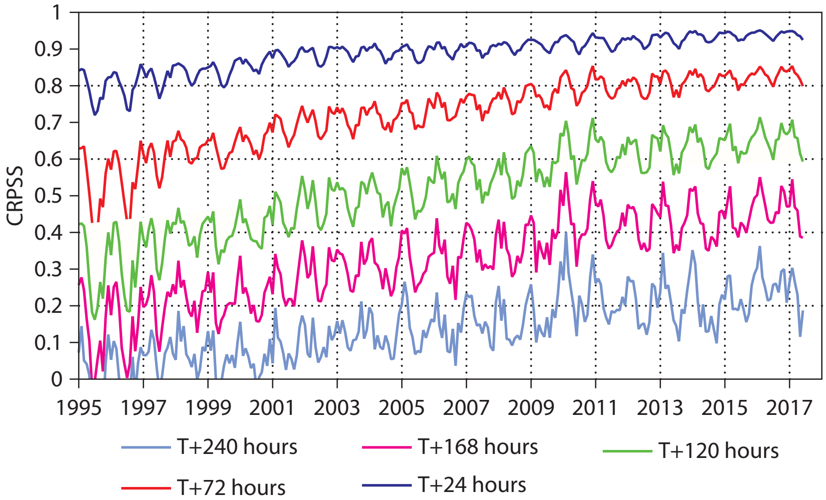

Figure 7 Time evolution, from 1 January 1995 to today, of the Continuous Ranked Probability Skill Score (CRPSS) of ENS forecasts of 500 hPa geopotential height over the northern hemisphere, for lead times of 24 hours, 72 hours, 120 hours, 168 hours and 240 hours. Forecasts are verified against operational analyses. The more or less regular pattern of peaks and troughs in each line stems from differences in predictability related to the seasons: winter weather tends to be more predictable than summer weather.

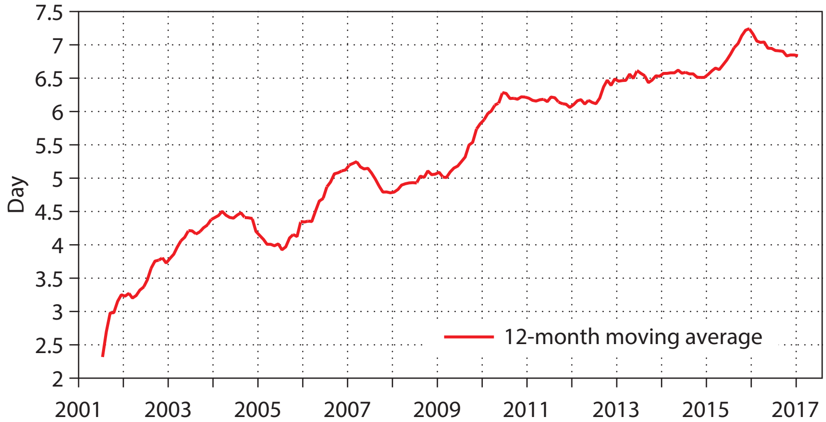

Figure 8 Time evolution of the forecast lead time when the CRPSS for the prediction of 24-hour accumulated precipitation drops below 0.1. Forecasts are verified against observations at SYNOP network weather stations in the extratropics.

Figure 7 shows the time evolution of the skill of ENS forecasts for 500 hPa geopotential height over the northern hemisphere from 1 January 1995 to today. Skill is measured by the Continuous Ranked Probability Skill Score (CRPSS), which compares the Continuous Ranked Probability Score (CRPS) of ensemble forecasts with that of a reference forecast, such as climatology. CRPS measures how close ensemble distributions are to observed values. The CRPS for a deterministic forecast is equal to the mean absolute error. CRPSS has a value of 1 for a perfect forecast, and is zero for a forecast that has the same skill as a statistical forecast based on climatology. Figure 7 shows that for 500 hPa geopotential height, a variable that describes the large scales in the free atmosphere, ENS forecasts have improved by about 1.5 days per decade. For example, today's 5-day forecasts (green line) are as skilful as 3-day forecasts (red line) were in 2001. This represents a predictability gain of about 2 days over a 16-year period.

If we look closer to the surface, results are even more striking. For precipitation, for example, Figure 8 shows that between 2002 and 2017 the lead time when the CRPSS dropped below 0.1 increased from about forecast day 3 to about forecast day 7, equating to a predictability gain of about 4 days over a 15-year period. Similar improvements are found for other variables and other regions (e.g. Europe, not shown).

Looking at longer forecast lead times, it is worth remembering that in 2006 the medium-range ensemble was extended to 15 days, and that in 2008 ENS was joined to the monthly ensemble and extended to 32 days. In 2015 it was further extended to 46 days. These extensions were justified by the fact that forecasts for these extended lead times had been improving as well. Clearly, for these forecast ranges only spatially large scales and time-averaged fields can be predicted with a certain level of skill. Results documented in *Buizza & Leutbecher* (2015) and *Vitart et al.* (2014) show that for these large-scale, low-frequency phenomena the forecast skill horizon has been extended to several weeks. The reader is also referred to *Vitart et al.* (2014) for a comprehensive overview of how the skill of ECMWF monthly forecasts evolved during the preceding 15 years.

Product development

Key to the successful use of ensemble forecasts is the ability to extract and communicate the information that is relevant to each user's decision-making process.

Alongside the development of the ENS perturbation methodologies, there has been substantial work and progress in the development of ensemble-based forecast products to address a range of different user requirements and enable forecasters to extract the appropriate information from the ENS.

When the ensemble was first introduced, the number of ENS-based products was limited. 'Stamp' maps showed each ENS member at forecast day 7, allowing the user to quickly assess by eye the range of possible weather states. This was complemented by cluster products that objectively grouped the set of ENS members into a small number of different scenarios that showed the different forecast evolution for 5 to 7 days ahead over Europe. For a small set of pre-defined locations, 'plume diagrams' showed the evolution of a small number of surface parameters through the forecast range. These products were issued to users by fax.

Nowadays users have access to a wide range of ENS data and products that process and present the ensemble information in different ways according to the needs of the user. The focus at ECMWF is to provide generic products that will be useful to assist operational weather forecasters. Many users complement the ECMWF products by doing their own post-processing to generate specific products tailored to their individual needs.

Extreme Forecast Index

The Extreme Forecast Index (EFI) was specifically designed to alert forecasters to occasions of potentially extreme weather (*Lalaurette*, 2003; *Zsoter*, 2006). The EFI compares the current ensemble forecast to the model climate distribution (generated by running a large set of re-forecasts over the last 20 years). It highlights areas where the current ENS forecasts are showing an enhanced likelihood of unusual weather. A large EFI indicates that the weather is likely to be extreme in the context of what can occur locally.

The EFI is one of the most popular ensemble products with forecasters. It can be especially useful in forecasting around the world, when forecasters may not have detailed knowledge of the regional climate. Since the EFI focuses on anomalies relative to the local climate, it is especially relevant for impact-based forecasting, where local extremes (or return periods) are more relevant than fixed event thresholds.

Storm tracks

Also relevant for severe weather forecasting are specific sets of products for extra-tropical and tropical cyclones. In both cases the cyclones are tracked in each ENS member and a range of products show the evolution of certain features along the forecast track, such as central pressure and maximum wind associated with the system. See the separate article on the hurricanes Harvey and Irma in this Newsletter for examples.

Both sets of products are designed to show information about the tracks and intensities of storms in the forecast and to help the forecasters quickly answer questions such as where and when severe storms will occur; how intense they will be; and where there may be a risk of a severe tropical cyclone in the coming days or weeks.

Many ENS products are available on the ECMWF website and many are now interactive, allowing the user to for example click on a location of high EFI to examine the details of the full ENS distribution at that location. ecCharts is an interactive web application that enables users to explore the ECMWF forecasts in even more detail. It allows them to zoom in on any area of interest; to select and overlay different forecast parameters; to compare and combine HRES and ENS forecasts; to compute probabilities for specific events of their own choosing (for example by selecting a precipitation threshold and time interval); and even to define combined events (such as the probability of both heavy precipitation and extreme wind).

There are several other products not mentioned here which highlight different aspects of the ensemble distribution, for use by forecasters in different situations. Each is designed to extract the most relevant information from the ensemble and to present it to the forecaster as clearly as possible, so the forecaster can focus on their job without having to spend time themselves trying to process the huge amount of information in the ensemble.

A look to the future

Looking to the future, three trends can be detected in the way ensembles are being upgraded:

i. A move towards an Earth-system approach to modelling and assimilation;

ii. A move towards a seamless approach in the design of the analysis, medium-range, sub-seasonal and seasonal ensembles;

iii. A move towards higher resolution.

The first trend is justified by results obtained in the past two decades that have shown that by adding relevant processes we can further improve the quality of the existing forecasts, and we can further extend the forecast skill horizon at which dynamical forecasts lose their value.

The second trend is partly motivated by scientific developments and partly by technical requirements. From a scientific point of view, there is evidence that processes that were thought to be relevant for the extended range are also relevant for the short range. An example is the introduction of a dynamic ocean in ECMWF ensembles. We started using a coupled ocean–land–atmosphere model for the seasonal and monthly timescales. We also introduced it in the medium-range ensemble once we realised that it could help to improve its reliability and accuracy. From a technical point of view, having an integrated approach whereby the same model is used in analysis and prediction mode, from day 0 to year 1, simplifies maintenance and the implementation of upgrades. Furthermore, it facilitates the diagnostics and evaluation of a model version, since tests carried out over different timescales can help to identify undesirable behaviour that could lead to forecast errors.

The third trend stems from the need to better resolve the smaller scales and their interaction with slightly less small scales, and so on. All scales from the microphysics within individual clouds to large-scale weather systems covering thousands of kilometres are relevant in weather prediction, and errors propagate from the smallest to the larger scales. If we consider the current operational ensemble, we should not forget that even if it uses a grid spacing of 18 km in the first 15 days, it can actually resolve in a realistic way only scales that are about 5 to 6 times the grid spacing. This is because the scales closest to the model grid spacing are not simulated in an accurate way. Thus, today the ECMWF global ensembles (EDA and ENS) have an effective resolution of about 100 km. Even the highest-resolution limited-area ensembles, such as the ones in operation at Météo-France and the UK Met Office, have a resolution of about 2 km, which corresponds to an effective resolution of about 10 km.

ECMWF's ten-year Strategy adopted in 2016 sets ambitious goals in line with these requirements. These include the introduction of a 5 km global ensemble by 2025. Even this will, however, not be the end of the road. If we want to be able to predict weather events such as intense wind storms or heavy precipitation at the scales at which they occur, it will be essential for model resolution to be increased to a few hundred metres for limited-area models and, in the long term, possibly to even finer resolutions than 5 km globally.

In conclusion … ensembles are the way forward!

The past 25 years have seen major advances in ensemble prediction, both in the way ensembles are generated and in ensemble products. Forecasts have become more accurate and reliable thanks to improvements in the initial conditions (i.e. in the use of observations and in the data assimilation system used to generate them); in the quality of forecast models; and in ensemble configurations. The introduction of additional relevant Earth system processes, such as the coupling to dynamic ocean and sea-ice models, has also led to improvements, and it has helped us to 'tame' the butterfly effect (*Buizza et al.*, 2015). The use of re-forecasts has made it possible to extract more meaningful signals from the raw forecast data.

We are confident that the future will see the use of ensembles also in areas where they are not yet used. Ensemble reliability and accuracy will continue to improve as a result of further advances in models, data assimilation methods, and the schemes used to simulate the initial and model uncertainties. Resolution will be increased to better simulate small-scale processes that are not currently resolved and to capture their important interactions with larger-scale processes. Ensembles of analyses and forecasts will be linked closer together to improve their performance. Physical processes that are not yet included in the models but that are relevant for weather prediction will be included, to make the forecasts more and more realistic.

The time is right: ensembles are the way forward!

FURTHER READING

Barkmeijer, **J.**, **R. Buizza**, **E. Källén**, **F. Molteni**, **R. Mureau**, **T. Palmer**, **S. Tibaldi** & **J. Tribbia**, 2013: 20 years of ensemble prediction at ECMWF. *ECMWF Newsletter* **No. 134**, 16–32.

Buizza, **R.** & **T.N. Palmer**, 1995: The singular-vector structure of the atmospheric general circulation. *J. Atmos. Sci.*, **52**, 9, 1434–1456.

Buizza, **R.**, **M. Miller** & **T.N. Palmer**, 1999: Stochastic representation of model uncertainties in the ECMWF Ensemble Prediction System. *Q.J.R. Meteorol. Soc.*, **125**, 2887–2908.

Buizza, **R.**, **M. Leutbecher** & **L. Isaksen**, 2008: Potential use of an ensemble of analyses in the ECMWF Ensemble Prediction System. *Q.J.R. Meteorol. Soc.*, **134**, 2051–2066.

Buizza, **R.**, 2014: The TIGGE medium-range, global ensembles. *ECMWF Technical Memorandum* **No. 739**.

Buizza, **R.**, & **M. Leutbecher**, 2015: The Forecast Skill Horizon. *Q.J.R. Meteorol. Soc.*, **141**, Issue 693, Part B, 3366–3382.

Buizza, R., M. Leutbecher & A. Thorpe, 2015: Living with the butterfly effect: a seamless view of predictability. *ECMWF Newsletter* **No. 145**, 18–23.

Houtekamer, **P.L.**, **L. Lefraive** & **J. Derome**, 1996: A system simulation approach to ensemble prediction. *Mon. Wea. Rev.*, **124**, 1225–1242.

Lalaurette F. 2003. Early detection of abnormal weather conditions using a probabilistic extreme forecast index. *Q. J. R. Meteorol. Soc.*, **129**, 3037–3057.

Molteni, **F.**, **R. Buizza**, **T.N. Palmer** & **T. Petroliagis**, 1996: The new ECMWF ensemble prediction system: methodology and validation. *Q.J.R. Meteorol. Soc.*, **122**, 73–119.

Palmer, **T.N.**, **R. Buizza**, **F. Doblas-Reyes**, **T. Jung**, **M. Leutbecher**, **G.J. Shutts**, **M. Steinheimer** & **A. Weisheimer**, 2009: Stochastic parametrization and model uncertainty. *ECMWF Technical Memorandum* **No. 598**.

Richardson, **D.S.**, 2000: Skill and economic value of the ECMWF Ensemble Prediction System. *Q.J.R. Meteorol. Soc.*, **126**, 649–668.

Toth, **Z.** & **E. Kalnay**, 1997: Ensemble forecasting at NCEP and the breeding method. *Mon. Wea. Rev.*, **125**, 3297–3319.

Vitart, **F.**, **G. Balsamo**, **R. Buizza**, **L. Ferranti**, **S. Keeley**, **L. Magnusson**, **F. Molteni** & **A. Weisheimer**, 2014: Sub-seasonal predictions. *ECMWF Technical Memorandum* **No. 738**.

Zsótér, **E.** 2006: Recent developments in extreme weather forecasting. *ECMWF Newsletter* **No. 107**, 8–17.

doi:10.21957/l384nb

Assimilating satellite data along a slanted path

NIELS BORMANN

Over the last two decades satellite radiances have come to have the greatest impact on forecasts compared to other types of observations used in numerical weather prediction (NWP). The satellite data are used to determine the initial conditions at the start of a forecast. However, their assimilation into forecast models has traditionally neglected the fact that, most of the time, satellite instruments view the Earth at an angle and therefore sound an atmospheric column that slants through the atmosphere. Data assimilation systems have instead essentially assumed that satellite instruments sense vertical profiles of the atmosphere, only taking into account the increased path length to determine to which vertical layers the radiances relate. With the upgrade of ECMWF's Integrated Forecasting System (IFS) in November 2016 (IFS Cycle 43r1), ECMWF became the first NWP centre to fully take the slanted viewing geometry into account in its operational system. The change has led to improvements in the assimilation of satellite radiances. This has resulted in improved forecast performance that is statistically significant in the short range, particularly in the stratosphere and at higher latitudes. This article gives an overview of what has changed and how forecasts have improved. The interested reader is referred to *Bormann* (2017) for further details.

Satellite viewing geometry

Close to 80% of the observations currently assimilated in the IFS come from passive sounding instruments on satellites, which measure radiation emitted naturally and do not actively transmit their own signals. These observations also have the largest impact on forecasts. The measurements are made by so-called nadir sounders. The viewing geometry of such sounders is illustrated schematically in Figure 1. Despite their name, most of the time nadir sounders do not look directly downwards towards Earth but view the Earth at an angle, as determined by the satellite's zenith angle (Figure 1). Cross-track scanners, for instance, scan the Earth at different viewing angles, with zenith angles varying between 0° (viewing directly downwards, i.e. a nadir view) to 50–60° either side from the nadir.

When the satellite instrument views the atmosphere at a zenith angle greater than 0° (see dashed red line in Figure 1), two effects occur: first, the viewing path through the atmosphere gets longer compared to the nadir view. This affects the position of the layers in the vertical to which a particular channel is sensitive. This effect has always been taken into account in radiative transfer models such as RTTOV (Box A) by appropriately scaling the optical depth of the atmosphere. Second, the atmospheric column sounded by the satellite instrument slants through the atmosphere. This effect has so far been ignored in data assimilation at NWP centres. Instead, the observed radiances from the slanted column have been compared with model equivalents (simulated radiances derived from a short-range forecast) calculated from a single vertical column (represented by the dashed black line in Figure 1). The vertical column's location is based on the geo-location information provided with the data, that is, the position where the instrument's view intersects the Earth's geoid (the surface corresponding to mean sea level). For a channel that senses the atmosphere at higher levels, this means that the model information is extracted in the wrong place. The displacement is most relevant for larger zenith angles and for channels that are mostly sensitive to higher layers

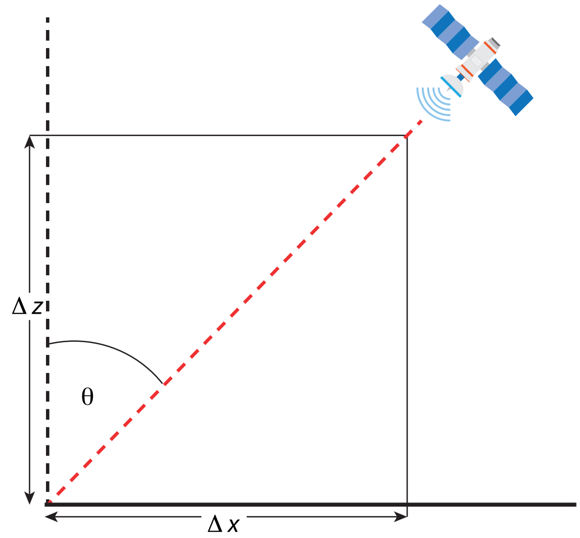

Figure 1 Schematic diagram of the satellite viewing geometry. The dashed black line represents the profile used until recently to describe the atmosphere for radiative transfer calculations in the ECMWF data assimilation system; the red dashed line shows the profile used in slant-path radiative transfer calculations. The angle θ is the satellite zenith angle.

Radiative transfer model

A

The purpose of a radiative transfer model is to determine what kind of radiances would be measured given a particular state of the atmosphere. The radiative transfer model can thus be used to simulate the radiances associated with the modelled state of the atmosphere in a short-range forecast. The data assimilation system compares the observed radiances with the simulated model equivalents and makes adjustments to the model atmosphere to better match the observations provided. A highly accurate radiative transfer model is important in order to make reliable adjustments. The radiative transfer model used at ECMWF is RTTOV, which has been developed by the EUMETSAT NWP SAF. RTTOV stands for Radiative Transfer for TOVS (TIROS Operational Vertical Sounder). TIROS refers to the US Television Infrared Observation Satellite programme.

of the atmosphere. For example, for a channel primarily sensitive to temperature in the lower stratosphere, at around 16 km, and viewing the Earth with a zenith angle of 60°, the displacement will be around 28 km. Such a channel will also have some sensitivity to levels even higher in the stratosphere, and for these levels the displacements will be even larger.

The displacement error can be avoided by making better use of the model atmosphere when we interpolate the model fields to the observation locations. Instead of interpolating to the dashed black line in Figure 1, we can interpolate to the dashed red line and then pass the resulting slanted profile of model information to the radiative transfer model. During the assimilation, the adjustments of model variables can then also be applied to the slanted column rather than the vertical column. These changes were implemented operationally in IFS Cycle 43r1 on 22 November 2016. The slant-path calculations are performed for all sounding radiances used in our clear-sky system. The calculations require knowledge of the satellite's zenith and azimuth angles, which together describe the orientation of the slanted atmospheric column in a three-dimensional atmosphere. This information is usually provided with the observations. The technical framework used is based on earlier developments for two-dimensional observation operators for limb-viewing instruments, such as radio occultation or passive infrared limb radiances (*Healy et al.*, 2007, *Bormann et al.*, 2007). The effect and the size of the displacement in the slant-path case are roughly similar to those encountered from radiosonde drift, which should be taken into account for some sondes in the IFS from Cycle 45r1 onwards.

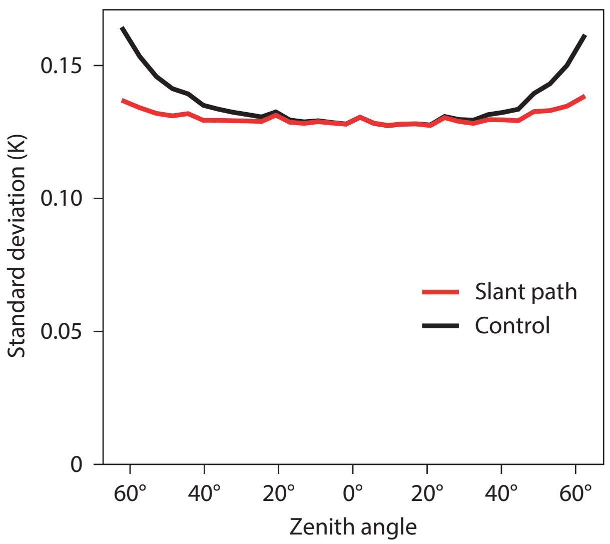

Figure 2 Standard deviations of the differences between observations and model equivalents (observations minus short-range forecast) for ATMS channel 9 as a function of scan-position (labelled here by zenith angle on the x-axis). The previous approach ('Control') leads to larger differences for higher zenith angles compared to when the slant-path effect is taken into account ('Slant path'). The statistics are based on observations after 3x3 averaging, and the same atmospheric background fields were used for the Slant-path and the Control calculations. Data cover a 1-month period in January/February 2015, over sea, after screening for clouds.

Better calculations of model equivalents

Taking the slanting viewing geometry into account leads to better simulations of the satellite observations from the model background (a short-range forecast). Figure 2 shows the standard deviation of the difference between observations and model equivalents as a function of the scan position for a particular channel of the Advanced Technology Microwave Sounder (ATMS). The channel is primarily sensitive to temperature around the tropopause. With the previous approach (black line in Figure 2), larger differences between the observations and the model equivalents are apparent for larger zenith angles. This can be attributed to the displacement errors inherent in the previous approach. The feature is very significantly reduced when we take the slant-path effect into account (red line in Figure 2). The influence of the slanted viewing path is particularly noticeable for this instrument as it is a cross-track scanner with a particularly wide swath, leading to some of the largest zenith angles at the swath edges. There has been a trend towards wider swaths for newer cross-track scanning instruments. This makes it increasingly important to take the slant-path effect into account.

Figure 3 shows the global overall effect of modelling the slanted viewing geometry for a range of instruments and channels. Note that these statistics include all zenith angles. The effect is of course stronger for observations with large zenith angles. The microwave instrument ATMS and the hyperspectral Cross-track Infrared Sounder (CrIS) benefit the most from the slant-path modelling. Relative reductions in the standard deviations of the differences between observations and model equivalents reach nearly 8% and about 2% globally for ATMS and CrIS, respectively. Aside from the relatively wide swath for these instruments, this is also a reflection of lower noise in the observations, so that the displacement errors are relatively more important compared to the instrument noise. While in absolute terms the effect is larger the higher the channel peaks in the atmosphere, this is not necessarily apparent in these relative statistics. This is because instrument noise and errors in the short-range forecasts, which also contribute to the differences between observations and model equivalents, tend to be higher for temperature channels sounding the higher stratosphere. Some effect is also clearly visible for humidity-sounding channels sounding the upper troposphere (e.g. ATMS channels 18–22).

For temperature-sounding channels, the effect is largest over mid- and higher latitudes. This can be seen in Figure 4a, which shows the differences between the previous and the new approach for an ATMS channel peaking around the tropopause. The reason is that the spatial gradients along the viewing direction tend to be particularly large for temperature in these regions. By contrast, for humidity sounding channels sounding the upper troposphere, the reduction in standard deviation is most noticeable around the mid-latitude storm tracks (Figure 4b).

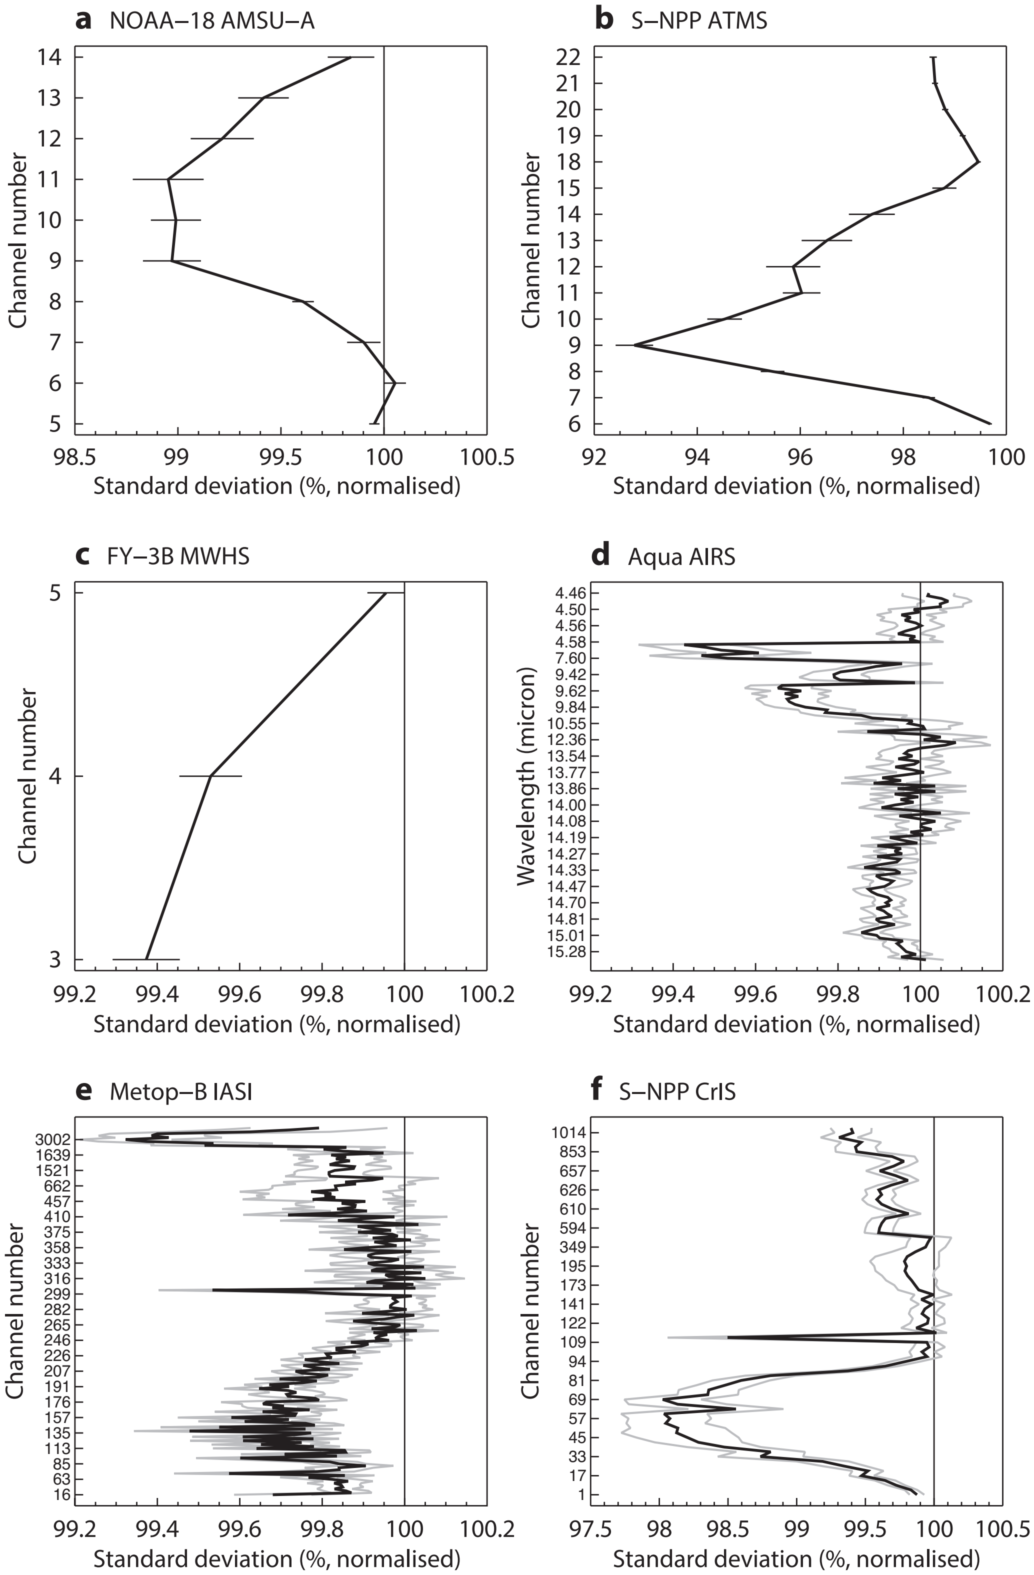

Figure 3 Improvements from slant-path calculations for a range of satellite instruments. The plots show the standard deviations of differences between observations and model equivalents, calculated with the slant-path effect included, normalised by the values obtained using the previous approach. Values under 100% therefore show a reduction in the displacement error in the calculations of the model equivalents. The atmospheric background is the same for both calculations and is taken from a 1-month period in January/February 2015. Statistics are shown for (a) the Advanced Microwave Sounding Unit (AMSU)-A on board the NOAA-18 satellite, (b) ATMS on board the Suomi-National Polar Partnership (S-NPP) satellite, (c) the Microwave Humidity Sounder (MWHS) on board the Feng-Yun-3B satellite, (d) the Atmospheric Infra-Red Sounder (AIRS) on board the Aqua satellite, (e) the Infrared Atmospheric Sounding Interferometer (IASI) on board Metop-B, and (f) CrIS on S-NPP.

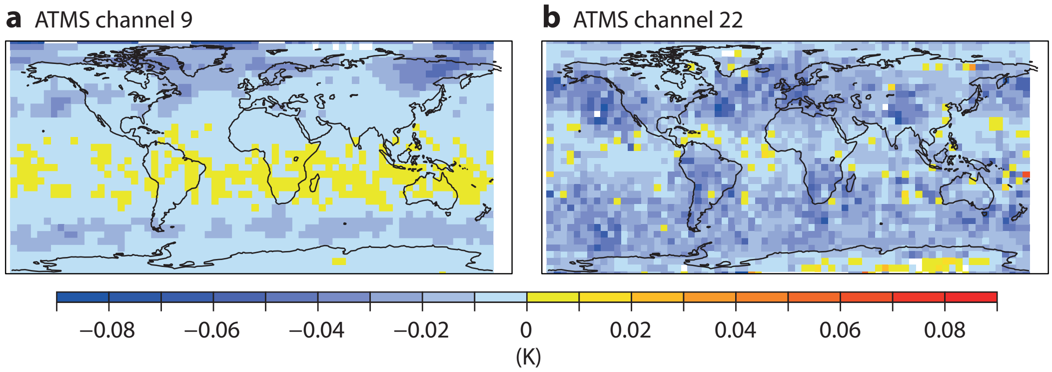

Figure 4 Geographical distribution of improvements in observation modelling. The maps show the change in the standard deviation of differences between observations and model equivalents when introducing the slant-path calculations for (a) ATMS channel 9, a channel primarily sensitive to temperature around the tropopause, and (b) ATMS channel 22, the uppermost tropospheric humidity-sounding channel of ATMS. Negative values (blue) show a closer agreement between model equivalents and observations when using slant-path calculations. For both channels, the statistics are based on observations after 3x3 averaging, and the same atmospheric background fields were used for the slant-path and the previous calculations. Data cover a one-month period in January/February 2015 after cloud screening.

Better forecasts

Assimilation experiments show that accounting for the slant-path effect improves forecast quality. Two experiments were conducted: a 'Control' experiment in which the slant-path effect is neglected, similar to the operational configuration before Cycle 43r1, and a 'Slant path' experiment, in which the slant-path effect is taken into account for all sounder radiances that are not treated in the all-sky assimilation. Radiances treated in all-sky are currently not considered since the treatment of viewing angles is more complex in cloudy conditions, and the interpolation to the slant path may also introduce undesirable smoothing of cloud details. This means the majority of microwave humidity sounders, with the exception of MWHS and ATMS, are excluded from the slant-path treatment. The experiments used ECMWF's 12-hour 4D-Var assimilation system, with a model resolution of TCo639 (16 km), an incremental analysis resolution of TL255 (80 km), and 137 levels in the vertical. Two four-month periods were considered: 2 June – 30 September 2014 and 2 December 2014 – 31 March 2015.

The most notable effect in the assimilation experiments is that overall the analysis makes smaller adjustments to

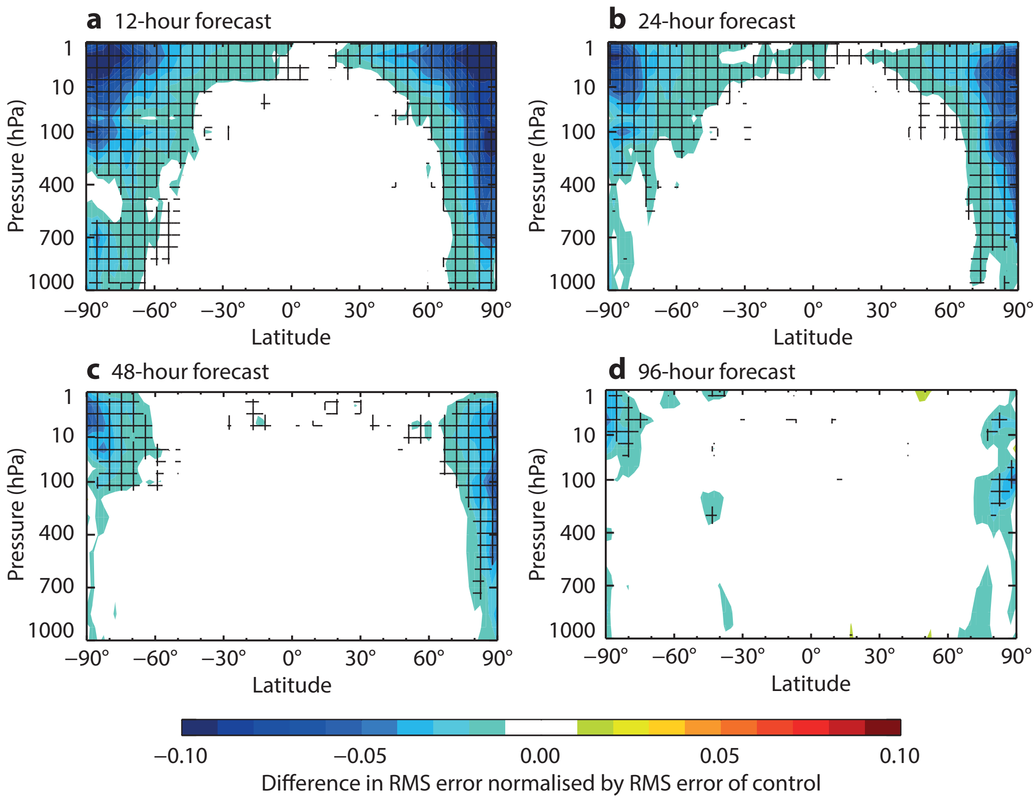

Figure 5 Better short-range forecasts and reduced increments. The panels show the relative change in the zonal mean root-mean-square vector wind error when the slant-path calculations are used in the assimilation for forecast ranges of (a) 12 hours, (b) 24 hours, (c) 48 hours and (d) 96 hours. The 12-hour forecast range also provides a measure of the change in the size of the analysis increments. Statistics are based on forecasts for a total of eight months over the two seasons considered, and each experiment was verified against its own analyses.

the background. The size of the analysis increments (i.e. the difference between the background and the analysis) is reduced by up to around 10% for higher latitudes and higher levels. This is illustrated in Figure 5a, which shows zonal (east–west) means of the relative change in the size of the analysis increments for wind. Other variables show a very similar pattern. As we did not change the observation errors in these experiments, the reduction in the analysis increments is an expected effect of the smaller differences between observations and model equivalents for the sounder radiances shown earlier. It could be argued that the reduction in the displacement error apparent from the previous section means that the assumed observation error used in these experiments should be reduced accordingly for the observations concerned. This would give them more weight and would most likely alter the size of the analysis increments. This has not been pursued here since the uncertainty in the assumed observation errors is probably larger than the error reduction obtained with the slant-path modification. However, taking the slanted path of satellite soundings into account may well make it possible to use a reduced observation error in the future, for instance for ATMS.

The reduced analysis increments lead to some statistically significant reductions in the size of forecast errors up to day 3, especially at high latitudes and in the stratosphere (Figure 5b–c). These improvements can be found in all geophysical variables. The reductions in forecast error beyond day 1 are, however, mostly small and generally less than 1% when averaged over the extra-tropics in the troposphere.

Outlook

Taking the slant-path geometry better into account is an example of the better use that can be made of the full three-dimensional model information that is available to us in an NWP system. Through this work, we improve the ability of the observations to identify and correct errors in short-range forecasts by eliminating a source of error arising from an unnecessary simplification in the interpretation of the full model information. Similar benefits were previously demonstrated in the limb-viewing context (*Healy et al.*, 2007; *a*, 2007), where taking horizontal structure into account also reduced the errors inherent in the assimilation of these observations. Making better use of the full spatial information will become even more important as spatial model resolution increases and uncertainties in observations and short-range forecasts are reduced. For example, changes to be implemented as part of IFS Cycle 45r1 will account for radiosonde drift. For satellite radiances, we intend to investigate how we can better take the extent of the spatial footprint into account: with nominal spatial model resolutions well below the footprint size of satellite radiances, simple interpolations of model variables to a single location (or a slanted line of sight) are likely not to be optimal and explicit modelling of the spatial footprint may be beneficial. Such an approach is expected to be particularly beneficial when assimilating cloud- and rain-affected observations from microwave imagers since it offers ways to better capture sub-footprint size model cloud variability. In addition, three-dimensional effects around clouds can play a significant role, and local three-dimensional radiative transfer calculations may offer benefits in certain situations. These developments will further improve the realism with which we can simulate observations and hence improve our interpretation and use of observations to specify the initial conditions at the start of a forecast.

FURTHER READING

Bormann, **N.**, 2017: Slant path radiative transfer for the assimilation of sounder radiances, *Tellus A*, **69**, 1272779, doi: 10.1080/16000870.2016.1272779.

Bormann, **N.**, **S.B. Healy** & **M. Hamrud**, 2007: Assimilation of MIPAS limb radiances in the ECMWF system. Part II: Experiments with a 2-dimensional observation operator and comparison to retrieval assimilation. *Q. J. R. Meteorol. Soc.*, **133**, 329—346.

Healy, **S.B.**, **J. Eyre**, **M. Hamrud** & **J.-N. Thépaut**, 2007: Assimilating GPS radio occultation measurements with two-dimensional bending angle observation operators. *Q. J. R. Meteorol. Soc.*, **133**, 1213—1227.

doi:10.21957/9fxea2

How to evolve global observing systems

ERIK ANDERSSON

The evolution of global observing systems is driven by user requirements as coordinated by the World Meteorological Organization (WMO). The WMO Integrated Global Observing System (WIGOS) is the combination of satellite-based and surface-based observations that contribute information for WMO programmes. The observations come from satellite programmes managed by space agencies and from a number of observing networks managed by government agencies or commercial operators (Figure 1).

Figure 1 The WMO Integrated Global Observing System (WIGOS) combines the observing capabilities of surface- and space-based platforms and stations to serve a variety of weather and climate needs. (*Graphic: WMO*)

ECMWF works closely with the WMO to make sure the best possible observations are available to the Centre. They are used for data assimilation in ECMWF's Integrated Forecasting System (IFS), which comprises atmosphere, ocean, land, river, sea-ice and atmospheric composition models, as well as for forecast verification. ECMWF is represented in the WMO Expert Teams that gather observation requirements and update them in line with the evolving needs for more complete real-time monitoring of the Earth system and higher accuracy.

As part of a rolling requirements review, Statements of Guidance are produced in several application areas. These feed into a vision document which gives high-level guidance on how observing systems should develop. A 'Vision for WIGOS in 2040' is currently under review. One of the developments acknowledged in the draft Vision is the increasing variety of organisations that are running observing systems of interest to WMO application areas. The aim, strongly supported by ECMWF, is to integrate these observations into one overall system where possible.

ECMWF's involvement

The IFS is observation hungry: observations are needed to initialise the various components of the Earth system model at increasing resolution, with the best possible accuracy and in a timely manner. ECMWF's Strategy 2016–2025 calls for improved Earth system modelling and data assimilation to enable better weather forecasts. This will require more observations of the atmosphere, the deep ocean, the ocean and land surfaces, rivers and lakes, atmospheric composition and sea ice. The associated observation requirements evolve in time as the models' realism improves. Continued progress towards improved weather forecasts thus depends on sustained investment in the capabilities to observe the Earth system.

ECMWF plays an active role in the WMO rolling requirements review process. In addition to its representation in WMO Expert Teams on observations, ECMWF's Deputy Director of Forecasts (the Author) was until recently the Point of Contact in the global numerical weather prediction (NWP) application area and the Rapporteur to the WMO on the scientific evaluation of the

Application area	Point of Contact for latest SoG	Date of latest SoG
Global NWP	Erik Andersson (ECMWF)	June 2016
High Resolution NWP	Thibaut Montmerle (France)	June 2016
Nowcasting and Very Short Range Forecasting	Paolo Ambrosetti (Switzerland)	June 2015
Sub-seasonal to longer predictions	Yuhei Takaya (Japan)	June 2016
Aeronautical Meteorology	Jitze van der Meulen (Netherlands)	June 2016
Forecasting Atmospheric Composition	Oksana Tarasova (WMO)	-
Ocean Applications	Guimei Liu (China)	June 2016
Agricultural Meteorology	Robert Stefanski (WMO)	June 2011
Hydrology	Silvano Pecora (Italy)	July 2014
Climate Monitoring	GCOS Secretariat	2010
Climate Science	WCRP	-
Space Weather	Terry Onsager (USA)	May 2012

Table 1 For each application area, the Statement of Guidance (SoG) provides an assessment of the adequacy of observations to fulfil requirements and suggests areas of progress towards improved use of space-based and surface-based observing systems.

impact of observations on NWP. He now chairs the WMO Inter-Programme Expert Team on the Observing System Design and Evolution (IPET-OSDE). ECMWF's participation is an important avenue to emphasise the importance of the WIGOS observing networks that weather forecasting relies on.

Rolling review

To make sure that all WMO programmes are served by relevant data, observation requirements are gathered in a systematic way. This is done in a similar manner for 12 application areas that each make direct use of observations. Examples of application areas are Global NWP, Regional NWP, Nowcasting and Very Short Range Forecasting, Sub-Seasonal to Longer Prediction, Climate Monitoring, and Hydrology. The full list of application areas is given in Table 1. A WMO application area comprises activities for which it is possible to compile a consistent set of observational user requirements agreed by community experts working operationally in this area.

An expert in each application area serves as the Point of Contact: the conduit between the stakeholder community for that application area to the rolling requirements review (RRR) and the requirements database (OSCAR). The stakeholder community includes national meteorological and hydrological services (NMHSs), WMO Regional Associations, and WMO Technical Commissions and their expert teams.

The RRR process serves to review the evolving requirements for observations and the capabilities of existing and planned observing systems. Through so-called 'Statements of Guidance' (SoG), the expert Point of Contact in each application area addresses the extent to which the capabilities meet the requirements, and they produce gap analyses with recommendations on how these gaps could be addressed. The SoGs are available online at: *www.wmo.int/pages/prog/www/OSY/GOS-RRR.html*.

For each application area, the process consists of four stages:

(i) a review of technology-free requirements for observations within an area of application covered by WMO programmes and co-sponsored programmes;

(ii) a review of the observing capabilities of existing and planned observing systems, both surface- and space-based;

(iii) a Critical Review of the extent to which the capabilities (ii) meet the requirements (i); and

(iv) a Statement of Guidance based on (iii).

This process is repeated in an approximately 18-month cycle. The aim of the SoGs is to inform WMO Members of the extent to which their requirements are met by present systems, will be met by planned systems, or would be met by proposed systems. An SoG is essentially a gap analysis with recommendations on how to address the gaps. The SoGs also serve as a useful resource for dialogue with observing system agencies on whether existing systems should be continued, modified or discontinued; whether new systems should be planned and implemented; and whether research and development is needed to meet unfulfilled user requirements.

For example, the most recent update of the SoG for Global NWP highlights the need for more wind observations,

Summary of Statement of Guidance for global NWP

A

Global NWP centres:

- make use of the complementary strengths of in situ and satellite-based observations;
- have shown strong positive impact from advanced microwave sounding instruments (such as AMSU-A);
- have shown strong positive impact also from high spectral resolution sounders with improved vertical resolution (AIRS, IASI and CrIS);
- have shown strong positive impact from radio occultation data in the upper troposphere and lower stratosphere in particular;
- use 4D data assimilation systems to benefit from more frequent measurements (e.g. from geostationary satellites, aircraft and automated surface stations) and from measurements of cloud, precipitation, ozone, etc.;
- benefit from the improved timeliness of key satellite data resulting from systems such as DBNet;
- would benefit from further increased coverage of aircraft data, particularly from ascent/descent profiles in the tropics;
- are beginning to see the benefits from global dissemination of high-resolution BUFR radiosonde measurements with detailed time-space information;
- would benefit from more timely availability and wider distribution of some observations, in particular several types of in situ measurement and radar that are made but not currently disseminated globally, such as soil wetness, snow depth, precipitation from rain gauges and radar and ground-based GPS;
- would benefit from more ice thickness data and surface salinity.

The critical atmospheric variables that are not adequately measured by current or planned systems are (in order of priority):

- wind profiles at all levels outside the main populated areas;
- temperature and humidity profiles of adequate vertical resolution in cloudy areas, particularly over the poles and sparsely populated land areas;
- satellite based rainfall estimates;
- snow equivalent water content.

particularly in the tropics and the Arctic. It notes that, over most of the Earth, observations of 3D wind fields are "marginal or poor", while coverage of surface wind is "marginal or absent" over some areas in the tropics and the Arctic. The Statement concludes that wind profiles at all levels outside the main populated areas are a top priority among variables that are not adequately measured by current or planned systems. A full summary of the SoG for global NWP is given in Box A.

Based on knowledge of the current and planned observing systems, the gaps identified by the SoGs, and an assessment of which future observing systems are likely to be feasible and affordable, the 'Vision for WIGOS in 2040' (currently under review) provides guidance on the component observing systems to which the WMO community should aspire. A plan to achieve this Vision will subsequently be developed. Currently WMO Members are working towards the 'Implementation Plan for the Evolution of Global Observing Systems', which is based on the vision document for 2020. The implementation plan is available at: *www.wmo.int/pages/prog/www/OSY/Publications/EGOS-IP-2025/EGOS-IP-2025-en.pdf*.

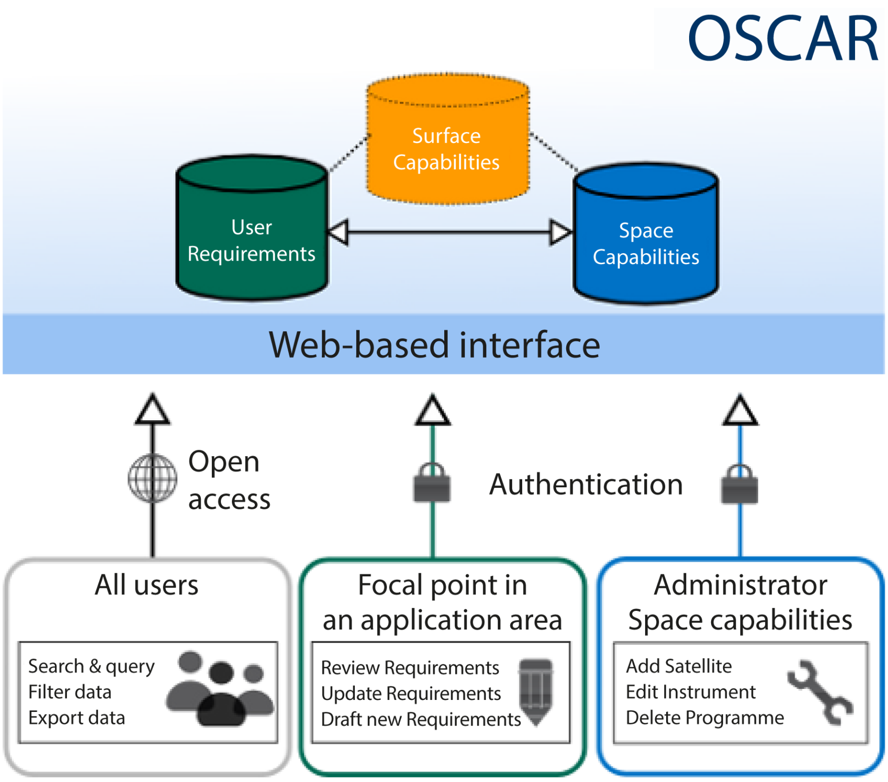

Figure 2 Schematic overview of OSCAR, the WMO Observing Systems Capability Analysis and Review Tool. OSCAR can be accessed online at: *www.wmo-sat.info/oscar*. (*Diagram: WMO*)

Defining user requirements

The user requirements are not system dependent and are intended to be technology-free. No consideration is given to what type of measurement characteristics, observing platforms or data processing systems are necessary (or even possible) to meet them. An online database has been constructed that can be viewed in the context of a given application via a convenient user interface (OSCAR, see Figure 2 and Box B). The requirements for observations are stated quantitatively in terms of five criteria: horizontal and vertical resolution; frequency (observation cycle); timeliness (delay in availability); and uncertainty (acceptable error and any limitations on bias). For each application, there is usually no abrupt transition in the utility of an observation as its quality changes; improved observations (in terms of resolution, frequency, accuracy, etc.) are usually more useful while degraded observations, although less useful, are usually not useless. Moreover, the range of utility varies from one application to another.

Therefore, for each of these criteria, the requirement includes three values determined by experts: the 'goal', the 'threshold', and the 'breakthrough' value.

The 'goal' or 'maximum requirement' is the value above which further improvement of the observation would not cause any significant improvement in performance for the application in question. This is deemed to be the case if the cost of improving the observations beyond the goal would not be matched by a corresponding benefit. The goals are likely to evolve as applications progress and develop a capacity to make use of better observations.

The 'threshold' or 'minimum requirement' is the value that has to be met to ensure that data are useful. Below this minimum, the benefit derived does not compensate for the

What is OSCAR?

B

OSCAR is a resource developed by the WMO in support of Earth observation applications, studies and global coordination.

OSCAR/Requirements contains quantitative user-defined requirements for observations of physical variables in WMO application areas (i.e. related to weather, water and climate). OSCAR/Space provides detailed information on all Earth observation satellites and their instruments and measurement capabilities, while OSCAR/Surface contains information about in-situ observing stations.

Through its three databases OSCAR targets all users interested in the status and the planning of global observing systems as well as data users looking for instrument specifications and the observing capabilities of each platform or station.

User requirements are recorded in tables which can be sorted and filtered, e.g. by variable, application or spatial domain. Variables are defined in a technology-free manner, i.e. without being constrained by space- or surface-based measuring capabilities. They do not necessarily overlap with the direct output of a specific instrument. Requirements for these variables are expressed in terms of the following criteria: uncertainty, horizontal and vertical resolution, observing cycle and timeliness.

For each WMO application area, requirements are directly maintained online by the designated Point of Contact and are regularly reviewed by groups of experts. This process is overseen by the Inter-Programme Expert Team on the Observing System Design and Evolution (IPET-OSDE).

additional cost involved in using the observation. Threshold requirements for any given observing system cannot be stated in an absolute sense; assumptions have to be made concerning which other observing systems are likely to be available.

Within the range between the threshold and the goal, observations become progressively more useful. The 'breakthrough' is an intermediate level between 'threshold' and 'goal' which, if achieved, would result in a significant improvement for the targeted application.

Vision for 2040

The 'Vision for WIGOS 2040' provides high-level goals to guide the evolution of observing systems in the coming decades. These goals are intended to be challenging but achievable. The Vision attempts to address the needs of all application areas with WMO programmes and co-sponsored programmes to which WIGOS responds. The Vision considers that future observing systems will build upon existing sub-systems, both surface- and space-based, while making use of existing, new and emerging observing technologies not presently incorporated or fully exploited. The Vision incorporates observations acquired by commercial operators (surface- and space-based) and acknowledges their importance as well as the challenges involved in ensuring the free and open exchange of such data between NMHSs.

The Vision document has recently been drafted and is currently going through an extensive review process amongst WMO Members and Expert Teams with the aim of final approval in 2019.

The draft Vision acknowledges that NMHSs are no longer the sole providers of ground-based meteorological observations. Instead, typically a variety of organisations are now running observing systems of interest to WMO application areas. These may be different government agencies operating under the ministries of agriculture, energy, transport, tourism, environment, forestry, water resources, etc. Especially in developing countries, they may be non-profit organisations or commercial entities. It is a principle of WIGOS to integrate these observations into one overall system as far as possible.

On the satellite side, the Vision retains a strong focus on operational geostationary and polar-orbiting platforms. In addition it considers the possibilities of instruments in Highly Elliptic Orbits (HEO) that would permanently cover the polar regions; Low-Earth Orbit (LEO) satellites with low or high inclination for a comprehensive sampling of the global atmosphere; and lower-flying platforms, for example small satellites serving as gap fillers or for dedicated missions which are best realised that way.

Amongst surface-based platforms, the Vision includes the traditional networks as well as new opportunities from automated low-cost observations collected from mobile phones and cars for example, which have the potential to provide a wealth of information in urban areas in particular.

Conclusion

The global observing system needs to evolve constantly to meet changing user requirements in all WMO application areas. One of the challenges WIGOS faces is the growing variety of potential observation providers, including commercial entities. ECMWF plays an active role in the WMO-led process to define user requirements and develop guidance on future observing systems. In this way the Centre can help to ensure that WIGOS provides the observations that are needed for continued improvements in global numerical weather prediction.

doi:10.21957/5t18tb

RMDCN upgrade nears completion

TONY BAKKER

The Regional Meteorological Data Communication Network (RMDCN) is currently nearing completion of an upgrade of the connections of many sites following negotiations between ECMWF and Interoute Communications Ltd in 2016 as part of a technical and commercial refresh exercise.

The RMDCN provides a computer network infrastructure for the meteorological community in World Meteorological Organization (WMO) Region VI and beyond. It was set up in 2000 and provides any-to-any connectivity between more than 50 sites (Figure 1). Among other things, the RMDCN serves to ensure the secure and timely delivery of ECMWF forecasts to its Member States and the exchange of meteorological observations between connected sites. In the framework of the WMO/ECMWF agreement on the RMDCN, the ECMWF project team manages the network and monitors the Quality of Service on a 24-hour basis for all participating centres.

Upgrade requirements

ECMWF's contract with Interoute Communications Ltd, authorised by the Council in December 2012, covers a nine-year term with the option to break after six years. To ensure ongoing value for money during this period, technical and commercial refresh (TCR) exercises are undertaken at 2½ and 5½ years following the start of the operational service on 30 June 2013. As part of the first of these TCR exercises, various elements of the service were reviewed according to technical and commercial criteria and overall value for money.

In order to support ECMWF in this process, the ECMWF Technical Advisory Committee (TAC) at its 47th session on 15 and 16 October 2015 established a subgroup on the RMDCN with Graham Mallin (United Kingdom) as Convener and with the following terms of reference:

1. to assist ECMWF with the specification of technical requirements to be considered for the Technical and Commercial Refresh of the RMDCN contract with Interoute;
2. to assist ECMWF in reviewing Interoute's offer for the Technical and Commercial Refresh;
3. to comment on any contractual changes required to implement the outcome of the Technical and Commercial Refresh;
4. to involve WMO/Regional Association VI and the WMO Secretariat as observers in the subgroup, as appropriate.

ECMWF, with the assistance of the TAC Subgroup, discussed various technical options to be included in the TCR, such as Internet backup (i.e. DMVPN), Internet Protocol version 6 (IPv6), Multicast, and Cloud Computing. It also reviewed the ECMWF-funded RMDCN Basic Package for ECMWF Member States. During the lifetime of the RMDCN, the Basic Package speed for ECMWF Member States has doubled roughly every three years while keeping the monthly recurring charges the same. When the Interoute service was deployed in 2013, the Basic Package was a Platinum service type with a 4 Mbps connection, which was a doubling of the speed compared to the Basic Package provided by the previous supplier, Orange Business Services Ltd. An upgrade of the Basic Package with a doubling of the speed to 8 Mbps while maintaining the current monthly charge was therefore in line with established practice concerning upgrades of the Basic Package.

The TAC Subgroup agreed on the following set of requirements for Interoute to review:

- investigate the current technical deployment of the network and report on any potential improvements (e.g. IPv6)
- investigate and propose the deployment of an IP multicast solution to meet a EUMETSAT requirement
- investigate and propose the deployment of an Internet backup solution to replace the current DMVPN pilot
- provide information on Interoute's cloud service portfolio available within the RMDCN footprint
- investigate the current monthly recurring charges for all sites and make proposals for both price reductions and increased bandwidth at the current charges
- provide pricing for an upgrade to 8 Mbps for ECMWF Member States where applicable.

All sites with unchanged configuration since the Operational Commencement Date (30 June 2013) were offered a discount of at least the minimum level as defined in the TCR clause of the contract. Discounts apply to the charges for Interoute service elements and exclude underlying access circuits.

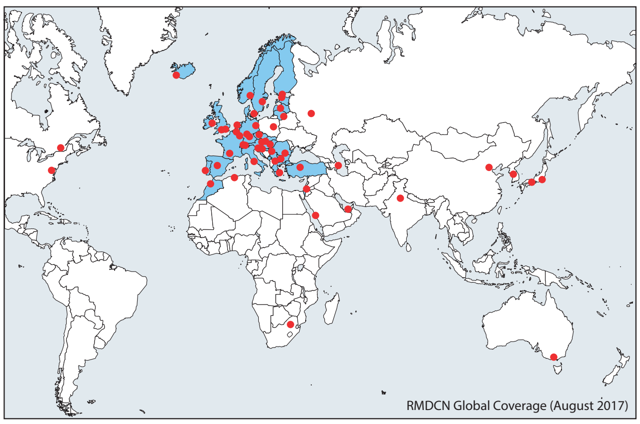

Figure 1 Currently 55 sites are connected to the RMDCN network. The shaded countries indicate ECMWF Member and Co-operating States.

Towards implementation

Following the receipt of Interoute's opening offer, ECMWF commissioned the consultancy The Network Collective (TNC) to conduct an assessment of Interoute's offer in the current market. Their findings can be summarized as follows:

- Overall the offer sits very comfortably in the lower to middle of TNC's market spread, so whilst there is room for improvement compared to some suppliers, it is nevertheless a relatively competitive deal and would beat much of the market.
- Due to the nature of the benchmarking clause within ECMWF's contract with Interoute, the focus for ECMWF should be on those elements considered to be materially out of sync with current market pricing.
- Commercially the areas in which to optimise pricing are China, Australia and, to a lesser extent, Japan. There are also more minor concerns for pricing in Italy, Israel, Morocco, Turkey, and EUMETSAT's connection in Germany.

ECMWF contacted Interoute and discussed the report. Interoute acknowledged the findings and was able to address most areas of concern. In addition to Interoute's TCR offer, the TAC Subgroup also discussed the term of the contract between ECMWF and Interoute. The contract with Interoute is for a nine-year term with a break clause at year six (Table 1).

The break clause was agreed in order to allow termination of the contract if the TCR did not show value for money for the service. Given the long lead time required to start a tendering process, preparation for this would have to start in early 2017 if it was decided to execute the break clause. The TNC report reassured the TAC Subgroup that the contract is still "a competitive deal and would beat much of the market". In addition, the RMDCN Operations Committee reported in the 22nd Operations Committee meeting that they were very content with the service provided by Interoute and the high level of service availability during the first 2½ years of service.

The TAC Subgroup on the RMDCN met on 20 and 21 September 2016 and concluded that:

1. The price/performance of Interoute's offer is a good package and represents value for money for the RMDCN community as a whole.
2. The existing contract may be amended at an appropriate point in the future as additional services become sufficiently mature for inclusion in the RMDCN, for example DMVPN.
3. Access to the RMDCN from Cloud Computing systems should be subject to a design and security review to avoid adverse impact on the RMDCN operational service.
4. ECMWF should continue with the following actions independently of the TAC Subgroup:
   (i) request that RMDCN user sites which have been offered a choice between a discount and an increase in bandwidth respond to this offer and identify their required changes to the current service based on this offer;
   (ii) consult with the RMDCN community to establish requirements for a 24/7 operational DMVPN service provided by Interoute and take forward development of this service;
   (iii) support EUMETSAT and Interoute in discussing potential solutions based on EUMETSAT's requirements for a Multicast service on the RMDCN;
   (iv) review Contract Clause 11 with a view to establishing a clearer framework for future service changes based on market rates.

The Subgroup recommended to the TAC that the service should continue for the full nine-year term, that the break clause should not be executed, and that preparations for returning to the market be initiated by the TAC in 2019. During its 48th session from 13 to 14 October 2016, the TAC unanimously endorsed the conclusions of its Subgroup on the RMDCN.

Implementation

ECMWF then started the process of implementing the conclusions of the TAC Subgroup, beginning with gathering the wishes from RMDCN members for either a discount or an upgrade of their connection. Of the 55 sites connected to the RMDCN, 19 sites opted for an upgrade of their connection speed and 23 sites opted for a discount on the monthly charge. For 12 sites there was no change as they had either changed their configuration after the Operational Commencement Date or connected to the RMDCN after this date. One site (USA) decided to terminate its connection to the RMDCN.

Once this was done the change orders were raised with Interoute by the end of December 2016. Monthly charge discounts were backdated to 13 September 2016, when

Timeline	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9
Full contract	Nine-year full contract term								
End of year 6 break clause	Six-year fixed contract terms						Break clause option to full contract		
Technology cycle	RMDCN three-year technology cycle			RMDCN three-year technology cycle			RMDCN three-year technology cycle		

Table 1 The full RMDCN contract term is nine years but a break clause enables termination of the contract after six years.

Interoute's original offer was made. Some sites requested additional upgrade requirements beyond what was presented in the TCR offer from Interoute. By 1 September 2017 about half of the upgrades had been completed with the remaining sites scheduled to be upgraded within the following few weeks. The overall delivery process of the TCR has taken much longer than expected. Discussions are taking place with Interoute to review the TCR and delivery issues.

The overall operational performance of the network remains very high, achieving 100% global availability for 10 of the last 12 months of service. In late September 2017 the 23rd RMDCN Operating Committee Meeting took place. This gave the whole user community an opportunity to discuss and reflect on this TCR exercise and discuss with The Network Collective consultancy the future of the service in general.

Table 2 shows the current configuration of the connected sites, with the bandwidth column showing the current speed in Mbps. The TCR upgrade column shows which option each site selected.

Outlook

Following the decision to continue for the full nine-year term, the contract with Interoute will now terminate on 30 June 2022. This means that a second TCR of the service will take place in late 2018 to early 2019 for implementation by mid-2019. Following this second review, ECMWF and the RMDCN community will have to start discussions on the future of the RMDCN beyond 2022.

One of the key changes that will take place in the near future is the move of the ECMWF data centre to Bologna, which means that ECMWF's RMDCN connection will move to Bologna as well. The move of the RMDCN connection to Bologna is scheduled for late 2019.

Questions about the Technical and Commercial Refresh or any other aspect of the RMDCN service can be sent to ECMWF's service manager, Tony Bakker, at *rmdcn@ecmwf.int*.

Country/Site	City	WMO GTS	WMO WIS	Bandwidth (Mbps)	Site Type	Joined Interoute service	TCR upgrade (Mbps)
ECMWF Member States & ECMWF							
Austria	Vienna	RTH		10	Gold	May 2013	From 8 to 10
Belgium	Brussels	NMC		10	Platinum	May 2013	Discount
Croatia	Zagreb	NMC		8	Gold	Jan 2014	From 4 to 8
Denmark	Copenhagen	NMC		50	Platinum	Jan 2014	Discount
Finland	Helsinki	NMC		50	Platinum	Jan 2014	No change
France	Toulouse	RTH	GISC	100	Platinum	Jan 2014	Discount
Germany	Offenbach	RTH	GISC	50	Platinum	Jan 2014	Discount
Germany-DR	Berlin			N/A	Iron-B	Mar 2016	Discount
Greece	Athens	NMC		4	Platinum	Jan 2014	From 4 to 8
Iceland	Reykjavik	NMC		2	Platinum	Jan 2014	From 2 to 5
Ireland	Dublin	NMC		100	Platinum	Jan 2014	Discount
Italy	Rome	RTH		15	Gold	Jan 2014	From 10 to 15
Luxembourg	Luxembourg	NMC		8	Platinum	Jan 2014	Discount
Netherlands	De Bilt	NMC		8	Platinum	Jan 2014	From 4 to 8
Netherlands-DR	Woensdrecht			4	Copper	Jan 2014	From 4 to 8
Norway	Oslo	NMC		25	Platinum	Jan 2014	From 20 to 25
Portugal	Lisbon	NMC		8	Platinum	Jan 2014	From 4 to 8
Serbia	Belgrade	NMC		4	Gold	Jan 2014	From 4 to 8
Slovenia	Ljubljana	NMC		8	Platinum	Jan 2014	From 4 to 8
Spain	Madrid	NMC		4	Platinum	Jan 2014	From 4 to 8
Sweden	Norrköping	RTH		10	Platinum	May 2013	Discount
Switzerland	Zurich	NMC		20	Platinum	Jan 2014	No change
Switzerland-CSCS	Lugano			20	Platinum	Jan 2014	No change
Turkey	Ankara	NMC		34	Platinum	Jan 2014	No change
United Kingdom	Exeter	RTH	GISC	20	Platinum	Jan 2014	Discount
ECMWF	Reading	WMC		500	Platinum	May 2013	No change

Country/Site	City	WMO GTS	WMO WIS	Bandwidth (Mbps)	Site Type	Joined Interoute service	TCR upgrade (Mbps)
ECMWF Co-operating States & EUMETSAT							
Bulgaria	Sofia	RTH		10	Silver	May 2013	Discount
Czech Republic	Prague	RTH		7	Gold	Jan 2014	Discount
Estonia	Tallinn	NMC		2	Silver	Jan 2014	Discount
EUMETSAT	Darmstadt			20	Platinum	Jan 2014	Discount
Hungary	Budapest	NMC		8	Platinum	Jan 2014	No change
Israel	Bet Dagan	NMC		15	Platinum	Jan 2014	From 10 to 15
Latvia	Riga	NMC		2	Silver	Jan 2014	From 1 to 2
Lithuania	Vilnius	NMC		1	Silver	Jan 2014	Discount
The former Yugoslav Republic of Macedonia	Skopje	NMC		N/A	Iron B	Jan 2014	No change
Morocco	Casablanca	NMC	GISC	2	Bronze	Jan 2014	Discount
Romania	Bucharest	NMC		15	Platinum	Jan 2014	From 10 to 15
Slovakia	Bratislava	NMC		1	Gold	Jan 2014	Discount

Country/Site	City	WMO GTS	WMO WIS	Bandwidth (Mbps)	Site Type	Joined Interoute service	TCR upgrade (Mbps)
Other RMDCN Members							
Algeria	Algiers	NMC		N/A	Iron B	Mar 2017	No change
Australia	Melbourne	WMC	GISC	4	Platinum	May 2014	Discount
Azerbaijan	Baku	NMC		2	Copper	Jul 2017	No change
Canada	Dorval	NMC		2	Copper	Jan 2014	Discount
China	Beijing	RTH	GISC	16	Platinum	Jan 2014	Discount
China-DR	Beijing			16	Copper	Apr 2015	No change
India	New Delhi	RTH	GISC	6	Platinum	Jan 2014	From 4 to 6
Japan	Tokyo	RTH	GISC	10	Platinum	May 2013	Discount
Japan-DR	Osaka			10	Copper	Dec 2014	Discount
Jordan	Amman	NMC		1	Iron A	Jan 2014	No change
Poland	Warsaw	NMC		1	Silver	Jan 2014	Discount
Russian Federation	Moscow	WMC	GISC	10	Platinum	Jan 2014	Discount
Saudi Arabia	Jeddah	RTH	GISC	2	Bronze	Aug 2014	From 2 to 4
South Africa	Pretoria	RTH	GISC	4	Bronze	Jan 2014	From 2 to 4 (primary only)
South Korea	Seoul	NMC	GISC	4	Platinum	Jan 2014	From 4 to 6
United Arab Emirates	Abu Dhabi	NMC		N/A	Iron B	Sep 2014	No change
United States of America	Washington	WMC	GISC	50	Platinum	Jan 2014	To be terminated Dec 17

Table 2 The current RMDCN configuration including the TCR upgrade as of 1 September 2017. Platinum, Gold and Silver site types have dual connectivity. Copper and Iron A/B site types have a single connection. N/A signifies an Internet connection of unknown speed.

Acronyms:

GTS = Global Telecommunication System
WIS = WMO Information System
RTH = Regional Telecommunication Hub
NMC = National Meteorological Centre
WMC = World Meteorological Centre
GISC = Global Information System Centre
DR = Disaster Recovery
CSCS = Swiss National Supercomputing Centre

ECMWF Calendar 2017/18

Nov 13–16	Workshop on shedding light on the greyzone
Nov 13–17	5th International Conference on Reanalysis (ICR5) (Rome)
Nov 28–29	Workshop on developing Python frameworks for Earth system sciences
Dec 4–6	ECMWF/ESA Workshop on Using Low-Frequency Passive Microwave Measurements in Research and Operational Applications
Dec 7–8	Council
Dec 8	Symposium for Adrian Simmons
Jan 22–25	Workshop on SST and sea-ice observations and analysis for NWP and climate applications
Jan 22–26	Computer user course: ecFlow
Jan 26	Symposium to mark 20 years of 4D-Var
Jan 29 – Feb 2	Training course: Use and interpretation of ECMWF products
Feb 5–9	Training course: Use and interpretation of ECMWF products
Feb 19–22	Computer user course: ecCodes, BUFR
Feb 26 – Mar 1	Computer user course: ecCodes, GRIB
Mar 12–16	NWP training course: Data assimilation
Mar 19–23	EUMETSAT/ECMWF NWP-SAF training course: Satellite data assimilation
Apr 9–13	Advisory Committee for Data Policy and data policy meetings of EUMETSAT and ECOMET
Apr 16–20	NWP training course: Advanced numerical methods for Earth system modelling
Apr 23–27	NWP training course: Physical parametrization of sub-grid scale processes
Apr 24	Policy Advisory Committee
Apr 25–26	Finance Committe
Apr 30 – May 4	NWP training course: Predictability and ensemble forecast systems
May 15–16	Security Representatives' meeting
May 16–18	Computing Representatives' meeting
May 12–24	Workshop on radiation in NWP models
Jun 4–7	Using ECMWF's Forecasts (UEF)
Jun 13–14	Council
Jul 10–12	Workshop on physics–dynamics coupling (PDC18)
Sep 10–13	Annual Seminar
Sep 24–28	Workshop on high-performance computing in meteorology
Oct 1–3	Training course: Use and interpretation of ECMWF products
Oct 8–10	Scientific Advisory Committee
Oct 11–12	Technical Advisory Committee
Oct 22–23	Finance Committee
Oct 24	Policy Advisory Committee
Dec 4–5	Council

ECMWF publications

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

Technical Memoranda

808 **Weaver**, **A.T.**, **S. Gürol**, **J. Tshimanga**, **M. Chrust** & **A. Piacentini**: "Time"-parallel diffusion-based correlation operators. *August 2017*

807 **Ingleby**, **B.**: An assessment of different radiosonde types 2015/2016. *August 2017*

806 **Buizza**, **R.**, **E. Andersson**, **R. Forbes** & **M. Sleigh**: The ECMWF research to operations (R20) process. *July 2017*

Index of Newsletter articles

This is a selection of articles published in the *ECMWF Newsletter* series during recent years.
Articles are arranged in date order within each subject category.
Articles can be accessed on ECMWF's public website – http://www.ecmwf.int/en/research/publications

Contact information

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

Telephone National 0118 949 9000

Telephone International +44 118 949 9000

Fax +44 118 986 9450

ECMWF's public website http://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).

Problems, queries and advice	Contact
General problems, fault reporting, web access and service queries	servicedesk@ecmwf.int
Advice on the usage of computing and archiving services	advisory@ecmwf.int
Queries regarding access to data	data.services@ecmwf.int
Queries regarding the installation of ECMWF software packages	software.support@ecmwf.int
Queries or feedback regarding the forecast products	forecast_user@ecmwf.int

Newsletter | Number 153 – Autumn 2017
European Centre for Medium-Range Weather Forecasts
www.ecmwf.int