ECMWF Feature article

from Newsletter Number 154 – Winter 2017/18

METEOROLOGY

ECMWF's new long-range forecasting system SEAS5

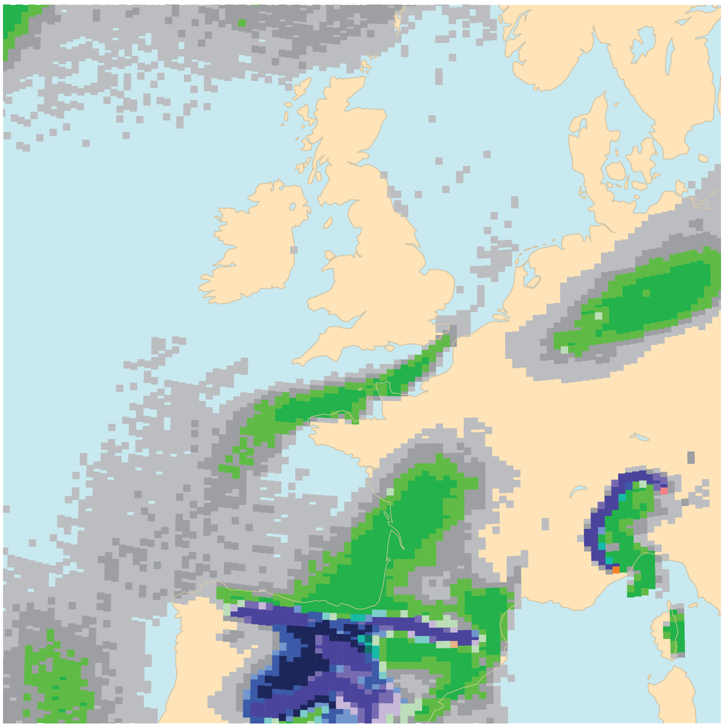

www.ecmwf.int/en/about/news-centre/media-resources
doi:10.21957/tsb6n1

This article appeared in the Meteorology section of ECMWF Newsletter No. 154 – Winter 2017/18, pp. 15-20.

ECMWF's new long-range forecasting system SEAS5

Tim Stockdale, Stephanie Johnson, Laura Ferranti, Magdalena Balmaseda, Simona Briceag

On 5 November 2017, the latest generation of ECMWF's seasonal forecasting system, SEAS5, became operational. As the name suggests, this is the fifth system we have run at ECMWF to produce real-time seasonal forecasts. Seasonal forecasts provide predictions of how the average atmospheric, ocean and land surface conditions over particular areas and periods of time are likely to be different from the long-term average. They are useful to a number of sectors, such as agriculture, water management, energy and health, and they can help to prepare for potential periods of extreme conditions.

In SEAS5 a number of upgrades have been implemented, in particular in the ocean model, atmospheric resolution, and land surface initialisation. The new configuration also represents a move towards a seamless approach to forecasting across timescales. SEAS5 forecasts show substantial improvements in the tropics, in particular for sea-surface temperature in the equatorial Pacific.

Unlike the configurations of ECMWF's Integrated Forecasting System (IFS) used to produce medium-range high-resolution and ensemble forecasts (HRES and ENS), which are typically upgraded at least once a year, SEAS is upgraded only occasionally, at intervals of four to six years. This slow refresh cycle is partly due to the resources needed to complete the large re-forecast sets required for calibration, and partly to offer users a more stable service. It is possible that this approach might change at some point in the future. We discuss the future evolution of seasonal forecasting at the end of this article.

From System 4 to SEAS5

SEAS5 replaces 'System 4', referred to here as S4. SEAS5 uses IFS Cycle 43r1 and represents six years of IFS development in terms of physics, numerics, new Earth system components and initialisation methods. Many of these improvements were focused on our medium-range forecasts and have been described elsewhere. We focus here on the aspects most important for SEAS5 and on improvements that have been made specifically for the long-range forecasts.

ORCA025, sea-ice model and ORAS5

As for S4, SEAS5 uses the community ocean model NEMO (Nucleus for European Modelling of the Ocean), but with an upgraded model version, ocean physics and resolution. The resolution has been increased from 1 degree and 42 layers in S4 to 0.25 degrees and 75 layers in SEAS5 (ocean model configuration ORCA025z75). The vertical resolution is particularly high in the uppermost part of the ocean, with an increase in the number of levels in the first 50 metres from 5 to 18. The increase in horizontal resolution improves the representation of sharp fronts and ocean transports. The vertical resolution increase means that the diurnal cycle of sea-surface temperatures (SST) is much better captured, with a 1-metre top level in the new configuration compared to the previous 10-metre top level. The high-resolution ocean model is used both by SEAS5 and the ENS medium-range forecasts.

An important innovation in SEAS5 is the inclusion of prognostic sea-ice. The sea-ice model is LIM2, part of the NEMO modelling framework. In S4, the sea ice was prescribed by sampling the recent history of sea-ice occurrences. The prognostic sea-ice model allows sea-ice cover to respond to changes in the atmosphere and ocean states. The intention is to capture interannual variability and trends in the sea-ice cover. Therefore SEAS5 provides seasonal outlooks of sea-ice cover, which is a product of interest for users.

SEAS5 ocean and sea-ice initial conditions are provided by the new ocean analysis and reanalysis ensemble (ORAS5). ORAS5 uses the same ocean model and sea-ice as the coupled forecasts in SEAS5 and is driven by ocean observations from floats, buoys, satellites and ships. It consists of five ensemble members covering the period 1975 to the present. The ensemble provides information on the uncertainty associated with the (re)analysis. Compared to its predecessor ORAS4, which was used for S4, ORAS5 has higher resolution and updated data assimilation and observational datasets. It provides sea-ice initial conditions by assimilating sea-ice concentration. The underlying SST analysis before 2008 comes from the HadISST2 dataset, the same as that used in the ERA5 climate reanalysis, which is currently in

 doi:10.21957/tsb6n1

production at ECMWF. From 2008 onwards, the SST and sea ice are given by the OSTIA product, which is used in ECMWF's operational analysis for numerical weather prediction (NWP).

An improved perturbation scheme is used to generate the ensemble of ocean reanalyses. The scheme consists of two distinct elements: perturbations to the assimilated observations, both profiles and surface observations, and perturbations to the surface forcing fields.

Atmospheric resolution upgrade

Horizontal resolution in the atmospheric component of SEAS5 is also significantly higher, increasing from TL255 in S4 to TCo319 in SEAS5. The corresponding grid-point resolution has increased from 80 km to 36 km. Although the spectral resolution increase is less dramatic than the change in grid-point resolution, the ability of the cubic grid to better represent the smallest spectral scales and the energy within them more than makes up for this. The wave model resolution has increased from 1 degree to 0.5 degrees to maintain a match to the atmosphere. The vertical resolution remains at L91. With these resolution changes, the IFS resolution in SEAS5 now exactly matches that used in the extended part of ENS, which covers the 15–45 day time range.

Land surface initialisation

A key requirement for seasonal forecasting is that the initialisation of the re-forecasts is consistent with that of the real-time forecasts, otherwise the calibration of the forecasts becomes invalid. This is a particular challenge for the land surface, where real-time analyses and reanalyses have very different resolutions of the heterogeneous land surface. Values from the ERA-Interim climate reanalysis cannot be used anyway because they come from an incompatible land surface model.

In S4 the problem was partially addressed by a custom offline recalculation of the land surface for the ERA-Interim period, which subsequently became known as 'ERA-Interim Land'. This provided initial conditions compatible with the land surface model used in S4, but the mismatch in horizontal resolution between real-time and reanalysis remained severe, and there were also problems with the precipitation forcing used. For SEAS5 a new offline recalculation of the land surface initial conditions was made, at the required TCo319 resolution and with a revised precipitation forcing. Comparison of test forecasts made using this dataset with others using the operational analysis for a recent overlap period showed a generally very good degree of consistency, while also demonstrating the superiority of the operational analysis in terms of the impact on 2 m temperature forecast anomalies. Thus this configuration was used in SEAS5, despite the fact that the consistency between past and present is still not perfect. Similar land surface initial conditions are also now used for the ENS re-forecasts. In future, we expect to be able to create consistent land surface reanalyses on demand, further improving the consistency and reliability of the land surface and snow cover initial state.

Working towards a seamless approach

SEAS5 is a move towards a seamless approach to forecasting across timescales at ECMWF. Our initial goal is, to the extent possible, to minimise the number of IFS configurations used and ensure that the model is run consistently across timescales. SEAS5 is thus configured almost identically to Cycle 43r1 of the extended-range ENS. There are a few small differences which were included because of their perceived relevance or importance for longer-range forecasting. In SEAS5 the tropospheric sulphate aerosol forcing has a decadal time variation, using the same values as in ERA5. Sulphate aerosol has greatly reduced over Europe in the last 30 years, while it has increased over regions such as China, and using a present-day climatology for all of the re-forecasts would slightly affect their accuracy, and thus the calibration of the real-time forecasts. In future it should be possible to find a unified treatment of time-varying sulphate aerosols for all our forecast configurations. SEAS5 also uses the same treatment of time-varying volcanic aerosol as was used in S4 – damped persistence of a highly simplified specification of an initial state. In principle, this method allows us to respond in real time to any large volcanic eruption: we simply add an estimate of the optical depth of the aerosol from a large volcano after it has happened. Although forecasts from before the eruption will be invalid, we have some capability to adjust the real-time system to predict the expected impact once the eruption has occurred. However, our ability to handle a future large eruption is still fairly rudimentary, and we would expect substantial errors to occur, particularly in the response of the northern hemisphere winter circulation.

A final difference from ENS is that the tropical non-orographic gravity wave drag was retuned to ensure a reasonable phase evolution of the stratospheric quasi-biennial oscillation (QBO). This retuning, which is resolution dependent, has now been applied to the ENS. This is an example of how we expect our seamless approach to work: even if the same configuration choices (e.g. horizontal and vertical resolution, or the balance of resources between different parts of the Earth system model) may not be quite optimal for different forecast timescales, by reducing the number of effective configurations we can become more efficient in identifying and fixing problems, to the benefit of forecast quality at all timescales.

Operational implementation

Important changes have been made to how seasonal forecasts are run operationally, so as to enhance the service given to users. The release date has been brought forward from 12 UTC on the 8th day of each month to 12 UTC on the 5th day of each month. The re-forecast period has been increased to cover 1981–2016, a 36-year period, compared to 1981–2010 for S4. When it comes to assessing past performance, the longer the period available the better. The re-forecast ensemble size has also been increased from 15 to 25, which reduces sampling uncertainty in assessing performance, especially in the mid-latitudes.

Users accessing the plots on our website will see a restructuring of the plots, with the verification plots now grouped in separate families. Full verification based on the 1981–2016 re-forecasts is available. There are also some minor enhancements to the range of forecast plots available, most notably the addition of SST anomaly plumes for the NINO1+2 region, which is important for Peru and Ecuador.

A final important change is that, although the verification is made using the full 1981–2016 re-forecast period, the charts of fields such as 2 m temperature and precipitation are presented as anomalies relative to the more recent 1993–2016 period. This is to ensure that the anomalies remain relevant in the context of climate change. Temperatures (and some other fields such as geopotential height) have changed sufficiently over the last 36 years for seasonal mean values to be almost always warm compared to the early years of the re-forecast period. The result is that all too often the temperature forecasts are that it will always be warm everywhere, relative to an increasingly distant past. This may be a correct probabilistic statement about next season's expected weather, but it is of limited use to a typical user who wants to know what to expect relative to a more recent past. We do not want to use too short a reference period, which would bring its own problems regarding stability and sampling issues, but the 24-year period 1993–2016 seems broadly appropriate. It is also consistent with how the EU-funded Copernicus Climate Change Service (C3S) being implemented by ECMWF will present its new multi-model forecasts (see next section). For users who want to calibrate and reference the SEAS5 forecasts in ways specific to their own application, the full 36 years of re-forecast data remains available.

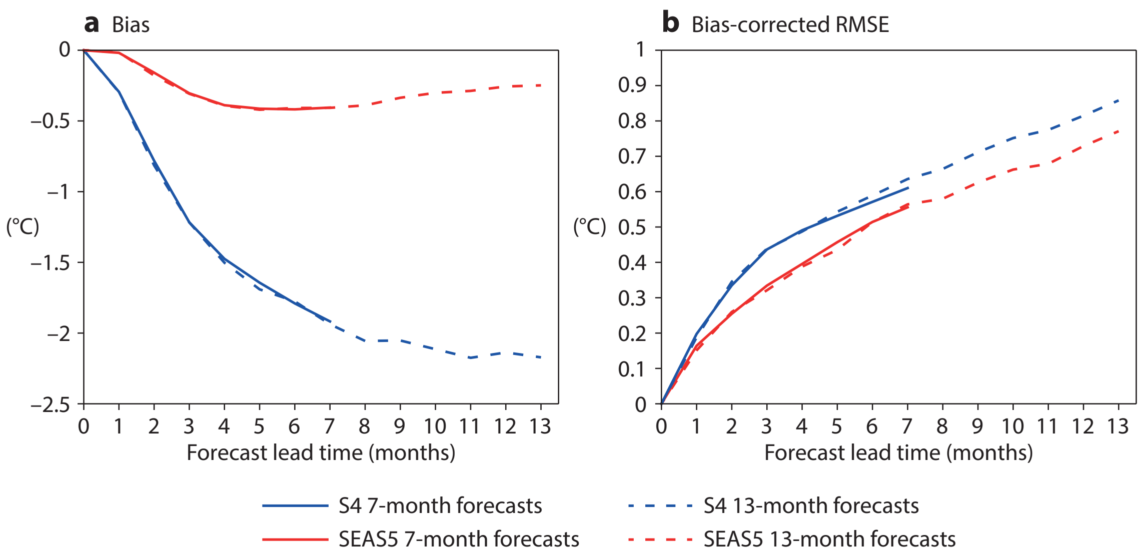

Figure 1 Verification charts for SST anomaly forecasts in the NINO3.4 region showing (a) bias and (b) bias-corrected RMSE for S4 and SEAS5. The long-range forecast (out to 7 months, produced monthly) is shown by the solid line, and the annual-range forecast (out to 13 months, only produced quarterly) is shown by the dashed line. All re-forecast start dates in 1981–2016 are included, a total of 432 dates for the 7-month forecasts and 144 dates for the 13-month forecasts.

 doi:10.21957/tsb6n1

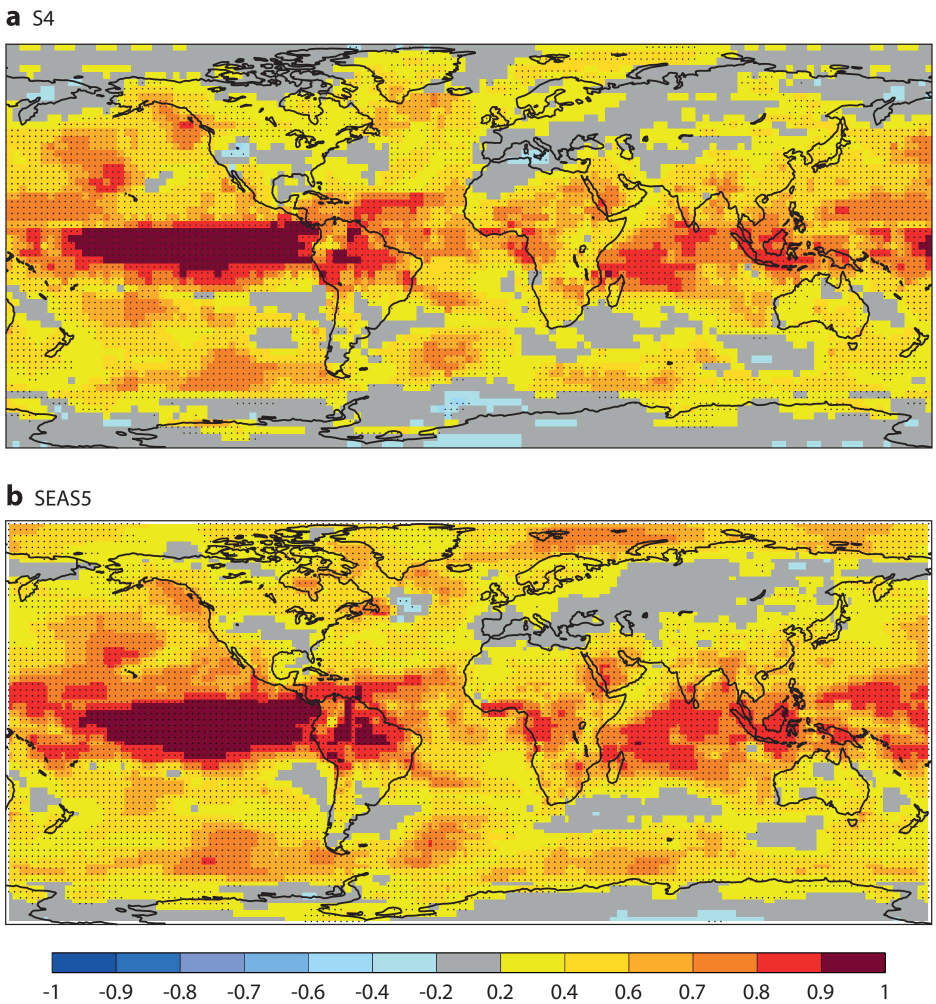

Figure 2 Anomaly correlation for ensemble mean December–January–February 2 m temperature predictions from 1 November for (a) S4 and (b) SEAS5. Measured skill in SEAS5 is higher partly due to the increased ensemble size, but beyond this there are real and statistically significant improvements in the tropics and in the Arctic. An anomaly correlation of 1 corresponds to a perfect deterministic forecast, while 0 means no skill.

SEAS5 and C3S

The combination of a major resolution increase and large increases in the number of re-forecast integrations has come at a very substantial computational cost. A significant part of the cost of SEAS5 production is met by C3S because SEAS5 is one of the core contributions to the new C3S multi-model seasonal forecasting service. The fact that C3S was willing to fund the computational costs of the re-forecasts, together with some human resources, enabled us to implement a substantially higher resolution system than we had otherwise planned for.

While ECMWF retains ownership and control of the full-resolution real-time forecasts, both the re-forecast dataset and a comprehensive 1 degree resolution dataset from the real-time forecasts will be publicly distributed by C3S. In its operational phase, the C3S release date will be the 10th day of each month, and any user anywhere in the world will be able to access both multi-model and individual SEAS5 plots on the C3S website, and download C3S multi-model and SEAS5 data for use in whatever product the user wants. The participation of SEAS5 in the open-access C3S multi-model system should bring major benefits, such as the increased use of our forecasts and enhanced feedback to us from the global community. C3S and other Copernicus services have the resources and community engagement to enhance and develop seasonal forecast products and applications way beyond what was possible as an ECMWF core activity.

SEAS5 performance

SEAS5 has brought consistent improvements in seasonal forecasts in the tropics while the picture in the extratropics is more mixed.

Tropical performance

SEAS5 tropical SST biases have substantially improved over S4, particularly in the equatorial Pacific. As shown in Figure 1, the bias in the NINO3.4 region has improved by nearly 2°C in the annual-range forecast, and the root-mean-square error (RMSE) has improved by approximately 0.1°C at forecast leads longer than two months. Accumulated improvements in physics since S4 have improved various aspects of the tropical mean climate, and the combination of the improved ocean model and improved atmospheric physics results in improved 2 m temperature prediction skill in the tropical Pacific, visible in Figure 2. Skill is slightly degraded in the tropical Atlantic, though.

Extratropical performance

In the extratropics, increased ocean horizontal resolution and improvements in ocean vertical mixing have improved some SST biases, for example in the North Pacific, while degrading others. In the North-West Atlantic, a region of decreased skill has appeared (visible as a small patch of blue in Figure 2). As shown in Figure 3, this is due to SEAS5 failing to capture the decadal variability of this region. In contrast, S4 was able to simulate the long-term oscillation. Initial investigations suggest this is due to the change from ORAS4 to ORAS5 initial conditions; investigations are ongoing.

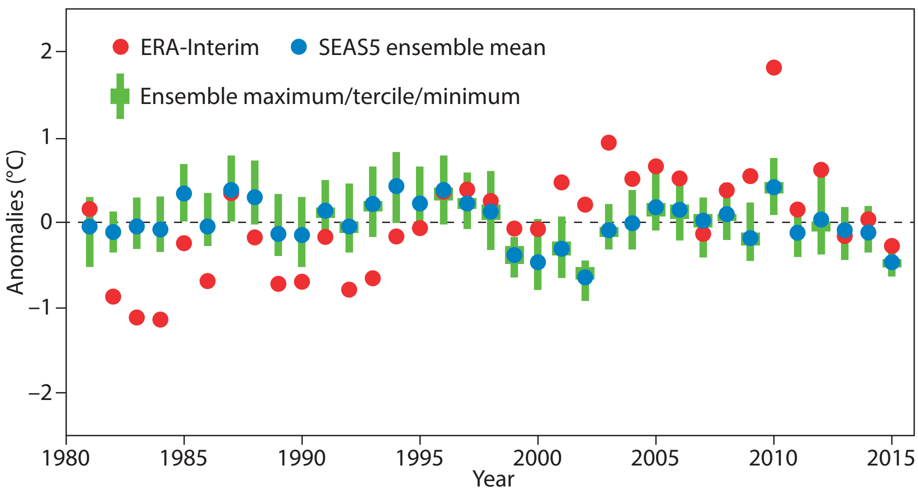

Figure 3 Time series of December–January–February SST anomalies in a small box in the North-West Atlantic from 1981 to 2015. Red dots are from the ERA-Interim reanalysis, while blue dots and green bars are SEAS5 forecasts from 1 November. The SEAS5 forecasts do not capture the shift from negative anomalies to positive anomalies in the late 1990s, compromising the skill in this region.

Improving the skill of seasonal forecasts over Europe is always challenging because average predictability is low, and scores are subject to considerable sampling uncertainty. Nonetheless, recent seasonal forecast systems have tended to show a fairly consistent picture of the pattern of grid-point skill over the European region, including marked seasonal variation. Figure 4 shows grid-point anomaly correlation skill for 2 m temperature forecasts for Europe for different seasons. Skill comes from capturing both interannual variations and long-term trends. This deterministic measure of skill can be complemented by probabilistic measures such as reliability, shown for summer and winter in Figure 5. This score is created by aggregating forecast performance over all grid points in Europe, so we lose spatial detail to gain a better probabilistic assessment. Such compromises are inevitable given the relatively small number of cases (36 years) that we verify. The reliability plot suggests that, although average skill is often low over Europe, particularly in winter, reliability is generally quite high. Note, though, that reliability aggregated across a region does not guarantee reliability at individual points. For full information, including all start dates, target seasons and multiple measures of skill, users should look at the extensive long-range verification charts provided online at: *www.ecmwf.int/en/forecasts/charts/catalogue/*

doi:10.21957/tsb6n1

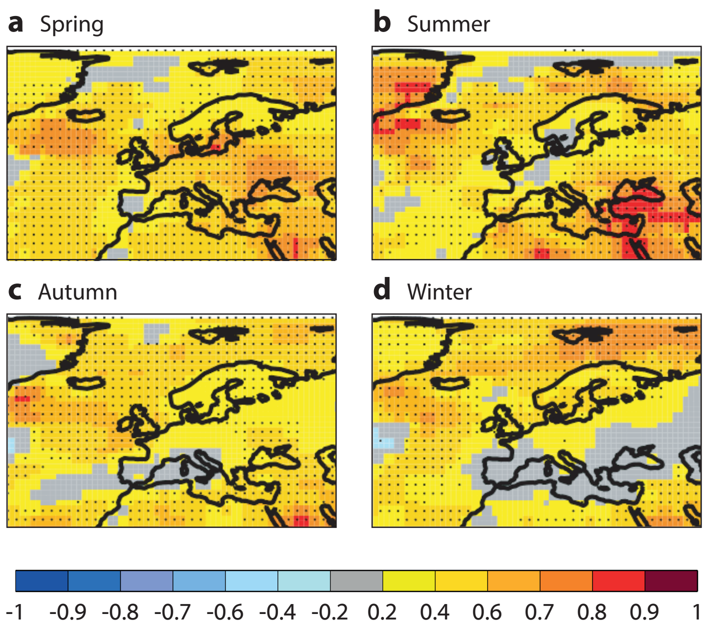

Figure 4 SEAS5 anomaly correlation skill for 2 m temperature in the European region, based on 1981–2016 re-forecasts, for (a) March–April–May, (b) June–July–August, (c) September–October-November, and (d) December–January–February, predicted from 1 February, May, August and November, respectively.

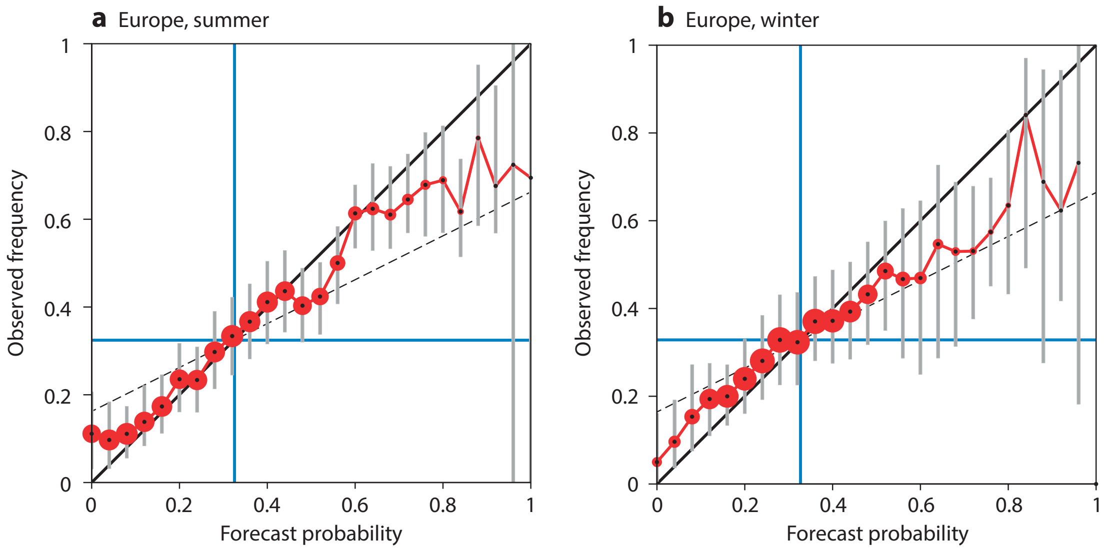

Figure 5 Reliability of forecasts of the probability that 2 m temperature anomalies will be in the upper tercile category for points in Europe (land and sea) for (a) 1 May forecasts for June–July–August and (b) 1 November forecasts for December–January–February. Both seasons have good overall reliability, as indicated by points lying close to the diagonal, but the June–July–August forecasts are sharper, i.e. more forecasts are far from the climatological value of 0.33. The distribution of forecast probabilities is indicated on the plot by the size of the circles, with bigger circles corresponding to more cases.

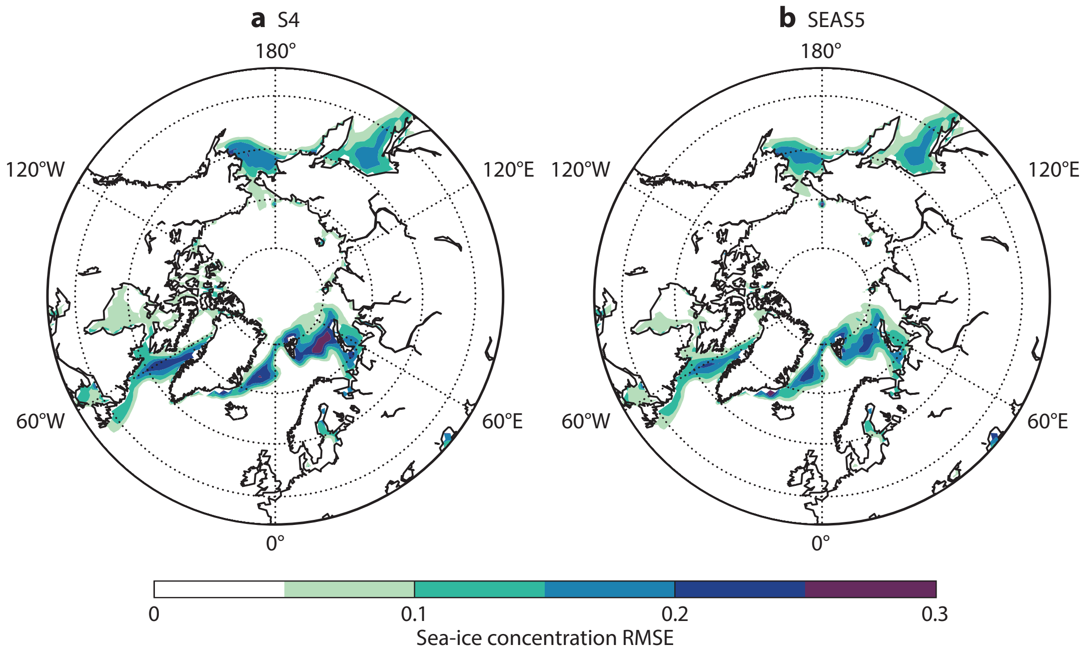

Figure 6 RMSE in predictions of December–January–February sea-ice concentration, for forecasts from 1 November, for the period 1981–2016 from (a) S4 and (b) SEAS5, showing the reduction in error due to the interactive sea-ice model.

Sea ice and stratosphere

The new prognostic sea ice gives SEAS5 the ability to predict sea-ice cover in the coming season. Although the forecasts are not completely unbiased, they result in more accurate sea-ice concentrations than S4, as shown in Figure 6. In the stratosphere, zonal wind and temperature biases have increased with respect to S4. Part of this is due to an unusual situation where higher horizontal resolution degrades the model in this area, for reasons that are still not fully understood. Stratosphere biases are currently being worked on at ECMWF.

Future strategy

SEAS5 was developed to be very close to the IFS configuration used for the monthly extension of the ENS. We plan to continue the path of scientific convergence, so that any improvements needed by SEAS are implemented on a continuous basis in the IFS, ready to be picked up when the next seasonal system is configured. The ENS and SEAS will be consistent in the sense that any desired changes important for the long range will already have been implemented in ENS. Each new IFS cycle will be tested for its long-range forecast performance, both in re-forecasts and a real-time configuration, so that once it is decided to update SEAS, this can be done easily.

This behind-the-scenes approach to making our forecast configurations more seamless leaves open the question as to whether more substantial changes might be made in how SEAS is run. Firstly, might SEAS be updated with every IFS cycle? This has some obvious drawbacks: a substantial increase in the cost of the re-forecasts; re-forecasts are only available a short time before the corresponding real-time forecast; a much more frequently changing system for users to work with. It also comes with a dilemma: if we insist that each new cycle at least maintains the long-range forecast performance, then the cost and difficulty of creating new IFS model cycles may substantially increase. If, on the other hand, we accept that the long-range forecast performance may fluctuate from cycle to cycle, then this could have a negative impact on seasonal forecast users.

doi:10.21957/tsb6n1

However, if frequent updates of the long-range forecasts were considered beneficial overall, for example because the IFS long-range performance was sufficiently stable for us to be confident that cycle updates would only ever have a small impact, then it might be feasible to better integrate the ENS and long-range ensembles. That is, it should be possible to design a single ensemble system where a range of resolutions and forecast lengths produce a truly seamless forecast system on timescales from days to seasons. Such a system would have some attractions beyond the purely aesthetic: cost savings might ensue from a well-designed ensemble (perhaps enough to compensate for the extra cost of more frequent long-range re-forecasts); long-range forecast products could be issued more frequently; and forecast products would be consistent across the different time ranges.

The future evolution of ENS and SEAS will in the end be determined by user requirements and priorities. Our scientific goal is to build a forecast model that fully exploits all sources of long-range predictability, and the necessary data assimilation systems to initialise the relevant parts of the Earth system. SEAS5 is but one step on the journey; there is still much ahead to look forward to.

For more information on SEAS5 and access to the SEAS5 User Guide, visit: *www.ecmwf.int/en/forecasts/documentation-and-support/long-range*.

Further reading

Tietsche, **S.**, **M.A. Balmaseda**, **H. Zuo** & **K. Mogensen**, 2015: Arctic sea ice in the global eddy-permitting ocean reanalysis ORAP5. *Climate Dynamics*, **49**, 775–789, doi:10.1007/s00382-015-2673-3.

Zuo, **H.**, **M.A. Balmaseda** & **K. Mogensen**, 2015: The new eddy-permitting ORAP5 ocean reanalysis: description, evaluation and uncertainties in climate signals. *Climate Dynamics*, **49**, 791–811, doi:10.1007/s00382-015-2675-1.

Zuo, **H.**, **M.A. Balmaseda**, **E. de Boisseson**, **S. Hirahara**, **M. Chrust** & **P. De Rosnay**, 2017: A generic ensemble generation scheme for data assimilation and ocean analysis. *ECMWF Technical Memorandum* **No. 795**.

European Centre for Medium-Range Weather Forecasts, Shinfield Park, Reading, RG2 9AX, England