



SEAS5 user guide

Version 1.2
(Description of updated tropical storm plots)

Date: Mar. 2021

Contents

1. Introduction to seasonal forecasting	4
1.1 The basis of seasonal forecasting	4
1.2 Methodology	5
1.3 How useful are today's seasonal forecasts?	6
2. The ECMWF Seasonal forecast system	6
2.1 Models	6
2.1.1 Ocean and sea ice model	6
2.1.2 Atmospheric model	7
2.1.3 Coupling	8
2.2 Initial conditions, re-forecasts and forecasts	8
2.2.1 Forecast system	8
2.2.2 Ocean analysis system	9
2.2.3 Atmosphere and land initial conditions	10
2.3 Post-processing and product generation system	11
2.3.1 Bias correction methodology	11
2.3.2 Interpreting anomalies: reference periods and long term variability	12
2.4 Known issues	14
2.4.1 Issues which can affect quality of forecast	14
2.4.2 Technical issues which do not affect forecast quality	15
3. ECMWF Seasonal forecast graphical products	15
3.1 Niño plumes	15
3.2 Forecast maps	17
3.2.1 Terciles	18
3.2.2 Ensemble mean	18
3.2.3 Probability of exceeding median	19

3.2.4 Probability of highest/lowest 20%	19
3.3 Forecast maps: sampling errors	19
3.3.1 Type I errors	20
3.3.2 Type II errors	21
3.4 Climagrams	24
3.4.1 Equatorial Southern oscillation	24
3.4.2 Northern hemisphere winter teleconnections	24
3.4.3 Area averages of 2 metre temperature and rainfall	26
3.4.4 Monsoon indices	26
3.5 Tropical storm forecasts	28
3.5.1 Tropical storm numbers	28
3.5.2 Track density	30
3.5.3 Verification	30
3.6 Annual range forecasts	31
4. Seasonal forecast data products	31
4.1 Data streams and MARS retrievals	31
4.1.1 Data streams in MARS	31
4.1.2 MARS retrievals	32
4.2 Archived atmosphere forecast data	35
4.2.1 Upper air fields	35
4.2.2 Surface fields	36
4.2.3 Model level fields	39
4.3 Archived wave forecast data	40
4.4 Ocean forecast data	41
5. Product interpretation	41
6. Operational history	42
7. References	43

1. Introduction to seasonal forecasting

1.1 The basis of seasonal forecasting

The ECMWF forecasts are created by using computational models to calculate the evolution of the atmosphere, ocean and land surface starting from an initial state based on observations of the Earth system. Due to limitations in the observing system, the initial state is not known perfectly. The evolution of the atmosphere is very sensitive to small errors in the initial conditions (it is a chaotic system) which limits the ability to forecast daily weather variations beyond 10 to 15 days in the future even if weather forecast models were perfect. Longer term predictions of the climate in the weeks, months and years ahead are possible due to a number of components in the Earth system, such as the ocean and cryosphere, which evolve more slowly than the atmosphere. Due to their slower evolution, they retain information from their initial state for longer and their evolution can be predicted on longer timescales. These Earth system components do not constrain the atmosphere enough to allow an accurate forecast of daily weather variations, but their influence can be detected in the average weather over a month or season. For example, monthly average rainfall in March to May in the Nordeste region of Brazil is related to sea surface temperatures (SSTs) in the tropical Pacific and Atlantic in the months before and during the rainy season. The seasonal forecast is a statistical summary of the daily weather calculated by the forecast model in the months ahead.

The coupled ocean-atmosphere system has several recurring phenomena that can be predicted on seasonal timescales. The most important of these is the El Niño Southern Oscillation (ENSO): the coherent, large-scale fluctuation of ocean temperatures, atmospheric circulation, air pressure and rainfall, across the tropical Pacific. Its influence is far reaching, with associated changes in sea-surface temperatures often across the whole width of the Pacific and the changes in tropical rainfall and winds spanning a distance of more than one-half the circumference of the earth. El Niño episodes (also called Pacific warm episodes) and La Niña episodes (also called Pacific cold episodes) represent opposite extremes of the ENSO cycle. The ENSO cycle is the largest known source of year-to-year climate variability.

Unusually warm or cold SSTs in the tropical Atlantic or Indian Oceans can cause major shifts in seasonal climate in nearby continents. For example, the sea surface temperature in the western Indian Ocean has a strong effect on the precipitation in tropical eastern Africa, and ocean conditions in the tropical Atlantic affect rainfall in northeast Brazil. In addition to the tropical oceans, other factors that may influence seasonal climate are snow cover and soil wetness, as

well as the winds in the stratosphere. All these factors affecting the atmospheric circulation constitute the basis of long-term predictions. Even so, relatively low predictability on seasonal time-scales is a feature of much of the globe, especially in the mid-latitudes and for smaller spatial scales (several hundred km, rather than several thousand).

1.2 Methodology

The ECMWF forecasts are created using numerical models. These models solve a complex set of hydrodynamic equations that describe the evolution of the atmosphere and ocean and include a set of parameterizations, which approximate how some processes such as convective precipitation affect this evolution. A description of the ECMWF forecast model is available in the [User Guide to ECMWF Forecast Products](#) and detailed information about the version of the model used in SEAS5 is in Section 2 of this document.

Forecasts are started from estimates of the initial state derived from observations of the Earth system. To estimate the impact of the error in these estimates on the forecast, many forecasts are initialized, each with slightly different, but plausible initial conditions, and each represents one way the atmosphere could evolve. In addition, because the forecast model is not perfect, the evolution of the forecast itself contains error. To represent the uncertainty caused by these errors, each individual forecast is perturbed by random numbers representing possible model errors. The combination of these individual forecasts is called an ensemble. Using ensembles, probabilistic statements can be made about the most likely atmospheric state, or weather, at a given point in time. An ensemble also provides information on the confidence in the forecast: the differences in forecast outcome generated by different ensemble members can be used as a measure of the precision of the forecast.

While the effect of random error can be assessed with ensembles, systematic error causes seasonal forecast models to develop characteristic errors that affects all forecasts in the ensemble as the forecast evolves. This means the average atmosphere and ocean state produced by the model differs from the observed average climate, which affects the accuracy of the forecast. These systematic errors are known as “bias” and can vary depending on season, region and forecast lead time. Since the magnitude of the bias can be comparable to year-to-year variation in seasonal average weather, this bias needs to be considered when interpreting the forecast. Bias is usually estimated by creating a set of seasonal forecasts for past years, called “re-forecasts” or “hindcasts” that can be compared to the historical record. These re-forecasts are created with a version of the forecast system as close as possible to that used for real-time forecasts, to ensure they provide a good estimate of the bias to be expected in the real-time forecasts. The difference between average seasonal climate in the model and observations can be accounted for when producing seasonal forecast charts.

1.3 How useful are today's seasonal forecasts?

As well as assessing the bias, re-forecasts are used to assess the skill and reliability of the seasonal forecasting system. Skill generally measures the ability of the re-forecast to reproduce past year-to-year variability, and reliability measures whether the forecast probability of events (such as a cold winter in Europe) matches the observed occurrence of those events. These assessments are [displayed together with ECMWF forecast products](#) and are described in Section 3. The quality of observations used for initialisation is key for the quality of the forecasts and re-forecasts. As a consequence, re-forecasts for periods before the availability of satellite observations - the 1980s - are not generally used to assess operational seasonal forecast systems and most state-of-the-art seasonal forecast systems create 20 to 40 years of re-forecasts. This limits the number of seasons available to assess the forecast system and thus the robustness of the estimates of forecast quality.

The benefits of seasonal forecasting are most easily established in forecasts for some areas of the tropics. This is because, in current state-of-the-art forecast systems, many tropical areas have more predictable signal than the mid-latitudes. In some parts of the world, and in some circumstances, it may be possible to give a relatively narrow range within which weather values are expected to occur. Such a forecast can easily be understood and acted upon; some of the forecasts associated with strong El Niño or La Niña events fall into this category. More typically, the probable ranges of the weather differ only slightly from year-to-year. Forecasts of these modest shifts might be useful for some but not all users. A more detailed comment on product interpretation is available in Section 5.

2. The ECMWF Seasonal forecast system

The system consists of an ocean analysis to estimate the initial state of the ocean, a global coupled ocean-atmosphere general circulation model to calculate the evolution of the ocean and atmosphere, and a post-processing suite to create forecast products from the raw numerical output. Detailed descriptions of the models, the analysis and the post-processing are given below.

2.1 Models

2.1.1 Ocean and sea ice model

As in System 4 (referred to hereafter as SEAS4), SEAS5 uses the NEMO (Nucleus for European Modelling of the Ocean) ocean model (Madec 2016), but with changes to the model version,

ocean physics and resolution. The resolution has been increased from a 1° ORCA grid to 0.25° ORCA grid. The increase in horizontal resolution improves the representation of sharp fronts and ocean transports. The number of vertical levels has increased from 42 layers to 75 levels. Vertical resolution is particularly refined in the uppermost part of the ocean; the number of levels in the uppermost 50 metres increases from 5 to 18. The depth of the surface layer of the ocean model has decreased from 10-metres to 1-metre, which improves the representation of the diurnal cycle of sea-surface temperatures (SST).

The ocean model configuration is based on that developed by the DRAKKAR international research network for the NEMO version V3.4 (SEAS4 used V3.3), and contains upgrades regarding aspects of ocean-surface wave interaction (Breivik et al. 2015) originally introduced at ECMWF. These aspects include a momentum flux estimated from the dissipation term (accounting for the intensity of breaking waves); the surface boundary condition of the turbulent kinetic energy equation, which now account for the energy flux from breaking waves (Craig and Banner 1994); and the Coriolis-Stokes forcing term is introduced in the momentum equation.

The Louvain-la-Neuve sea Ice Model (LIM2, Fichefet and Maqueda 1997), originally developed at the Belgian Université Catholique de Louvain, is introduced in SEAS5. LIM2 is part of the NEMO modelling framework. This prognostic sea-ice model allows sea-ice cover to respond to changes in the atmosphere and ocean states, enabling SEAS5 to provide seasonal outlooks of sea-ice cover.

2.1.2 Atmospheric model

The atmospheric component of SEAS5 is the [ECMWF IFS](#) (Integrated Forecast System) version [43r1](#). This model version was introduced for medium-range forecasting on 22 November 2016. The SEAS5 configuration of the IFS is based on the configuration used for the 43r1 ENS extended forecasts, including the horizontal and vertical resolution. The spectral horizontal resolution used for the main dynamical part of the model calculations is T319, but all of the model physical parameterization (including clouds, rain and the land surface) are calculated in physical space on a reduced O320 gaussian grid, which has grid spacing of approximately 36 km. There are 91 levels in the vertical, with a model top in the mesosphere at 0.01 hPa, or around 80 km. The atmospheric model uses a two-time-level semi-Lagrangian scheme for its dynamics, and has a 20 minute time-step.

The SEAS5 configuration has a few settings and forcings which are different from the ENS extended settings to improve the representation of processes that affect seasonal skill. The tropical amplitude of the non-orographic gravity wave drag was considerably reduced compared to the default settings in 43r1 in order to improve the modelling of the Quasi-Biennial Oscillation and the climate mean stratospheric winds. This setting has been adopted for future cycles of the

IFS. Tropospheric sulphate aerosol follows the decadal varying CMIP5 climatology, rather than the time invariant climatology that is default in 43r1. Volcanic stratospheric sulphate aerosol continues to be treated as in SEAS4; the initial load of volcanic aerosol is prescribed using [GISS data](#) (2012 update), the horizontal distribution is approximated by three numbers (northern hemisphere, tropical and southern hemisphere amounts) and the vertical distribution follows a prescribed profile that is applied globally. The forecast is initialized using the GISS values from the month before the forecast starts, and then evolved in time with damped persistence. SEAS5 cannot predict volcanic eruptions, but after a major eruption occurs, manual estimates of the volcanic aerosol, based in part on Copernicus Atmosphere Monitoring Service (CAMS) SO₂ analyses, will be included in future real-time forecasts. It would be preferable to have a better characterization of volcanic aerosol distribution and properties, and eventually real-time analysis systems should be able to provide such information. Prognostic ozone is calculated using the a new scheme (Monge-Sanz et al. 2011), but unlike in SEAS4, ozone is not radiatively interactive. Instead, the radiation scheme sees the same ozone climatology as used in the 43r1 ENS extended forecasts.

2.1.3 Coupling

A gaussian method is used for interpolation in both directions, primarily due to the complexity of the 0.25° ORCA grid. The gaussian method automatically accounts for the inevitably different coastlines of the atmosphere and ocean models - values at land points are never used in the coupling, since these can be physically very different to conditions over water. The coupling interval is 1 hour, which allows resolution of the diurnal cycle.

2.2 Initial conditions, re-forecasts and forecasts

2.2.1 Forecast system

The seasonal forecast consists of a 51-member ensemble. The ensemble is created using a combination of SST and atmospheric initial condition perturbations (described in Sections 2.2.2 and 2.2.3) and the activation of stochastic physics. The stochastic physics settings are identical to those used in the medium-range ENS and use both the SPPT3 scheme and stochastic backscatter.

A seasonal forecast is produced each month. The forecasts have an initial date of the 1st of each month, and run for 7 months. Forecast data and products are released at 12Z UTC on a specific day of the month. For SEAS5, this is the 5th. While SEAS5 is expected to be operational from 1 Nov 2017, forecasts for 1 Jan 2017 through to 1 Oct 2017 have also been completed and archived for reference. In addition, on 1st February, 1st May, 1st August and 1st November, 15 of

the 51 forecast members are extended a further 6 months for a total forecast length of 13 months. These annual range forecasts are designed primarily to give an outlook for ENSO.

Every seasonal forecast model suffers from bias: the forecast mean climate produced by the IFS differs from the observed mean climate. Since interannual variation in forecasted seasonal climate can be comparable to the bias, this bias needs to be considered to interpret the forecast. We estimate the bias from a set of retrospective seasonal forecasts for past years that can be compared to the historical record. As well as quantifying the bias, these “re-forecasts” are used to assess the skill of the seasonal forecasting system and thus inform decisions based on the forecast.

The set of re-forecasts (also sometimes known as hindcasts or back integrations) start on the 1st of every month for the years 1981-2016 and have 25 ensemble members. The data from these re-forecasts is available to users of the real-time forecast data to calibrate real-time forecast products. To calibrate the annual range forecasts, on 1st February, 1st May, 1st August and 1st November, 15 of these 25 ensemble members are extended a further 6 months.

For start dates on the 1st February, May, August and November, the re-forecast ensemble size will in due course be extended to 51 members, to allow a better assessment of skill of the system. These 26 additional ensemble members are not considered part of the operational system. They will be archived in MARS together with the first 25 members, and are available for use in studying SEAS5 performance. Similar extended re-forecast sets were made for System 3 and SEAS4.

2.2.2 Ocean analysis system

SEAS5 ocean and sea-ice initial conditions for forecasts and re-forecasts are provided by the new [operational ocean analysis system](#) (OCEAN5), made up of the historical ocean reanalysis (ORAS5) and the daily real time ocean analysis (ORTA5). OCEAN5 uses the same ocean and sea-ice model as the coupled forecasts in SEAS5. Compared to its predecessor ORAS4 (Balmaseda, Mogensen, and Weaver 2013), OCEAN5 has higher resolution, updated data assimilation and observational data sets and, most importantly, provides sea-ice initial conditions.

ORAS5 is based on Ocean Reanalysis Pilot 5 (ORAP5) (Tietsche et al. 2017; Zuo, Balmaseda, and Mogensen 2017), but using updated observational data sets. The ocean in-situ temperature and salinity comes from the recent quality-controlled EN4 (Good, Martin, and Rayner 2013), which has higher vertical resolution and fuller spatial coverage than the previous version EN3. The altimeter sea-level data have also been updated to the latest version (DUACS2014) from CMEMS (Copernicus Marine Environmental Monitoring Services). The underlying SST analysis before 2008 now comes from the HadISST2 dataset, the same used in the ERA5 reanalysis. The sea-ice concentration before 1985 comes from ERA-40 and from 1985 to 2008 it comes from an

OSTIA reprocessed product. From 2008 onwards the SST and sea-ice are given by the OSTIA product delivered in real-time, which is also used in the ECMWF operational analysis.

In order to sample the uncertainty in our knowledge of the ocean state, ORAS5 contains a 5-member ensemble analysis. The perturbation scheme used to generate the ensemble of reanalyses consists of two distinct elements: perturbations to the assimilated observations, both at the surface and at depth, and perturbations to the surface forcing fields (Zuo et al. 2017). Prior to starting the coupled model forecasts, the ocean analysis temperatures are further perturbed so that all ensemble members of forecasts and re-forecast have different ocean initial conditions. Using the ORAS5 HadISST2 pentad analysis error repository (Zuo et al. 2017, Section 4), perturbations are applied to the upper 22 levels of the sea temperature and reduce with depth.

2.2.3 Atmosphere and land initial conditions

For the re-forecast period (1981 to 2016), the atmospheric initial conditions come from ERA-Interim (ERA-I, Dee et al. 2011) and for the forecasts (1st January 2017 and later) they come from ECMWF operational analyses. SEAS5 includes prognostic ozone, and requires ozone initial conditions. The interannual variability of ozone in ERA-I is affected by changes in satellite instruments over time, and does not represent the true interannual variability of ozone in the atmosphere (Dee et al. 2011). Since these spurious changes have been found to drive substantial temperature errors in the stratosphere, they cannot be used as initial conditions. In SEAS5, a seasonally varying climatology is derived from the ozone model within a run of the 43r1 IFS nudged to ERA-I vorticity (12h timescale) and tropopause temperature (5 day timescale, needed to control biases in lower stratosphere temperature). The use of climatological initial conditions precludes any initial data on ozone anomalies, but the ozone field evolves during the forecast to values consistent with the predicted stratospheric state.

Initial conditions for the land surface and lakes are created differently than the atmosphere initial conditions. For the re-forecast period, the HTESSEL land surface model used in Cy43r1 is run in offline mode, with forcing data (precipitation, solar radiation, near surface temperature, winds and humidity) coming from ERA-I. For the real-time forecasts, ERA-I is not available. Thus from 1 Jan 2017 onwards, the land surface initial conditions are taken from the ECMWF operational analyses.

The real-time analyses must be interpolated from O1280 to the O320 grid. This interpolation can result in locally large differences compared to initial conditions prepared directly at the lower resolution. Future changes in the operational analysis may also introduce further incompatibilities in the land initial conditions. Consequently, a limiter is used to prevent the real-time land surface values taking unrealistic values relative to those used in the re-forecasts, which might otherwise occur in mountainous regions and/or poorly observed areas. The limit fields define the maximum

and minimum permitted values of the field at each grid-point in the initial conditions of the real-time forecast. Limit fields are defined for each surface variable and for each calendar month. The limits are defined as the maximum and minimum values observed at that point and calendar date for the 36 year re-forecast period, plus an additional margin specified as a global constant for each field. This margin is generally chosen as 0.75σ (where σ is the globally averaged standard deviation of the field), which would correspond to a 500-year return period value. By capping the real-time initial conditions at this level, we ensure that any unphysical initial values are replaced with values which, while still extreme, are still feasible. The limiter thus acts as a safety limit, and has little or no impact on correctly estimated anomalies. Equally, because it acts so rarely, some level of inconsistency is still possible in the real-time initial conditions. For snow depth and lake variables, the margins are specified differently – for snow, we take a margin of 2 cm water equivalent beyond the previous range, thus for areas where snow was not observed on a particular date in the previous 36 years, the depth is limited to 2 cm water (up to 20 cm snow). For lakes, where there are known incompatibilities between the offline and real-time analyses, the margin is set to zero and the real-time initial conditions are simply limited to the previously observed min/max values.

Ensemble member 0 is initialized from unperturbed atmospheric initial conditions. Initial conditions for all other ensemble members have perturbations applied to some fields to represent uncertainty in the initial atmosphere state. The perturbed fields include all upper air fields and a limited set of soil moisture, soil temperature, snow, sea-ice temperature and skin temperature fields. As in operational ENS, perturbations from an ensemble of data assimilations (EDA) and perturbations constructed from the leading singular vectors are applied. A detailed description of the initial condition perturbation methodology is available in Part V of the [43r1 documentation](#). EDA perturbations are not available for the earlier years in the re-forecast set, so to preserve consistency across the hindcast set, the EDA perturbations from 2015 were applied to all re-forecast years.

2.3 Post-processing and product generation system

2.3.1 Bias correction methodology

Seasonal mean climate anomalies are typically of a similar magnitude to model biases, so some form of post-processing to remove model bias is needed. SEAS5 forecast products are generally corrected for mean biases in the forecast system, but no other corrections are applied. For example, the spatial plots of ensemble mean forecasts are not normalized to match observed variance, and probability forecasts are not calibrated according to past forecast skill.

ECMWF SEAS5 products available in graphical format are calibrated using a 25-member ensemble hindcast set over 1993-2016, to align with the Copernicus C3S calibration period. Two slightly different methods are used for bias correction, although in both cases the a posteriori correction is based on the assumption of a quasi-linear behaviour of the atmosphere and ocean anomalies. For Niño indices (e.g. Niño 3), the mean bias of the model relative to observations is estimated as a function of lead-time and calendar month from the difference between model re-forecasts values and verifying analyses. At present we use the NCEP Olv2 SST analysis, and estimate the model bias using re-forecasts from 1993-2016. This bias is then used to correct the model output and produce a forecast of absolute values of SST. To issue a forecast anomaly, this absolute value is then referenced against a specified observed climatology (1981-2010). Note that the years of the reference climatology and the years used to estimate the model bias do not have to be the same, and that this approach requires a high quality observational analysis.

For all other predicted variables, biases are removed from consideration by considering only model anomalies with respect to the model mean state. Specifically, the values of the forecast ensemble are compared to the values of a climate reference ensemble (made up of model re-forecasts with the same lead time and calendar month, and covering 1993-2016), and the differences between model forecast and model climate are assessed and plotted. Probabilities are calculated by first using the 600 member re-forecast ensemble to define the relevant percentile boundaries (e.g. the model climate median is defined as the average of the 300th and 301st rank-ordered values), and then counting the fraction of the real-time forecast ensemble members that exceed the percentile (so 37/51 members would be taken to indicate a forecast probability of 0.725).

Whichever bias correction methodology is used, there will be some inaccuracy in the estimation of model bias or model and observed climate due to the limited sample of past cases. In the case of the Niño indices, the length of the calibration period and the average accuracy of the forecasts (r.m.s. errors at several months are of order 0.4 °C) result in an uncertainty in the bias correction of just below 0.1 °C. Thus bias uncertainty is a small contributor to the overall uncertainty in a forecast, as long as the bias in the real-time forecasts is consistent with the re-forecasts.

2.3.2 Interpreting anomalies: reference periods and long term variability

The advantage of using anomalies with respect to model climate for graphical products is that they are independent of observational datasets. The disadvantage is that the climate base period for anomalies cannot be chosen independently of the re-forecast period. For example, a model climatology based on forecasts in the 1993-2016 period would differ from a climatology based on forecasts in the 1981-2010 period, if there have been real low frequency changes in the climate system between these periods that the model captures. The observed difference between these

periods is likely to be a combination of systematic changes (which we hope the model can capture) and an element of chance (which the model will not reproduce). The tricky issue is to relate differences in model climate between different periods to differences in observed climate.

Suppose we know the model climate for the 1993-2016 period, based on a total of 600 integrations (a 25-member ensemble for 24 years). Suppose we also know that for the observed variable of interest, the 1993-2016 period was, for example, 0.2°C warmer than the 1981-2010 period. If the forecast for the coming season shows a slight cooling of 0.1°C in the ensemble mean compared to the model climatology for 1993-2016, how do we interpret this if we are asked to produce a forecast relative to the 1981-2010 period? Do we allow for the 0.2 difference in observations between the two periods, and predict a 0.1°C warming, or do we simply insist that the model gives a 0.1°C cooling? (Of course the spread of the ensemble will in any case be bigger than this difference, but the choice will certainly affect the probability distribution).

For fields which are close to being deterministic, i.e. whose value consists of a large seasonally predictable signal and a small amount of unpredictable noise, then it is reasonable to suppose that the difference in the observed climate between the periods is due to a real change in the system, and is not just an artefact of sampling. If we further assume that the model is capable of reproducing the observed low-frequency variability (which is not certain), then we would expect the model climate for 1993-2016 to be shifted relative to the model climate for 1981-2010 by the same amount as the observations, and so we can apply the correction of 0.2°C to adjust the base period of the forecast.

However, in other cases the level of noise in seasonal means may be large enough that the difference in observed values between two periods is likely to be largely due to chance. That is, the difference between observed temperatures in 1993-2016 and 1981-2010 might be due to chance variations. In this case, it might be more appropriate to take the (600 member) model climate as the best estimator of what the model would have produced for the specified thirty-year period. In this case we might not make a correction to allow for the different base period.

For SEAS5, anomalies are plotted relative to 1993-2016 model climate, both for consistency with Copernicus C3S and with a view that anomalies relative to a more recent “past” are likely to be more relevant to most users. However, the re-forecasts are produced from 1981-2016, allowing users to explore the choice of different reference and calibration periods.

Temperature is a field where it is clear that there are substantial trends to warmer values over recent decades, and this is reproduced in the seasonal forecast system and needs to be accounted for in some way when considering different base periods. Nonetheless, the proper calibration of low frequency (decadal or longer) variability within the ECMWF system is not fully

understood. For noisy fields without strong trends, such as precipitation, adjusting for base period differences is likely to be inappropriate.

2.4 Known issues

We list here known issues with SEAS5. These are additional to the general limits of seasonal forecast systems and the limits inherent in the specific design of SEAS5, which are discussed in the rest of this document.

2.4.1 Issues which can affect quality of forecast

Inconsistent lake temperatures in Lake Superior and the Caspian Sea.

The initial conditions for lakes (including the Great Lakes and the Caspian Sea) for the re-forecasts are obtained from an offline run of the land surface/lake model. Due to limitations in the lake model, there are some biases in the resulting simulation, particularly visible in the surface of Lake Superior becoming too warm in summer due to excess stratification. For the operational analysis, from which the real-time forecasts start, a nudging scheme is used for the Great Lakes and the Caspian Sea to prevent such biases developing. Thus the real-time forecasts do not have a surface temperature bias in their initial state, and as a consequence the real-time forecasts of Lake Superior are cooler in the summer than the re-forecasts were. This manifests as a cold signal in the lake temperature in the initial conditions of the real-time forecasts, which is artificial and should be disregarded. The limiter (discussed above) prevents the problem being excessive, and tests have shown that the impact remains very local.

Although the nudging scheme in the operational analysis successfully controls the surface temperature of the lakes to which it is applied, in some situations it leads to unrealistic temperatures at the bottom of the lake. This is affecting both Lake Superior and the Southern Caspian Sea in the present (2017) operational analysis. This is an error which is present in the real-time forecasts, but not the re-forecasts, and can manifest as unrealistic warming in the lake surfaces during winter. The limiter is quite successful in controlling the deep lake temperatures (because the interannual variability is small), but some residual affect in the Caspian Sea in winter may remain. Because this is a problem with the operational analyses, we hope that it can be addressed during the lifetime of SEAS5.

Difficulty reproducing trends in Atlantic tropical storm numbers

Although SEAS5 represents the year-to-year variation of Atlantic tropical storm numbers as well as might be expected, in the last few years the trends in model predictions and observed storm numbers have started to move in opposite directions. Investigations have shown that this is due

to errors in the way the model is handling recent trends in some aspects of tropical climate, in particular wind shear. Similar difficulties in handling low frequency variability of Atlantic storms have been noted in earlier periods (e.g. from the 1980's to the 1990's). To improve the robustness of the tropical storm forecast products, it was decided to switch the calibration method from using a fixed calibration period (1993-2015) to a moving calibration period (a moving interval of the 10 previous years). This change was made in March 2021, to coincide with the start of the issuing for forecasts for the 2021 hurricane season. The moving interval method is used only for the northern hemisphere basins.

2.4.2 Technical issues which do not affect forecast quality

No issues have been identified.

3. ECMWF Seasonal forecast graphical products

There are two classes of product produced by the Seasonal Forecasting system at ECMWF. The first is a set of graphical products. These are designed to show the main features of the model predictions for the forthcoming seasons, in an easily understood way. They are usually accompanied by a set of verification products, to show the SEAS5 skill relevant to that product. We describe these graphical products and their verification below.

3.1 Niño plumes

Forecasts of Equatorial Pacific sea surface temperature anomalies averaged over Niño 1+2 (0°-10°S, 80°-90°W), Niño 3 (5°N-5°S, 90°-150°W), Niño 3.4 (5°N-5°S, 120°-170°W) and Niño 4 (5°N-5°S, 160°E- 150°W) areas are shown for seasonal and annual-range [charts](#). Predicted monthly mean anomalies from each individual ensemble member are shown as red dots joined by thin red curves, and the verifying analysis, where available, is represented by a thick dotted dark blue curve. Forecasts start on the 1st of a month, and the monthly mean anomaly for that month is the first value plotted. This is joined to the preceding (observed) monthly mean anomaly with a dashed line to represent the continuity of the forecast with the analysis. Note that the lines do not represent the continuous time evolution of the SST, they simply connect the monthly mean values.

The Niño plumes show a spread in predicted values - sometimes the spread is large, sometimes it is relatively small. The spread in the first month is largely controlled by the perturbations applied to the ocean initial conditions, in particular the SST perturbations. The growth of the spread in

later months is dominated by the inherent unpredictability in atmospheric behaviour within the coupled system. The spread is observed to depend on both the time of year and, to some extent, the state of ENSO.

The Niño indices are used as indicators of El Niño activity. The predicted anomalies are defined with respect to the NCEP OIv2 climatology adjusted to a base period of 1981-2010, the most recent standard climate reference period available. Note that El Niño forecasts made elsewhere may be with respect to other base periods.

Since the equatorial oceans have been warming in recent decades, the size of positive anomalies will be larger if an older base period is used. There is no universally agreed definition of "El Niño" and "La Niña", although one common approach is to use a threshold of 0.5 °C applied to the Niño 3.4 SST index, with anomalies of 0.5 to 1 °C as a weak ENSO event. ENSO events can differ substantially in their spatial structure. The four different SST indices provided here give a fair description of how the SST anomalies are distributed in an east-west direction along the equator.

Together with the Niño time-series plots, we show verification statistics based on past forecasts/re-forecasts. The root mean square error (r.m.s. error) plot shows the cross-validated r.m.s. error for the ensemble mean of forecasts made with the same calendar start date in previous years. Also shown is the error obtained by persisting the initial anomaly (black dot-dashed line), and the r.m.s. spread of the ensemble. Comparison of the size of the spread with the forecast error shows the extent to which the forecast plume tends to be over- or under-dispersive. SEAS5 has both reduced error and reduced spread compared to SEAS4, and tends to be under-dispersive, particularly in Niño 4.

The mean square error skill score (MSSS) relative to climatology shows the skill of the forecast in a range between 1 (a perfect deterministic forecast) and 0 (no better than climatology). It shows how much of the variation of observed SST is being correctly forecast, and gives a sense of the lead time over which the forecast retains useful skill. MSSS is related to anomaly correlation, but unlike correlation is sensitive to errors in amplitude. If the amplitude is correct, then anomaly correlation is simply the square root of the MSSS.

The Mean Absolute Error (MAE) time-series plot shows a time history of forecast errors for forecasts starting at the given calendar month. The MAE of a forecast for a given month at a given lead time is the absolute difference between the forecast ensemble mean and the verification; this is then averaged across the different forecast leads for a given start date to give the MAE for that start date. On this plot we also show what we call the Best Absolute Error (BAE), which for a given forecast date is the average across lead times of either zero (when the verification lies within the predicted range) or the absolute difference between the verification and

the outer limits of the predicted range. For a perfect forecasting system, the BAE will be zero or close to zero most of the time. The MAE and BAE time-series plots give a sense of how much variation there is in the errors, and may suggest whether or not errors have tended to decrease with time or have tended to be associated with certain phases of ENSO. This is designed to complement the other scores, which are averaged across the whole of the re-forecast period.

The amplitude ratio shows the ratio of the amplitude of the anomalies from the re-forecasts to the amplitude of the corresponding observed anomalies. The amplitude at a given month is measured as the standard deviation of the index for that month, calculated from the set of individual ensemble members from all the re-forecasts. If the model is indistinguishable from reality, then the model and observed anomalies should have the same value, and the ratio would be one. If the ratio is, for example, greater than one, then it means that the model forecasts have bigger anomalies than are observed. This should be taken account of in interpreting the plume plots – if at a certain time of year the model is known to overestimate the amplitude of ENSO variability, then any anomaly it is predicting should be scaled appropriately before it is taken to be a forecast of the real world. In SEAS4 such a variance scaling was applied to the plumes, but in SEAS5 it is not. The plots also show the amplitude ratio of an anomaly persistence forecast – this differs from one because observed ENSO amplitudes vary according to the time of year.

3.2 Forecast maps

Spatial maps are produced showing the model-predicted anomalies in seasonally averaged quantities. In most cases both global and regional plots are produced, although not all plots are publicly available. Each plot is labelled with the period for which it is valid, e.g. DJF 2011/12 is the three-month period December 2011 - February 2012. The start date of the forecast is given, as is the number of model integrations in the forecast ensemble and the number used to define the climate, together with years for which re-forecasts are produced.

For SEAS5 the forecast products are released on the 5th day of each month, so the "usable" lead times are slightly less than their nominal values. Plots for lead times of 1, 2, 3 and 4 months are produced each month, and can be selected using the errors and time-line below the maps. It is good practice to compare the forecast charts for a given target period at different lead times as they become available. The major forecast signals are usually fairly stable, but not always.

Weaker signals are subject to appreciable sampling error, and even if the model signal remains unchanged, plots from different months vary just because of the sampling. The colour scale depends on the field plotted: in most cases blue is used for lower values and red for higher values of a field or probability, but for precipitation brown is used for drier and green for wetter conditions. For individual tercile and outer quintile (20%ile) categories, high probabilities are in red regardless of the field or category being plotted.

3.2.1 Terciles

For each forecast parameter, lead time and calendar start date, there is a set of 600 re-forecasts (a 25 member ensemble for 24 years). For each grid point, the 600 re-forecasts are analysed to determine the terciles of the model climate distribution at the specified lead time. The lower tercile is the value below which the outcome occurs in 1 out of 3 cases in the model climate, and the upper tercile is the value which is exceeded in 1 out of 3 cases. In the absence of any other information, and assuming the climate to be stationary, we would take the probability of a future value exceeding the upper tercile to be 1/3. Using the forecast we can calculate the fraction of ensemble forecast members which exceed the upper tercile of the model climate distribution. If there is no particular "forcing" acting on the system, then the proportion of forecast members exceeding the upper tercile will be about 1/3, and indeed this is often the case. However, if there is something in the climate system that "pushes" the forecast in a particular direction, then the predicted probability can be very different from 1/3, and these situations are typically of particular interest.

Plots of the probabilities of the individual tercile categories (ie below the lower tercile, between the lower and upper tercile, and above the upper tercile) are produced, with contour intervals which show both where there is an unusually high chance of a particular category and also where there is an unusually low chance of a particular category occurring. We also produce a tercile summary plot, which shows in a single figure the areas which have an increased probability of being either below the lower tercile or above the upper tercile. This plot gives a good overview of a seasonal forecast, and is listed first in the choice of plots offered to the user on the website.

3.2.2 Ensemble mean

The ensemble mean anomaly represents the shift in the first moment of the predicted probability distribution - it is not intended as a deterministic forecast of the actual value. Tercile and other percentile category probability plots give information on what the model is predicting relative to the typical amplitude of variation of the quantity concerned - for example, the chances of it being "unusually" warm. The ensemble mean plots give information on what the model is predicting in absolute terms - °C or mm of rainfall.

The values at each point on each map are subjected to a significance test before plotting. The significance tests are made using the Wilcoxon rank-sum test, which is non-parametric and very efficient at detecting shifts in the mean of a distribution; results are generally very similar to a 't' test. Points where differences between the forecast and climate distributions are not significant at the 10% level are blanked out (for most fields), and appear white on the plot. This is quite a lax test, and allows both areas of modest signal strength and some areas without a signal to be shown. A second significance test is made at the more stringent 1% level, and a solid contour line

encloses areas which are significant at this level. The significance test does not inform us about what confidence we should have in the eventual outcome, or say anything about skill or reliability of past forecasts. It simply informs the user about the likelihood of an apparent signal being due to sampling errors in the forecast ensemble, in the case when no signal is present. For more on sampling errors, see the sub-section on sampling below.

3.2.3 Probability of exceeding median

Probability maps show the probability of a given model variable (e.g. precipitation) being greater than the model climate median. As with the terciles, the climate median is estimated from the set of 600 re-forecasts made for the same calendar start date and lead time during the 24 year period 1993-2016. The probabilities are shaded symmetrically above 60% and below 40%. The probability plots do have not a significance masking applied, but as for the ensemble mean plots the 1% significance level contour is shown for guidance.

3.2.4 Probability of highest/lowest 20%

We also show probability maps for exceedance of the upper and lower 20th percentiles. These are useful for providing information in regions where the distribution of likely outcomes is shifted substantially from the climatological average. The probabilities here are calculated in the normal way, by counting the number of forecast members in the relevant interval of the climatological distribution. Detailed statistical examination of the tails of the forecast would require different analysis techniques. Although this could be done, verification of how well the forecast system predicts low probability events would be a challenge, given the very limited samples available.

3.3 Forecast maps: sampling errors

Information on the likely impact of sampling errors can be given in different ways. One traditional way is to test a null hypothesis that the forecast distribution is the same as the climatological distribution, using a significance test which is efficient at detecting any shift in the forecast distribution. Such a test can be helpful for screening out situations where an apparent forecast 'signal' is due to a chance fluctuation in the sampling. The results of such significance testing are shown on the ensemble mean and probability of exceeding median charts.

However, this sort of significance testing is very limited, even in telling us about sampling errors. It is not directly relevant to e.g. tercile probabilities, where the user is interested in the sampling accuracy of the probability of a particular event. And although it can warn us about the potential presence of "false positive" signals, it tells us nothing about "false negatives".

3.3.1 Type I errors

First we consider the possibility of a chart showing a signal when in fact none is present. In statistical terms this is a Type I error. If we assume the forecast and climatology distributions are identical, then we can calculate the probabilities of the calculated “forecast probability” of an event falling within a certain range, allowing for the sampling errors in both the model forecast and model climatology distributions. Such calculations are made using the binomial distribution, and the bigger our ensemble sizes, the lower the chance of obtaining the 'wrong' forecast category. The following tables give the probabilities of various signals being plotted for SEAS5 (S5) on the tercile, quintile and median plots on the web, when no forecast signal is present, i.e. when the forecast and climatological probability distribution functions are in fact the same. Equivalent sampling probabilities for previous systems (S4, S3 and S2) are also given. Sampling errors depend on both real-time ensemble size and the size of the re-forecast climatology (51 and 600 for SEAS5).

Tercile plots: Probability of different plotted ranges when model signal = 33.3%					
Plotted range	S5	S4	S3	S2	Colour
0-10%	0.0001	0.0002	0.0010	0.0030	
10-20%	0.03	0.03	0.05	0.07	
20-40%	0.81	0.80	0.76	0.70	BLANK
40-50%	0.15	0.16	0.18	0.19	
50-60%	0.01	0.01	0.02	0.03	
60-70%	0.0001	0.0002	0.0010	0.0030	
70-100%	0.0000	0.0000	0.0000	0.0000	
20%ile plots: Probability of different plotted ranges when model signal = 20%					
Plotted range	S5	S4	S3	S2	Colour
0-10%	0.05	0.05	0.07	0.10	
10-30%	0.91	0.90	0.86	0.79	BLANK
30-40%	0.04	0.05	0.07	0.10	
40-50%	0.001	0.001	0.003	0.01	

50-70%	0.0000	0.0000	0.0001	0.0004	
70-100%	0.0000	0.0000	0.0000	0.0000	
Median plots: Probability of different ranges when model signal = 50%					
Plotted range	S5	S4	S3	S2	Colour
10-20%	0.0000	0.0000	0.0001	0.001	
20-30%	0.003	0.004	0.009	0.02	
30-40%	0.09	0.09	0.11	0.15	
40-60%	0.82	0.82	0.77	0.69	BLANK
60-70%	0.09	0.09	0.11	0.12	
70-80%	0.003	0.004	0.009	0.02	
80-90%	0.0000	0.0000	0.0001	0.001	

From this table, we see that if the model forecast is equal to climatology, there is a high probability that the map will correctly show no signal (81% in the case of the tercile plots, 91% in the case of the 20%ile plots, 82% in the case of the median plots). There is a small but not negligible chance of a weak signal being shown (e.g. one colour band either side of the blanked area). Strong signals are very unlikely to occur by chance. The table also shows how the sampling properties have systematically improved with successive forecast systems. Remember that the probabilities in this table apply locally. If we look at a plot for a hypothetical case in which no signal is present, we would expect to see a moderate amount of colour overall, even if the a priori probability of it occurring at any given location is fairly small. If many degrees of freedom are present in the plot, even locally improbable events are likely to occur somewhere

3.3.2 Type II errors

The risk of a chart falsely showing a signal to be present is not the only concern. We also face the situation in which a signal is present, but the chart does not show it. In statistical terms this is a Type II error. We can calculate the probabilities of these errors in the same way as we handled the Type I errors, again allowing for sampling error in both the forecast and the climate ensembles. This time we must specify an assumed true level of signal in order to calculate the effect of sampling errors upon it. Some example tables for SEAS5 probabilities are given below:

Tercile plots	Model signal: 5%		Model signal: 15%		Model signal: 45%		Model signal: 55%		Model signal: 65%		Model signal: 85%	
Plotted range	Prob		Prob		Prob		Prob		Prob		Prob	
0-10%	0.93	CORRECT	0.20		0.0000		-		-		-	
10-20%	0.07		0.65	CORRECT	0.0002		0.0000		0.0000		-	
20-40%	0.0003	NULL	0.16	NULL	0.26	NULL	0.02	NULL	0.0004	NULL	-	NULL
40-50%	0.0000		0.0000		0.50	CORRECT	0.23		0.02		0.0000	
50-60%	-		-		0.22		0.50	CORRECT	0.23		0.0001	
60-70%	-		-		0.02		0.22		0.51	CORRECT	0.01	
70-100%	-		-		0.0003		0.02		0.24		0.99	CORRECT

20%ile plots	Model signal = 5%		Model signal = 35%		Model signal = 45%		Model signal = 60%		Model signal = 85%	
Plotted range	Prob		Prob		Prob		Prob		Prob	
0-10%	0.94	CORRECT	0.0001		0.0000		0.0000		-	
10-30%	0.06	NULL	0.27	NULL	0.03	NULL	0.0001	NULL	-	NULL
30-40%	0.0000		0.51	CORRECT	0.25		0.006		-	
40-50%	-		0.21		0.49	CORRECT	0.10		0.0000	
50-70%	-		0.02		0.23		0.82	CORRECT	0.01	
70-100%	-		0.0000		0.0003		0.08		0.99	CORRECT

Median plots	Model signal = 5%		Model signal = 15%		Model signal = 25%		Model signal = 35%	
Plotted range	Prob		Prob		Prob		Prob	
0-10%	0.92	CORRECT	0.19		0.01		0.0000	
10-20%	0.08		0.64	CORRECT	0.22		0.01	
20-30%	0.001		0.16		0.56	CORRECT	0.23	
30-40%	0.0000		0.006		0.20		0.52	CORRECT
40-60%	-	NULL	0.0000	NULL	0.01	NULL	0.23	NULL
60-70%	-		-		0.0000		0.0003	
70-80%	-		-		-		0.0000	

Tables not shown for signals of 65%, 75%, 85% and 95% because results are symmetrical with those given here.

From these tables it is apparent that for signals in the middle of the range represented by a colour band, there is a high probability that either the correct colour band OR an adjacent one will be shown. Roughly speaking, the sampling resolution of our system is +/- one colour band. In some case the chances of appearing outside of this range are not negligible. For example, in a tercile plot where the model signal is 45%, there is a 0.02 probability of showing a signal in the 60-70% range, and where the model signal is 55%, there is a 0.02 probability of the signal being estimated as being in the 20-40% range, and will thus be plotted as if it were the climatological probability. Thus the risk of substantial model signals being mis-interpreted as "no signal present" is real, given that a global map contains many degrees of freedom.

In overall terms, our system has a moderate sampling resolution. To give more globally reliable estimates of the model signal would require a substantial increase in ensemble size, which would be expensive. An alternative strategy is to increase the ensemble size by pooling results from several different forecasting models. This has the advantage of starting to sample over errors in the models themselves, which are typically more serious than the sampling errors discussed in this section. This multi-model approach is implemented at ECMWF in the EUROSIP project, and is also the basis of the C3S seasonal forecast service which is replacing EUROSIP.

3.4 Climagrams

The climagrams show a monthly time-series of percentiles of the forecast pdf of an index, together with the corresponding percentiles of the model and observed climatology. Climagrams are created both for indices of atmospheric variability (including the Southern Oscillation, the PNA and the NAO), and for area-averaged temperature and rainfall indices.

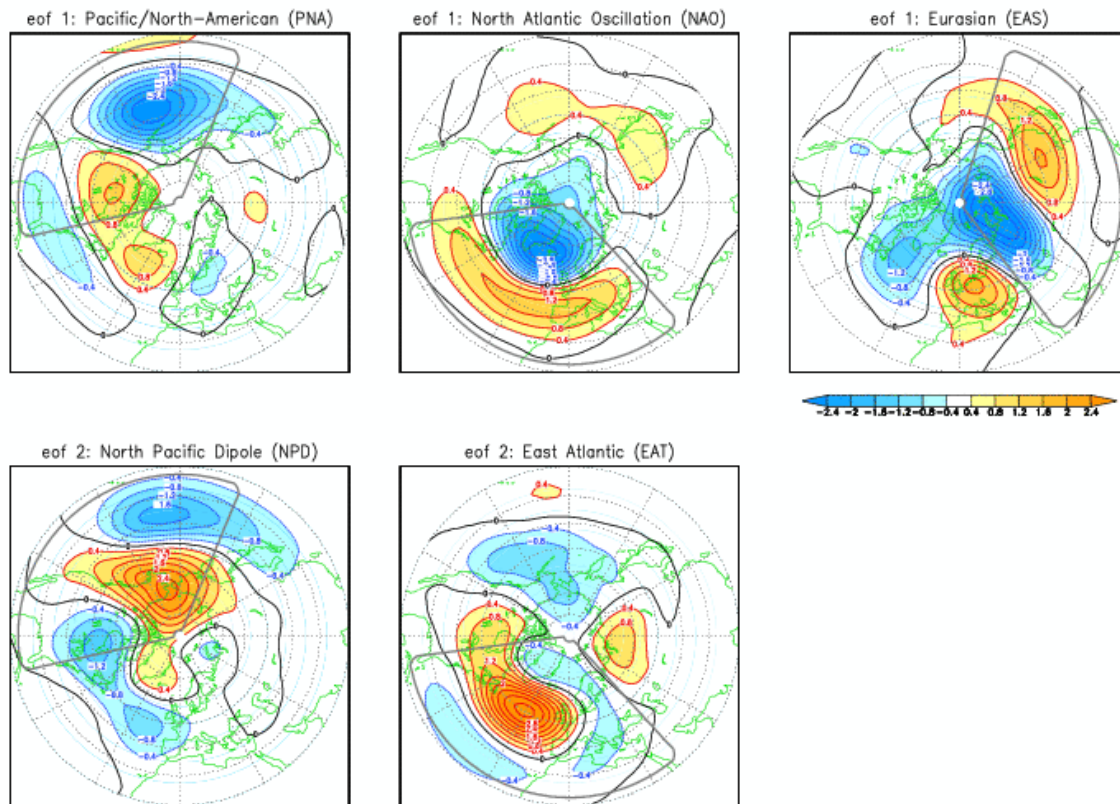
3.4.1 Equatorial Southern oscillation

The Equatorial Southern Oscillation index is defined as the difference between the standardized monthly anomalies of sea level pressure averaged over an area of the eastern Pacific (80°W - 130°W, 5°N - 5°S) and over Indonesia (90°E - 140°E, 5°N - 5°S).

3.4.2 Northern hemisphere winter teleconnections

A variety of statistical methods have been used in the literature to define Northern-Hemisphere (NH) teleconnection patterns. For the climagrams, leading variability patterns for the NH winter are defined by an EOF analysis of reanalysis monthly-mean geopotential height at 500 hPa in the December-to-March season, in the 30-year period (1981-2010). Geopotential monthly means are from ERA-Interim reanalysis until winter 2009/10, and from operational ECMWF analyses afterwards. The EOF analysis has been applied to three sectors in the latitude belt 25°-85°N: the

Pacific/North American sector (160°E-80°W), the Atlantic/European sector (80°W-40°E), the Asian/Pacific sector (40°E-160°E).

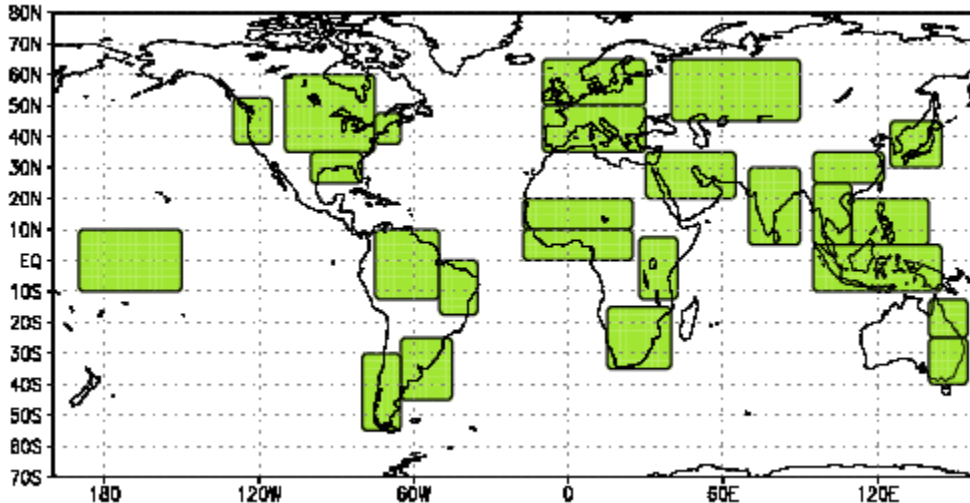


The first EOF defined in each of the three sectors corresponds to a well-documented teleconnection pattern: respectively, the Pacific/North American pattern, the North Atlantic Oscillation, and the Eurasian pattern. The positive phase of these patterns correspond to an intensification of the westerly winds over the central and eastern parts of the North Pacific and North Atlantic, and over central and eastern Siberia respectively.

The second EOF of the Pacific/North American and the Atlantic/European sector are also retained, since they modulate the intensity and position of the stationary-wave ridges over the north-eastern parts of the Pacific and Atlantic oceans. The sign convention for these EOFs (referred to as the North Pacific dipole and the East Atlantic pattern) is such that positive projections correspond to an amplification of the respective stationary-wave ridge.

3.4.3 Area averages of 2 metre temperature and rainfall

Area averages of 2m temperature and rainfall anomalies are computed over a set of 25 'grid boxes', shown in the figure below:



For 2m temperature, averages are computed using land fraction as a weight, in order to isolate temperature variations over land (2m T over sea is strongly constrained by the underlying SST, which SEAS5 reproduces). For rainfall anomalies, averages are computed over the whole area in each box. An exception is made for the 'Central tropical Pacific' grid box, which has no land points at the model resolution, and where no weight is applied for either variable.

The grid boxes were chosen to correspond to fairly homogeneous regions for seasonal-mean anomalies of both temperature and rainfall, with area-average values being positively correlated with anomalies at individual grid points over at least 80-90% of the area. However, for some areas, the box definition may be more suitable for one of the two variables, or for one particular season of the year (for example, the 'Sahel' box is optimised for summer rainfall).

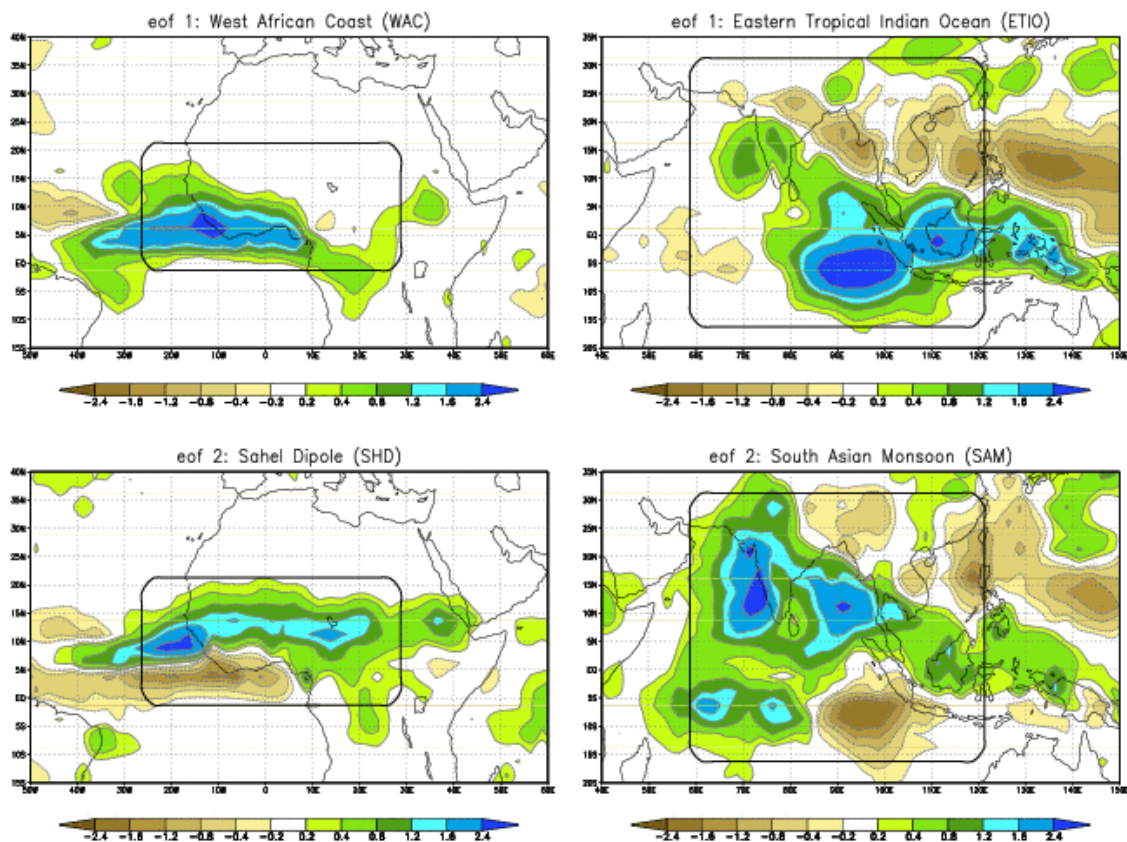
3.4.4 Monsoon indices

Indices of the large-scale distribution of monthly rainfall in the regions affected by the West African and South Asian monsoon were defined by means of EOF analysis. The two leading EOFs of monthly-mean rainfall anomalies from the GPCP (Global Precipitation Climatology Project) 2.5° dataset have been computed in the following domains:

- West Africa: 0°-20°N, 25°W-27.5°E
- South Asia - Indian Ocean: 15°S-30°N, 60° E-120°E

using June, July, August and September data from 1981 to 2005. The EOF patterns for the two regions are shown below, where the portion within the grey boundaries corresponds to the

regionally-normalized EOFs (for continuity, regressions of rainfall anomalies onto the corresponding PC are shown over a larger domain):




For West Africa, the first EOF represents rainfall anomalies along the Guinea Coast, while the second EOF has a dipole structure with opposite signs over Sahel and the south-western coastal regions. Projections on these EOFs are referred to as the West Africa Coast index and the Sahel Dipole index.

In the first EOF of the South Asia - Indian Ocean region, anomalies over Indonesia and the eastern tropical Indian Ocean have opposite sign to those over South-East Asia, the Bay of Bengal and the South China Sea. This pattern strongly resembles the rainfall response to summer ENSO episodes (the sign in the figure corresponding to cold events). Rainfall anomalies in the Indian sub-continent and the adjacent seas are dominant in the second EOF for this region, which also shows rainfall anomalies over the Indian Ocean resembling the response to the Indian Ocean SST dipole. According to the regions with largest anomalies, projections on these EOFs are referred to as the Eastern Tropical Indian Ocean and the South Asian Monsoon indices.

An All India Rainfall index is defined here as a weighted average of rainfall anomalies in the region 5° - 30° N, 70° - 90° E. The weights are proportional to the fraction of low-altitude land in each 2.5° box of the GPCP grid. This fraction is computed from the full-resolution land-sea mask and

surface topography of the ERA-Interim dataset, low-altitude points being defined as grid points with surface height less than 1000 m. The low-altitude criterion excludes points corresponding to the mountain regions of Nepal and Tibet. The land fraction is additionally set to 0 over Sri Lanka, values less 0.2 are discarded, and weights are finally normalized to have area average = 1. The final mask is shown below:



3.5 Tropical storm forecasts

The SEAS5 generates synoptic features analogous to tropical storms. These tropical storms are identified as tropical depressions with a "warm core" structure and a threshold strength - it is the warm-core structure which makes them dynamically equivalent to observed tropical storms, despite the fact that the model does not have sufficient resolution to produce the intensity of winds seen near the centre of a hurricane or typhoon. A tracking algorithm is applied to the 12-hourly upper air fields produced by each model integration, to locate and track individual tropical storms in the various ocean basins where they occur. Forecast products are then created by comparing the statistics of the tracked tropical storms between the forecast ensemble and the climatology derived from the re-forecasts.

3.5.1 Tropical storm numbers

For most seasonal forecast products (Niño SST plumes, rainfall and temperature anomalies), forecast values are calculated as an anomaly by comparison with the model climate in an additive manner. For some ocean basins, the model climate of tropical storm numbers differs substantially from the observed numbers (e.g. by a factor of 2). To produce forecasts of absolute numbers of storms (rather than e.g. just the number of storms as a percentile of the climatological distribution), the number of model storms is scaled multiplicatively such that the model climate

matches the observed climate. Precisely, a scaling factor is chosen by ordering the climatological distribution of tropical storm numbers, and taking the mean of those included in the range of the 25th-75th percentiles. The mean is calculated in this way from the observed climate over the calibration period (either a fixed period, or a moving interval) and the model climate (the same set of years, but 25 realizations from 25 ensemble members), and the scaling factor is the ratio of the two means. This approach is taken to avoid the means being unduly influenced by outliers in the distribution of tropical storm numbers. To estimate the standard deviation of the predicted number of tropical storms, the standard deviation of the original model forecast ensemble is scaled by the square root of the scaling factor applied to the mean number of storms. (If we were to model the number of storms in a season as a Poisson process, this square root scaling would be exactly the right way to scale the spread of the ensemble. We are not claiming that this 'model' is precisely the correct one to use, but we consider the derived scaling to be both reasonable and robust). From the 1st March 2021, forecast products for the northern hemisphere basins use a moving interval of the previous 10 years for the calibration. In the southern hemisphere, where there are no issues with low frequency variability, the calibration uses all available years. The tropical storm data used for calibration and verification comes from IBTrACSv4 (Knapp et al 2018). We use a combination of data sources: for the Atlantic and East Pacific basins, NHC best track reanalyses when available, and operational track data when not; and for other basins, JTWC data. (See Schreck et al. 2014 for the background to this choice). We download an updated version of the file every month, so both verification and the running 10 year calibration interval are kept up to date. Note that we assign tropical storms to the basin in which they first form and count them only once, so on rare occasions our numbers may not match those from other sources (e.g. in 2020, Tropical Storm Amanda-Cristobal formed in the East Pacific as Amanda, but crossed to the Atlantic as Cristobal: we count it as an East Pacific storm only).

The plots of tropical storm frequency show both the expected number of storms in each relevant basin for the coming six months, and the model-estimated standard deviation described above. Also shown is the result of a Wilcoxon rank-sum test on whether the 51 member ensemble forecast is or is not significantly shifted relative to the model climatological distribution. The period from which the observed climatology is derived, and for which the model climatology used in the w-test is derived, is stated on the plot labelling, and is in general the longest available period. It is not necessarily the same as the period used for calibration.

Forecasts of hurricane/typhoon numbers are shown in the same format, and are calculated using the same methodology, but with a higher threshold on model wind speeds to distinguish the stronger storms.

Forecasts are also provided of ACE, or Accumulated Cyclone Energy. This is calculated by accumulating the kinetic energy of each storm across its area of influence and its lifetime - so

large, intense, long-lived storms will contribute much more than small scale or short lived storms. These are shown as model forecast values normalized by the mean climatological model value from the re-forecast. An ACE greater than 1 means more energy than usual, less than 1 means less. Unlike storm numbers, which are counted depending on where a storm forms, ACE is assigned to the basin in which the relevant part of the storm track is located.

3.5.2 Track density

There are important interannual variations in the tracks of tropical storms, driven by large scale SST anomalies and other predictable factors. These changes in tracks are important for assessing the risk of landfall in various regions, and are just as relevant as variations in the total number of storms. Earlier forecast systems were not able to produce very good track information, in particular because the low resolution meant that the lifetime of model storms was too short. With the improvement of resolution over time, it has been possible to provide track density information since SEAS4.

The tropical storm density anomaly map shows the anomaly in the number of tropical storms expected to pass within about 300 km of a point. The numbers have not been calibrated, and so the map should be considered as a qualitative indicator of the expected anomaly. This plot naturally emphasizes areas where large numbers of storms usually occur - it is these areas that have large absolute variations.

The standardized tropical storm density standardizes the predicted anomaly against the variability in the model ensemble mean as measured in the re-forecasts. This plot shows where there are signals in both higher and more moderate storm density areas. To avoid over-interpretation of results, we only use three simple categories: "reduced expected value", where the anomaly in the ensemble mean is below minus one standard deviation of the ensemble mean, "enhanced expected value" where the anomaly is above one standard deviation, and "usual expected value" for values in between. Areas where the model prediction and climate are both for a track density of less than 0.5 are blanked out in white. Note that this is not a tercile probability map - it takes no account of how much variation there is between re-forecast ensembles in different years, but only standardizes using the variability of the ensemble mean.

3.5.3 Verification

For verification purposes, plots are provided showing the time-series of re-forecasts and forecasts of tropical storm numbers for each basin from the relevant start month. These plots also show the observed tropical storm numbers, extracted from IBTrACS4 as described above. The plots also include some verification statistics: the anomaly correlation between the predicted ensemble mean and the observed number of storms, and the cross-validated r.m.s. error in the tropical

storm number. These time-series plots are generated dynamically, and will include progressively more years of data during the lifetime of SEAS5.

Time series plots and verification information are provided separately for forecasts of tropical storm number, hurricane/typhoon number and ACE.

For the track density products, maps are shown of ACC between ensemble mean predicted track density and the observed values, for the 24 years of the hindcast set.

3.6 Annual range forecasts

Four times a year, annual range forecasts are run as an extension of the usual 7-month long seasonal forecasts. The primary purpose of the annual-range forecasts is to give an outlook for El Niño. The ENSO time-series plots are produced in the same way as the seasonal forecast. Verification information is also provided. Climagrams and tropical storm forecasts are not produced for annual range forecasts.

4. Seasonal forecast data products

Data products consist both of data produced by the forecast model, and various derived products which are calculated from this data and then encoded in GRIB for archiving and distribution. The data products give access to quantitative forecast values, and allow the creation of a range of user-specific forecast products. ECMWF member states can access all of the data directly from the archive system, and/or obtain real-time atmosphere forecast data via dissemination. Only a subset of the real-time data, as defined in the [Catalogue of ECMWF Real-Time Products](#), is designated for commercial use and available for dissemination to commercial customers. For the full set of archive data, including for the re-forecast period, see [Accessing forecasts](#).

4.1 Data streams and MARS retrievals

4.1.1 Data streams in MARS

IFS direct output (MMSF): The atmosphere model outputs many fields at 6-, 12- or 24-hour intervals, and a subset of these fields are archived in the MMSF stream. Accumulated fields contain the accumulated value of the field from the start of the forecast.

IFS monthly output (MSMM): Monthly means of the all output fields are automatically calculated and archived in stream MSMM for forecasts and re-forecasts. For surface fields, in addition to the monthly average (MEAN) field, fields consisting of the minimum (MIN) and maximum (MAX)

values occurring during the month at each grid point are formed for most variables. The MEAN, MIN and MAX values are calculated from the set of instantaneous output fields, with a 6h or 24h sampling interval, and so do not sample variations occurring on time-scales shorter than this. (For a few fields such as 2 metre temperature, minimum and maximum values on a time step by time step basis are tracked and archived). The standard deviation (SD) of the values used to calculate the monthly MEAN is also calculated. Since the available input data can be either 6h or 24h, for some fields the standard deviation includes the diurnal cycle, while for others it does not. These are archived under several types within the MSMM stream: fcmean (forecast mean), fcmax (the maximum value of the field occurring during the month), fcmin, and fcstdev (the standard deviation).

For accumulated fields in the MSMM stream, the monthly mean rate of accumulation is calculated. Since archived data are generally in SI units, monthly mean fluxes have convenient units, $W m^{-2}$. For rainfall, data is archived using the SI unit of $m s^{-1}$, and can be scaled by the user to a unit such as $mm day^{-1}$. Finally, for precipitation related fields MIN values are not calculated since in reality they are generally zero.

For the real-time forecasts, ensemble means of the 51-member forecast ensemble are calculated for all of the monthly mean fields, and archived in the MSMM stream as type EM.

For each forecast monthly mean field, for a given start date and lead time, the climate mean of the corresponding 1993-2016 re-forecasts is calculated. The climate means are archived as a new type HCMEAN in the MSMM stream. The date of the HCMEAN data is the date of the real-time forecast with which they are associated.

Wave model direct output (WASF): Direct output from the wave model at 24 hour intervals.

Wave model monthly output (SWMM): Monthly mean wave model output.

Real-time forecast anomalies (MMSA): Forecast monthly mean anomalies are calculated relative to a climate mean formed from the appropriate 1993-2016 re-forecasts. The anomalies are calculated for each ensemble member and for all of the monthly mean fields. These data are archived in the MMSA stream as type FCMEAN. The ensemble mean of the anomalies is also calculated for each monthly mean field, and archived in the MMSA stream as type EM.

4.1.2 MARS retrievals

All production data (re-forecasts, pre-operational and operational real-time forecasts) are archived as ORIGIN=ECMF, SYSTEM=5, EXPVER=0001. The month(s) to be retrieved are specified in terms of time into the forecast with fcmonth. Note that all seasonal forecasts start at 00 UTC. The annual range integrations (out to 13 months) are archived separately from the 7

month integrations, and are accessed by specifying METHOD=3 instead of METHOD=1. The first 7 months of METHOD=3 data for each extended integration is a simple copy of the corresponding METHOD=1 data.

If atmosphere data is retrieved on the archived grid, then the resolution will differ from that of SEAS4. An example showing how to modify a MARS request for SEAS4 data according to the details given above in order to retrieve the equivalent SEAS5 data is given in the table:

Mars request - SEAS4	MARS request - System 5
<pre> RETRIEVE, STREAM = MSMM, ORIGIN = ECMF, SYSTEM = 4, METHOD = 1, NUMBER = 0/TO/50, CLASS = OD, EXPVER = 1, DATE = 20110501, TIME = 00, TYPE = FCMEAN, LEVTYPE = SFC, PARAM = 51, FCMONTH = 1/2/3/4/5/6/7, TARGET = 2m_tmax_monthly </pre>	<pre> RETRIEVE, STREAM = MSMM, ORIGIN = ECMF, SYSTEM = 5, METHOD = 1, NUMBER = 0/TO/50, CLASS = OD, EXPVER = 1, DATE = 20170501, TIME = 00, TYPE = FCMEAN, LEVTYPE = SFC, PARAM = 51, FCMONTH = 1/2/3/4/5/6/7, TARGET = 2m_tmax_monthly </pre>
<pre> RETRIEVE, NUMBER = 0/TO/14, DATE = 19810501/19820501/ 19830501/19840501/ 19850501/19860501/ 19870501/19880501/ 19890501/19900501/ 19910501/19920501/ 19930501/19940501/ 19950501/19960501/ 19970501/19980501/ 19990501/20000501/ 20010501/20020501/ 20030501/20040501/ 20050501/20060501/ 20070501/20080501/ 20090501/20100501, TARGET = 2m_tmax_monthly_climate </pre>	<pre> RETRIEVE, NUMBER = 0/TO/24, DATE = 19930501/19940501/ 19950501/19960501/ 19970501/19980501/ 19990501/20000501/ 20010501/20020501/ 20030501/20040501/ 20050501/20060501/ 20070501/20080501/ 20090501/20100501/ 20110501/20120501/ 20130501/20140501/ 20150501/20160501, TARGET = 2m_tmax_monthly_climate </pre>

4.2 Archived atmosphere forecast data

The following tables detail the archived output of the atmosphere model.

4.2.1 Upper air fields

The following upper-air fields are available every 12 hours on the indicated **pressure levels**:

Parameter number	Parameter name	Pressure levels (hPa)
129	Geopotential	1000/925/850/700/500/400/300/200/100/70/50/30/10
130	Temperature	1000/925/850/700/500/400/300/200/100/70/50/30/10
131	U wind	1000/925/850/700/500/400/300/200/100/70/50/30/10
132	V wind	1000/925/850/700/500/400/300/200/100/70/50/30/10
138	Vorticity (relative)	1000/925/850/700/500/400/300/200/100/70/50/30/10
155	Divergence	1000/925/850/700/500/400/300/200/100/70/50/30/10
133	Specific humidity	1000/925/850/700/500/400/300/200/100/70/50/30/10
203	Ozone	1000/925/850/700/500/400/300/200/100/70/50/30/10

The pressure level data are spectral, apart from humidity and ozone which are grid point fields.

Monthly mean values are also calculated for these fields at the pressure levels listed in the table, and also the additional levels of 250/150/20/5/2/1 hPa. For all upper air fields, only the monthly MEAN is calculated, not MIN, MAX or SD.

The annual-range forecasts have a much reduced archive of pressure level data, since they are primarily designed for an ENSO outlook. The annual-range archive only contains pressure levels 850, 700, 500, 400, 300 and 200 hPa data, which (together with corresponding surface fields) allows tropical cyclone tracks to be calculated. 12-hourly data are not available for ozone or specific humidity. Monthly mean data for annual range forecasts are in principle available for all parameters at all levels.

Available every 12 hours as grid point data on selected **isentropic surfaces**:

Parameter number	Parameter name	Isentropic levels	Monthly mean output
60	Potential vorticity	315K/330K	MEAN

Available every 12 hours as grid point data on a constant **PV surface**:

Parameter number	Parameter name	PV level	Monthly mean output
3	Potential temperature	2000 (PV=2)	MEAN

4.2.2 Surface fields

The following four surface fields are output and archived at **step 0 only**:

Parameter number	Parameter name	Output frequency	Monthly mean output
26	Lake cover	step 0 only	
129	Surface geopotential	step 0 only	
172	Land-sea mask	step 0 only	
22008	Lake depth	Step 0 only	

The following fourteen surface fields are output and archived every 6 hours. Fields marked with an asterisk are accumulated fields.

Parameter number	Parameter name	Output frequency	Monthly mean output
34	Sea surface temperature	6h	MEAN/MAX/MIN/SD
139	Soil temperature level 1	6h	MEAN/MAX/MIN/SD
144	Snow fall*	6h	MEAN/MAX/SD

151	Mean sea level pressure	6h	MEAN/MAX/MIN/SD
164	Total cloud cover	6h	MEAN/MAX/MIN/SD
165	10m u wind component	6h	MEAN/MAX/MIN/SD
166	10m v wind component	6h	MEAN/MAX/MIN/SD
167	2m temperature	6h	MEAN/MAX/MIN/SD
168	2m dewpoint temperature	6h	MEAN/MAX/MIN/SD
169	Surface solar radiation downwards*	6h	MEAN/MAX/MIN/SD
175	Surface thermal radiation downwards*	6h	MEAN/MAX/MIN/SD
228	Total precipitation	6h	MEAN/MAX/SD
229	Instantaneous eastward turbulent surface stress	6h	MEAN/MAX/MIN/SD
229	Instantaneous northward turbulent surface stress	6h	MEAN/MAX/MIN/SD

The following 38 surface fields are output and archived **every 24 hours**. Fields marked with an asterisk are accumulated fields. Parameters 49, 51, 52 and 55 are not archived at step 0.

Parameter number	Parameter name	Output frequency	Monthly mean output
8	Surface runoff*	24h	MEAN/MAX/MIN/SD
9	Sub-surface runoff*	24h	MEAN/MAX/MIN/SD
31	Sea ice cover	24h	MEAN/MAX/MIN/SD
33	Snow density	24h	MEAN/MAX/MIN/SD
39	Volumetric soil water layer 1	24h	MEAN/MAX/MIN/SD
40	Volumetric soil water layer 2	24h	MEAN/MAX/MIN/SD

41	Volumetric soil water layer 3	24h	MEAN/MAX/MIN/SD
42	Volumetric soil water layer 4	24h	MEAN/MAX/MIN/SD
49	Maximum 10m wind gust	24h	MEAN/MAX/MIN/SD
51	Max 2m temperature in previous 24h	24h	MEAN/MAX/MIN/SD
52	Min 2m temperature in previous 24h	24h	MEAN/MAX/MIN/SD
55	Mean 2m temperature in previous 24h	24h	-
78	Total column liquid water	24h	MEAN/MAX/MIN/SD
79	Total column ice water	24h	MEAN/MAX/MIN/SD
137	Total column water vapour	24h	MEAN/MAX/MIN/SD
141	Snow depth	24h	MEAN/MAX/MIN/SD
142	Large scale precipitation*	24h	MEAN/MAX/SD
143	Convective precipitation*	24h	MEAN/MAX/SD
146	Surface sensible heat flux*	24h	MEAN/MAX/MIN/SD
147	Surface latent heat flux*	24h	MEAN/MAX/MIN/SD
170	Soil temp level 2	24h	MEAN/MAX/MIN/SD
176	Surface solar radiation*	24h	MEAN/MAX/MIN/SD
177	Surface thermal radiation*	24h	MEAN/MAX/MIN/SD
178	Top solar radiation*	24h	MEAN/MAX/MIN/SD
179	Top thermal radiation*	24h	MEAN/MAX/MIN/SD
180	East/West surface stress*	24h	MEAN/MAX/MIN/SD
181	North/South surface stress*	24h	MEAN/MAX/MIN/SD
182	Evaporation*	24h	MEAN/MAX/MIN/SD
183	Soil temperature level 3	24h	MEAN/MAX/MIN/SD

186	Low cloud cover	24h	MEAN/MAX/MIN/SD
189	Sunshine duration*	24h	MEAN/MAX/MIN/SD
205	Runoff*	24h	MEAN/MAX/MIN/SD
206	Total column ozone	24h	MEAN/MAX/MIN/SD
212	Top incoming solar radiation	24h	MEAN/MAX/MIN/SD
236	Soil temperature level 4	24h	MEAN/MAX/MIN/SD
243	Forecast albedo	24h	MEAN/MAX/MIN/SD
228008	Lake mixed-layer temperature	24h	MEAN/MAX/MIN/SD
228014	Lake ice depth	24h	MEAN/MAX/MIN/SD

The following derived field is not archived at daily resolution, but the monthly statistics are calculated and archived:

207	10m scalar wind speed	MEAN/MAX/MIN/SD
------------	-----------------------	-----------------

The annual range forecasts have a much smaller archive of daily values of surface fields, being SST, total precipitation, OLR, daily Tmax and Tmin, and 12 hourly values of MSLP and 10m winds. Monthly means are however available for the full set of parameters, just as for the seasonal forecasts.

4.2.3 Model level fields

For a limited number of integrations model level data is archived for the fields listed below:

Parameter number	Parameter name	Output frequency	Levels
129	Surface geopotential	12h	1
152	Log surface pressure	12h	1
130	Temperature	12h	1/TO/91/BY/2

131	U wind	12h	1/TO/91/BY/2
132	V wind	12h	1/TO/91/BY/2
133	Specific humidity	12h	1/TO/91/BY/2
138	Vorticity (relative)	12h	1/TO/91/BY/2
155	Divergence	12h	1/TO/91/BY/2
246	Cloud liquid water content	12h	1/TO/91/BY/2
247	Cloud ice water content	12h	1/TO/91/BY/2

This is primarily intended for use in dynamical downscaling. Because both the production and archiving of model level data can be very expensive, only a limited set of data is available. For the re-forecasts, only the first 5 ensemble members have model level data. For the forecasts, the first 11 members have model level data. In all cases, model level data is produced at 12h intervals from step 0 to step 4416 (i.e. the first 6 months only). Only every second level is archived (1,3,5,7,.....,89,91), a total of 46 levels. Two "surface" fields that are needed to reconstruct the model state are archived as level=1, following ECMWF convention. For the ensemble members and time step range for which model level data is output, sub-surface soil temperatures (fields 170, 283, 236) are output with the increased frequency of every 12 hours instead of every 24 hours. Monthly means are not calculated for model level data.

4.3 Archived wave forecast data

The wave model archives the following ten fields every 24 hours on the wave model grid (0.5° resolution) under stream WASF:

140220	Mean wave period based on first moment	MEAN/MAX/MIN/SD
140221	Mean wave period based on second moment	MEAN/MAX/MIN/SD
140229	Significant wave height	MEAN/MAX/MIN/SD
140230	Mean wave direction	-
140231	Peak period of 1d spectra	MEAN/MAX/MIN/SD
140232	Mean wave period	MEAN/MAX/MIN/SD

140233	Coefficient of drag with waves	MEAN/MAX/MIN/SD
140244	Mean square slope of waves	MEAN/MAX/MIN/SD
140245	10m wind speed	MEAN/MAX/MIN/SD
140249	10m wind direction	-

Monthly means are archived under stream SWMM. Wave direction cannot be meaningfully averaged as a scalar quantity (e.g. the average of 359° and 1° would be 180°!), so no monthly means are formed.

4.4 Ocean forecast data

Ocean data from the NEMO model is produced in netCDF format, with a complex native grid. The data are not available from MARS. Some raw data from the forecasts is archived to tape, but there are no immediate plans to make ocean data from the seasonal forecasts available to users.

5. Product interpretation

The ECMWF seasonal forecast system is a numerical system and the products are simply a statement of how the numerical calculations behave. Numerical products contain information on what is to be expected on seasonal timescales, but they also contain errors. Use of the raw numerical forecast products without interpretation is not recommended. Actual forecasts for users should be carefully prepared, perhaps combining data from several empirical and/or numerical sources. Creating and issuing properly prepared forecast statements is not a task undertaken by ECMWF, but is left to others, such as National Meteorological Services or appropriate international organizations. The probability maps on the web pages are uncalibrated - that is, they directly represent model output, and no adjustments to the probabilities have been made to account for model errors.

Correctly interpreting the seasonal forecast graphical products depends both on understanding the plot, and on understanding the characteristics of the forecasting system as a whole. In particular it is essential to use information about the past performance of the seasonal forecast, the spatial distribution of the forecast skill and the forecast reliability. Remember that the number of past cases is limited and sampling errors mean that it is easy to either over- or under-estimate model skill.

The significance tests shown on the spatial maps (see Section 3), measure how confident we are that the model forecast distribution differs from the model climate distribution. The significance

test is not a measure of confidence in the skill of the forecast, because it takes no account of observations. Skill estimates are given for forecast products based on re-forecast performance. In most cases, however, there is also a large sampling error on this estimate - in some cases skill will be underestimated, and in some cases it will be overestimated.

The ECMWF seasonal forecast model is global, has a surface grid with a 36 km spacing, and at its best can only hope to represent larger scale weather patterns. Local weather and climate is much influenced by features too small to be included in the model (hills, coastlines, land surface properties). Simply trying to read off local values from the maps could be very misleading. There are various objective methods which might in principle be useful for transforming the global-scale numerical model output into improved regional or local scale products. Study of patterns of variability (EOFs, covariance statistics etc.) may enable erroneous shifts in model variability to be corrected - if data records are long enough to be confident of the statistical robustness of any purported shifts. Downscaling techniques may be useful for obtaining local values from direct model output.

Interpreting rainfall anomalies in low rainfall regimes needs to be done carefully. There are substantial desert areas where the median rainfall in a month or season is zero, but heavy rainfall does sometimes occur. In such cases the climatological rainfall is positive, the model and observed climate means will be different, model forecasts will usually have ensembles where most members have zero rainfall, and "rainfall anomaly" as a deviation from mean climate may not be the most helpful concept to use.

6. Operational history

SEAS5 becomes operational in November 2017, replacing SEAS4. SEAS5 was run in pre-operational mode from January 2017, and a full 51-member ensemble exists for all following dates. SEAS4 ran in non-operational mode after the end of its operational phase until Feb 2018.

SEAS4 became operational in November 2011.

System 3 was operational from March 2007 to October 2011.

System 2 was operational from to January 2002 to February 2007.

In the early years of seasonal forecasting at ECMWF, the systems were run as experimental real-time forecast systems, and were not fully operational. Full operational status was achieved during the lifetime of System 2. System 1 produced forecasts routinely from January 1997 to December 2001.

7. References

[IFS cycle 43r1 documentation](#)

Balmaseda, Magdalena Alonso, Kristian Mogensen, and Anthony T. Weaver. 2013. "Evaluation of the ECMWF Ocean Reanalysis System ORAS4." *Quarterly Journal of the Royal Meteorological Society* 139(674): 1132–61.

Breivik, Øyvind et al. 2015. "Journal of Geophysical Research: Oceans Coupled Experiments." *Journal of Geophysical Research: Oceans* 120: 2973–92.

Craig, Peter D, and Michael L Banner. 1994. "Modeling Wave-Enhanced Turbulence in the Ocean Surface Layer." *Journal of Physical Oceanography* 24(12): 2546–59.
[https://doi.org/10.1175/1520-0485\(1994\)024<2546:MWETIT>2.0.CO;2](https://doi.org/10.1175/1520-0485(1994)024<2546:MWETIT>2.0.CO;2)

Dee, D. P. et al. 2011. "The ERA-Interim Reanalysis: Configuration and Performance of the Data Assimilation System." *Quarterly Journal of the Royal Meteorological Society* 137(656): 553–97.

Fichefet, T., and M. A. Morales Maqueda. 1997. "Sensitivity of a Global Sea Ice Model to the Treatment of Ice Thermodynamics and Dynamics." *Journal of Geophysical Research: Oceans* 102(C6): 12609–46.

Good, Simon A., Matthew J. Martin, and Nick A. Rayner. 2013. "EN4: Quality Controlled Ocean Temperature and Salinity Profiles and Monthly Objective Analyses with Uncertainty Estimates." *Journal of Geophysical Research: Oceans* 118(12): 6704–16.

Knapp, K. R., H. J. Diamond, J. P. Kossin, M. C. Kruk, C. J. Schreck, 2018: International Best Track Archive for Climate Stewardship (IBTrACS) Project, Version 4. NOAA National Centers for Environmental Information. <https://doi.org/10.25921/82ty-9e16>

Madec, Gurvan. 2016. NEMO Ocean Engine.

Monge-Sanz, B. M., M. P. Chipperfield, D. Cariolle, and W. Feng. 2011. "Results from a New Linear O3 Scheme with Embedded Heterogeneous Chemistry Compared with the Parent Full-Chemistry 3-D CTM." *Atmospheric Chemistry and Physics* 11(3): 1227–42.

Schreck, C. J., III, Knapp, K. R., & Kossin, J. P. (2014). The Impact of Best Track Discrepancies on Global Tropical Cyclone Climatologies using IBTrACS, *Monthly Weather Review*, 142(10), 3881-3899. <https://doi.org/10.1175/MWR-D-14-00021.1>

Tietsche, Steffen, Magdalena A. Balmaseda, Hao Zuo, and Kristian Mogensen. 2017. "Arctic Sea Ice in the Global Eddy-Permitting Ocean Reanalysis ORAP5." *Climate Dynamics* 49(3): 775–89.

Zuo, H. et al. 2017. “A Generic Ensemble Generation Scheme for Data Assimilation and Ocean Analysis.” ECMWF technical memorandum (795).

Zuo, Hao, Magdalena A. Balmaseda, and Kristian Mogensen. 2017. “The New Eddy-Permitting ORAP5 Ocean Reanalysis: Description, Evaluation and Uncertainties in Climate Signals.” *Climate Dynamics* 49(3): 791–811.

Comments

Comments and feedback on this document can be submitted using [this questionnaire](#).