Known IFS forecasting issues

Please note that numbering/ordering does ***not*** indicate/imply any sort of priority. Recent entries/changes/updates are shown in green. Greyed out means no longer current, but these issues can be relevant when examining archived forecasts.
Any enquiries related to the content of this page should be emailed to servicedesk@ecmwf.int (mentioning the "Known IFS forecasting issues web page").

Topic / title	Description	Related activities
2m Temperature		
T1. 2m temperature in the presence of inversions	In common with all models, 2m temperature forecasts from the IFS tend to have much larger errors, on average, during low level inversion situations, which are particularly common at high latitudes in winter. The basic physical explanation is that a set change in atmospheric energy content has a much larger impact on screen temperature in inversion situations than in unstable situations, because the energy change is commuted through a much smaller depth of the atmosphere (e.g. metres rather than kilometres). The lower the inversion, the larger is the potential error. There is also sensitivity here to the method we use to interpolate between air temperature at the lowest model level (~10m) and skin temperature (2m temperature is a diagnostic, not direct model output).	New reporting practices for radiosonde data ("BUFR" messages), slowly being introduced around the world, may alleviate this problem slightly, by providing model analyses with a much more detailed representation of the near surface layers. In regions with snow cover, where issues are often most apparent, a change in IFS formulation, from using a single-layer to a multi-layer snow scheme in June 2023 helped a bit (see also item S5 below).
T2. City temperatures too low	Due to the urban heat island effect not being represented, screen temperatures in large urban areas, particularly cities, are commonly too low compared to observations. The problem can be accentuated in winter by snow cover.	'Urban tiles' to be introduced in land surface scheme in due course.
T3. Screen temperatures fall too much near coasts	As a consequence of the radiation grid being larger than the model grid (due to computational constraints) night-time radiative cooling over land near to the coast is often too rapid. This is because cooling progresses according to T^4, and at near-coast points T is approximately the average temperature of the land and (warmer) ocean. As a result screen temperatures drop too much - related errors can sometimes exceed 10C. The problem is enhanced (i) when there is snow cover, (ii) at high latitudes, and (iii) where coasts have a convex shape (land-relative).	Improvements due to radiation code 'fixes' were introduced with cycle 41R2 in March 2016. In example cases the impact of these changes has been very positive. More substantial radiation code changes are likely in the longer term.
T4. Meteogram temperature issues in complex topography	In addition to the normal problems of representing screen temperatures in complex topography in current-generation global models, the user should be aware that the method by which screen temperatures on Meteograms are generated from model screen temperatures assumes a standard lapse rate (6.5^{o}C drop per km increase in altitude), and so if the difference in height between the site chosen, and the nearest model gridpoint (as shown in the ENSgram title) is large, the scope for large errors /biases increases. This is especially true in winter-time when inversions are more common: by definition an inversion implies a temperature increase with height, not a decrease, so the temperature correction applied could even be in the wrong direction. This issue is compounded by 2m Temperature issue T1 above.	Resolution upgrade in March 2016 (41R2) and June 2023 helped. Users can mitigate the impacts, in certain circumstances, by judicious selection of representative gridpoints. To help, guidelines on how meteogram data relates to model gridpoint data (with the MIR interpolation scheme) have been comprehensively updated in the Forecast User Guide - see here.

T5. China "cold spot"	There is a semi-permanent winter-time 'cold spot' over parts of central/eastern China. This can be most apparent in products that intrinsically display 2m temperature output in some 'anomaly' form - such as monthly forecast anomalies, seasonal forecast anomalies, and in the shorter ranges EFI and SOT. This is caused by a number of overlapping issues, whose (relative) importance can vary from day to day and from case to case: a) Sometimes the cold spot may not correctly reflect the IFS output in the sense that 2m temperatures are not always 'below normal' in this area when they are shown to be. In such cases the cold spot can owe its existence to incompatibilities between the current forecasting system, and ERA-Interim (ERA-I). ERA-I-based re-analyses are used to start the re-forecasts which form the 'model climatology' against which current forecasts are compared. So whilst these re-forecasts are rightly performed with the latest model version, they also inherit, as a starting point, auxiliary data such as snowfall from ERA-I, which feeds the ERA-I 'offline fields'. In turn this offline snow depth inevitably derives, in part, from what the ERA-I model puts on the ground in the way of snowfall, and this model's climatology is such that there is less snowfall in this area, on average, than in the current HRES. So HRES and other IFS components are inclined to have deeper snow cover in their analyses through the winter than the re-forecasts have in theirs, which encourages the development of 2m temperatures that are similarly lower - ie the 'cold spot' anomalies. b) Sometimes the 2m temperature output from the IFS is genuinely and continually too cold, e.g. by 10C or more. In such instances the cause can be insufficient melting, on the ground, in the autumn, of sleet/snow that fell in the model (see items S2 and S3 below). The snow remains on the ground, the air cools as a result, an inversion arises, more insolation is reflected out to space, and so on. This unwanted feedback loop is difficult to interrupt, in this part of the world, with mild air incursions, and the problem can persist for weeks or even months (as in late 2018). c) The IFS does not account for the drifting of snow on the ground. Snow that drifts will of course be redistributed, and will also sublimate more readily. Both can act to reduce snow depths in reality in a way that is not captured in the IFS, and real 2m temperature errors again result via the feedback loop referred to in (b). d) Issues (a), (b) and (c) are compounded by a historical dearth of snow depth observations in this area which might have helped bring things back on track, and also by an IFS snow analysis scheme 'feature' that anyway excludes all observations above 1500m (the said area is around 4000m). NOAA's daily northern hemisphere snow *cover* analysis, that does have complete spatial coverage, and that the IFS does utilise, is also discounted in this particular region because of the altitude. The 1500m cut-off can help avoid problems in complex terrain, though could be improved by using instead a measure of sub-grid orography to incorporate data over high level plateaus.	Understanding of the various factors contributing to this issue has improved markedly during 2018 through close monitoring at ECMWF. ERA5 started being used to initialise ECMWF re-forecasts in June 2019 (when cycle 46r1 went live). This helped alleviate issues relating to initialisation incompatibilities between re-forecasts and actual forecasts (mainly aspect a). Note however that this did not impact upon SEAS5 because the re-forecasts for that were generated before ERA5 was completed, and so had to use again ERA-I. On the other hand recent checks have shown that there are no clearcut incompatibilities between winter-time snow cover in the few SEAS5 forecasts that we have, and the SEAS5 re-forecasts. So the issues are complex and it is also possible that cold spot forecasts in this particular area have some integrity. Improvements to the snow analysis and snow physics are being worked on: multi layer snow was introduced operationally in June 2023 (aspects b, c and d). There is no longer a blanket exclusion of snow depth observations above 1500m (aspect d).
T6. Persistent "hot spots" in Central Africa	EFI and SOT values for temperature over some central parts of Africa are commonly strongly positive, and the SOT can also be extraordinarily high - it has exceeded 8 on occasion. This should be indicative of persistent extreme heat, relative to local climatology. Suggestions from various sources are that local conditions are mostly not extreme. Thus the signal is probably spurious (most of the time). Physically, this may relate to there being strong inter-dependence between soil moisture, cloud cover, precipitation and temperature in these areas. In the re-forecasts, which define the model climate (M-Climate), we use an offline-generated soil moisture analysis which may be in error - i.e. too moist - due to errors in precipitation in the driving ERA-Interim model in these regions. Another, less likely possibility is that soil moisture in operational runs, which is derived in a somewhat different way, is too dry. Perhaps both possibilities are contributing.	Since June 2019 (with cycle 46r1) the new re-analysis ERA-5 has been driving the re-forecasts: this helps as it is much more compatible with the current model. Longer term, increases in African observational coverage should also (1) improve the actual forecast, and (2) facilitate objective verification.
T7. High temperature spikes	Very rarely, under certain conditions and in certain locations, the IFS generates extreme, transient, 2m temperature spikes, which may differ by >10C beyond what is reasonable. Such extremes may last for just one time step. The impacts can be seen in various 2m temperature fields: the minimum, the maximum or a value for a given time. Impacts include (i) unrealistically high maxima and (ii) snapshot values for a given time that are below the minimum or above the maximum for a period that includes that time. The cause relates to fast switching to wet "tiles" by the surface exchange scheme and numerical instabilities that arise in conjunction, as well as to the way we post-process to get 2m temperature.	A partial fix has been created, by changing the thermal conductivity parameter for wet tiles. This will reduce the number of instances by almost half. It was introduced in 2019 (46r1). Work continues.
T8. Temperatures in the vicinity of deep lakes.	Deep lakes such as Lake Superior and Lake Malawi can present substantial problems for the single layer FLake model, because a change to the surface temperature is commonly applied, via assimilation, across too great a depth of water. This can impart undue inertia, such that the surface temperature, and 2m temperatures above and downwind, are too persistent during the forecast integrations. Forecasts at all lead times, from medium range through to seasonal, can be affected. There will also be impacts on other weather parameters. Less deep lakes, such as other 'Great Lakes' in North America, can also be affected, albeit in a less dramatic way.	

T9. Temperature errors related to vegetation	In quiescent conditions in springtime in particular, but perhaps also in autumn, forecast daytime temperatures can repeatedly be far too low in some extra-tropical regions. Errors of 5-10C have been relatively commonplace in short range forecasts in such conditions in recent years (e.g. in SW Russia). Initial investigations suggest that the error is because the IFS has too much leaf coverage (strictly the leaf area index or LAI), compared to reality/climatology, in the transition seasons. Although LAI in the IFS is nominally based on climatology, it is actually held artificially high during these seasons, because to do otherwise would lead to other larger errors, on average, for complex reasons that are not fully understood. With large LAI, as in the IFS when these errors occur, more insolation goes into latent heat at the surface (evapotranspiration), whilst in reality more should be available for sensible heating to help increase 2m temperature. There is also a positive feedback from the anomalous evaporation, which increases cloud cover too much, which in turn reflects back insolation and reduces solar heating.	Experiments show that in these particular situations correcting the LAI does correct the 2m temperature, even though in general it gives a degradation in scores. Investigations continue.
Precipitation		
P1. Marine convection propagation	In reality shower cells have a finite lifetime, so precipitation associated moves with the showers, as one can see on radar. In the IFS showers are instantaneous (as they are parametrised) and the related precipitation does not propagate. So showers triggered over the sea do not generally move inland in the model as they should. This can lead to under-prediction errors of several mm in inland locations, 10mm or more in extremis. The degree to which the error extends inland depends on the windspeed at the steering level for showers. For stronger winds the errors extend further inland. For snow showers the errors can be worse still, compounded by the relatively slow fall speed of snowflakes (up to say one tenth of that of raindrops). So a snowflake starting its descent at the coast might end up on the ground 100km inland, if winds are strong, whereas a raindrop in equivalent summer conditions might only propagate 20km before reaching the ground. Snow issue S1 relates.	2017 IFS changes (43r3) included detrainment of some hydrometeors from convection into large scale precipitation, bringing a small positive impact as those hydrometeors now drift with the wind. The new "moist physics", introduced in October 2021, changed many aspects of precipitation output (see here), but has not materially affected this issue.
P2. Underestimation of orographically-enhanced precipitation	As a consequence of topographical barriers being too low, in general (due to resolution), both the orographic enhancement of precipitation and the rain shadow effect tend to be underestimated in the IFS (more so in ENS than HRES, and more so in ENS after 10 days when resolution changes).	Resolution upgrades in March 2016 (41r2) and June 2023 (48r1) have helped. Integrated moisture flux diagnostics introduced in June 2018 (45r1), and EFI and SOT fields for those introduced in June 2019 (46r1), may also provide some assistance for predicting orographic rainfall. Removal of the 10-day resolution change, and later than that, in June 2023 (48r1), harmonisation of ENS and HRES resolutions to both be 9km improved the situation here.
P3. Underestimation of convective precipitation extremes	As a consequence of resolution, and the related parametrisation of convection, localised extreme values in precipitation totals will be systematically "underestimated" in IFS output. Differences equal to about one order of magnitude are possible. However this is not as bad as it seems, because when verified over areas that are the same size as the effective model gridbox size the agreement is generally much better. Note also that one should expect point maxima to be systematically higher in HRES than in ENS, due to resolution differences.	Resolution upgrades in Mar 2016 (41r2) and June 2023 (48r1) helped a bit. The ECMWF ecPoint precipitation downscaling initiative incorporates the estimation of sub-grid variability, and verifies very well, delivering much better predictions of rainfall at points, including extremes. Related 'point rainfall' products became available in April 2019 in ecCharts (documented here), and after that, in more limited form, in OpenCharts . See also P8. One impact of the new "moist physics", introduced in October 2021, has been to increase localised convective rainfall maxima at the gridscale (see here). Users should still expect maximum rain gauge values to be well in excess of a raw model forecast however, if sub-grid variability is large. Following the harmonisation of ENS and HRES resolutions to 9km in June 2023 (48r1) we no longer expect maxima to be systematically higher in HRES. One can sometimes get the impression from 10-day meteograms that this is not true, but that seems to relate to the practice of creating such a meteogram by clicking on an HRES precipitation maximum in OpenCharts or ecCharts.
P4. Tropical rainfall extremes greatest on day 1	If one examines the distribution, in forecasts, of daily rainfall totals for locations in the tropics, the (wet) tails tend to be longer for very short lead times (eg T+0 to T+24), implying that ENS has a greater propensity to generate extreme rainfall in short range forecasts than they do in medium range forecasts. For example the 99th percentile of daily rainfall at some locations at day 1 is twice what it is at day 3. This would appear to be a 'spin down' issue, of sorts, related to the handling of convection. Formulation of the EFI and SOT is such that they should intrinsically account for this (though note Miscellaneous issue M1 below), so the problem arises for the user particularly when referencing the direct model output.	ECMWF examined this issue closely during summer 2017. The causes are complex. The issues are still present in the 46r1 M-Climate (that uses ERA5). Further investigations in 2020 found a net spin down of average precipitation in the ENS, on day 1, globally, of about 20% in the first 12h, whilst HRES and Control did not exhibit any such spin down behaviour. These characteristics were much the same in cycles 46r1 (introduced 2019) and 47r1 (introduced 2020).

P5. Extreme rainfall at certain gridpoints ("rain bombs")	At particular gridpoints, that lie in areas of complex topography, IFS forecasts (notably HRES) can occasionally generate extreme localised precipitation totals in a matter of hours, say well over 100mm, when at neighbouring locations the amounts are far less. These extremes are incorrect; the error occurs in convective situations with light winds. Only a small number of gridpoints around the world are affected - mainly these lie in the following areas: southern China, parts of Eastern Africa, Papua New Guinea, along the Andes, southern Mexico. The cause is understood though is too involved to describe in detail here; in short it relates to a weakness in the semi-Lagrangian scheme (part of the model numerics). Tests have shown that the error will go away when the grid structure and model resolution are changed with the next cycle.	This issue was resolved with the introduction of cycle 41r2 in March 2016 (however the problem is still present in ERA5 output, because although that is also based on 41r2, it does not use the octahedral grid).
P6. Large-scale precipitation gradients match the land-sea mask	In some large-scale "warm rain" situations water bodies - sea and lake - can sometimes experience much more precipitation than adjacent landmasses in the IFS. This is unrealistic. The cause is complicated, relating to pragmatic historical tuning of cloud physics parameters, which have had some relationships with the underlying land-sea mask. In effect precipitation rains out much more readily over water bodies. Even small regions, such as the IJsselmeer in the Netherlands, or Lake Vanern in southern Sweden can exhibit these problems. In extremis rainfall rates can vary, unrealistically, by an order of magnitude across the land-water boundary (e.g. 0.2 versus 2.4mm/hr). Very warm moist air seems to be most prone to such issues.	A complex raft of modifications to the cloud physics schemes, which include reverting some earlier tuning, has delivered clearcut improvements in tests. Implemented in early June 2018 (45r1).
P7. Anomalous convective rings and squall lines in tropical rainfall	Animating HRES total precipitation rate in the Tropics one can detect, on occasion, instance(s) where a burst of localised high rates spawns an annulus, of increasing radius, of more modest rates, like a ripple from a stone in a pond. The initial burst tends to be large scale rain, effectively resolved convection, whilst the rings themselves are parametrised convection. Similar features do appear in nature, but the amplitude in the IFS is too large. The cause may be an over-active convective outbreak at the outset that triggers an over active adjacent downburst in response, which itself propagates outwards as a gravity wave, and which in turn helps trigger the convection scheme preferentially at its leading edge, where there is more forced ascent. Evidence has mainly been seen over tropical oceans, but also in the Sahel region of Africa. Spurious squall lines can also be created in IFS forecasts in a similar way (e.g. over Africa). These may propagate in a direction that is opposite to the propagation direction ordinarily observed when squall lines do occur in a given region.	Introduction of the new moist physics package, in October 2021, has alleviated some instances of this type. Squall lines can look much more realistic than hitherto (see here).
P8. Biases when forecasting rainfall at points - 'drizzle problem'	The IFS predicts gridbox average rainfall. Users often compare such gridbox forecasts with rainfall measured at points. Verification in this way reveals "biases", which are not always true biases but instead representivity issues, due to sub-grid variability. These apparent biases consist of over-prediction of small totals, and under-prediction of large totals, and are most apparent for parametrised convective precipitation which has greater sub-grid variability than large scale precipitation. Some users have referred to apparent over-prediction of small totals as the "drizzle problem", which is somewhat misleading if it relates to convective precipitation, which does not imply drizzle. In 2018 focussed verification at ECMWF, using high density observations, has shown that as well as the above representivity issue, in certain circumstances there can also be IFS over-prediction on the gridscale.	New ECMWF precipitation downscaling initiative incorporates the estimation of sub-grid variability, and also corrects for many weather-type-dependant gridscale biases. This verifies very well, delivering much better predictions of rainfall at points. This includes alleviation of the so-called "drizzle problem". The related "point rainfall" products became available in April 2019 in ecCharts, documented here, and later, in more limited form, in OpenCharts. See also P3.
P9. Validity of very small totals.	Due to grib packing issues for total precipitation fields, and related discretization, very small precipitation totals may not be accurately represented in ECMWF output. Accumulation values <0.04mm, in the period up to day 10, including those computed by subtraction, are unsafe; we recommend setting values below this threshold to zero. At longer lead times an even higher threshold may be appropriate.	GRIB 3 may help.
P10. Occasional convective rainfall in arid regions under-predicted.	Occasionally, particularly in areas and/or at times of year that are climatologically arid, the IFS correctly predicts daytime convective activity - as represented by the ECMWF lightning diagnostic - but zero rainfall at the surface. In such situations localised rainfall, at sub-grid level, can sometimes be observed. Whilst the gridbox average should not and will not capture localised sub-grid maxima, to be unbiased it needs to be greater than zero in such situations. There are two main candidates to explain the bias: (i) the main cause is believed to be the dependence of the convective parametrization on CAPE (convectively available potential energy). As CAPE can be relatively low in these situations, the amount of precipitation produced from the convection scheme is small, which subsequently evaporates before reaching the surface. (ii) a second potential contributor could be the evaporation of the rain in the dry air below cloud base. The evaporation is very dependent on the assumed drop size distribution for the rain and subgrid relative humidity variations in the sub-cloud layer. Larger drops in more humid air will penetrate further downwards before evaporating.	ECMWF's "point rainfall" post-processed output could in principle adjust for the bias, but in the present formulation the model gridbox rainfall forecast acts as a multiplying factor, so zeros are never altered. The new moist physics package introduced in October 2021 has helped to address this issue. It delivers rainfall more often in arid regions.
Snow		

S1. Snow drift in convective situations	When snow falls through cloud or beneath cloud it drifts with the wind. For large-scale (dynamic) precipitation IFS physics accounts for this. For convective precipitation however it does not; there is no drift, the precipitation arrives at the surface instantaneously once the convection is diagnosed, in the place that it is diagnosed. As a result snow arising from convective processes may be misplaced in the model (too far upwind), and the errors will be larger if winds along the snowflake path are stronger. Errors can be of order 100km. Precipitation issue P1 relates. The same issues exist for rain, but given the faster fallspeed of raindrops relative to the IFS model resolutions these errors are negligible. Clearly one also has to take account of the melting level.	2017 IFS changes (43r3) included detrainment of some hydrometeors from convective into large scale precipitation, bringing a small positive impact as those hydrometeors then drift with the wind.
S2. Snow on the ground takes too long to melt	In both ENS and HRES small amounts of snow on the ground tend to take too long to melt, even if the temperature of the overlying air is well above zero. This is because, for melting purposes, the snow that there is is assumed to be piled up high in one segment of a gridbox. For *smaller* nominal depths, the pile becomes higher, though at the same time covers a much smaller fraction of the box. The reason this is used is to improve the handling of screen temperature; by confining the snow to gridbox segments the impact on the temperature of that snow is reduced, and on average we find smaller errors and biases in 2m temperature as a result. The main downside is that snow cover pictures can look misleading, particularly at longer leads (when they can not of course be rectified by observational data). The cut-off above which snow is assumed to cover the full grid box is a 10cm depth - this is why a green hue used on standard snow depth charts on the web, which suggests to the eye the presence of some vegetation, disappears at 10cm.	Cycle 48r1 which went live in June 2023 had a minor positive impact on this problem, via introduction of a multi-layer snow scheme; previously there had been a single layer. Increments applied through the snow analysis scheme tend to be slightly smaller since then.
S3. Mixed rain /snow leads to snow accumulation	In marginal snow situations, when precipitation at the surface comprises both rain and snow, the snow component accumulates as lying snow. In the vast majority of cases this is wrong - it should melt instantaneously. This behaviour occurs because small snow depths within the model are assumed to be piled up into a small segment of a gridbox, and as such it is very difficult for them to melt quickly (as in Snow issue S2 above).	One related coding bug was identified and removed during the winter 2017/18, which helped a bit, but the problem of correctly representing the physics remains.
S4. Spurious snowfall in freezing rain situations	In certain winter situations, when snow descends through the atmosphere and melts to rain in a warm layer, before descending again through a cold (sub zero) layer, the model turns the precipitation back to snow far too readily. So surface precipitation in freezing rain situations commonly appears as snow, and that snow also accumulates on the ground. HOWEVER, it seems that where this precipitation is diagnosed as convective, this re-freezing problem does not exist.	Resolved with physics changes implemented in May 2015, though monitoring still required.
S5. Multiple snow layers	The model assumes that all snow on the ground has the same density (though that density does vary - e.g. increasing with age). So layers of different density, which arise in the real world, are not catered for. This can impact on several things, such as total snow water content, and upward heat conductivity, which in turn has the potential to adversely affect 2m temperature. In addition, *when new snow falls onto old, the change in snow depth is commonly less that it should be*, because the density assigned to the fresh snow depends in part on the density of the pre-existing lying snow, and so tends to be greater than it should. The magnitude of the error (in snow depth change) increases when the pre-existing snow is deeper and/or has a greater density. Example: if 10cm of new snow (ratio 12:1) fell onto 10cm of old lying snow (ratio 2.5), snow depth in the model would increase by only 3.5cm.	The multi-layer snow scheme introduced operationally in June 2023 (48r1) alleviated the majority of these problems.
S6. Analysed snow depths can oscillate between runs.	In certain scenarios - e.g. where deep snow is melting - the snow depth analysis can show much more snow at one of the main data times (e.g. 12UTC) than at the other (e.g. 00UTC). This can have a large detrimental impact on the 2m temperature forecast (compounded by issue S2). The cause is erratic reporting practice (in certain countries in particular) wherein zero depth reports may only be available at certain times of day. When they are not available the analysis scheme tends to interpolate between the available (non-zero) values, making both the spatial coverage and integrated snow volume across a region too large.	Improved reporting practices would address this issue. Assimilation-related solutions may also be possible, although creating and evolving different strategies for different countries is challenging and time consuming.
Tropical Cyclones		

TC1. Tropical cyclone intensity	Resolution limits our ability to fully capture the depth of some TCs, errors can be over 50hPa in extremis. The problems are larger for smaller systems, with a smaller eye - Super-typhoon Haiyan was one such example. Often minimum pressure in HRES will be lower than in all the ENS members, and likewise winds stronger than in all ENS members; this is because of the higher resolution of HRES. In such situations HRES guidance may be better, but not always.	With cycle 41R2 introduced in March 2016 came a resolution upgrade, and a substantial improvement to resolution used in the EDA. Key impacts were a marked reduction in positive depth bias in ENS analyses and forecasts, and more spread in ENS depth forecasts. A further resolution upgrade in June 2023 (48r1), to the ENS, meant that HRES and ENS ran at the same (9km) resolution thereafter, so whilst resolution remains a limitation the systematic resolution-related differences between HRES and ENS have naturally gone away.
TC2. Relatively slow-moving TCs can deepen too much in HRES	There is no coupling with the ocean in HRES. So for relatively slow moving TCs, when fluxes and mixing might in reality lead to a reduction in SST, the SST will remain at an elevated level and this can give the TC extra impetus to deepen too much (provided other factors such as shear remain favourable). For fast moving TCs the affected ocean is left behind, and so the problem is less acute or non-existent. Whilst issue TC1 above generates errors in the opposite sense that might sometimes fortuitously cancel, there are nonetheless recorded cases where TCs have been over-deepened, by as much as 50mb, because of the lack of coupling.	Coupling of HRES with the ocean began in June 2018 (cycle 45r1).
TC3. 10m wind maximum around a tropical cyclone	For tropical cyclone forecasts with accurate prediction of minimum pressure, the maximum wind speed is underestimated compared to estimates from official tropical cyclone warning centres. The difference can be partly explained by different wind speed definitions, and also the sub-grid scale nature of a local wind speed maximum. However, it is probably also related to the over-ocean drag parametrisation in cases of extreme wind speeds.	A change to the over-ocean drag parametrisation during extreme winds was implemented on June 30th 2020 with cycle 47r1; this increased the maximum wind in the vicinity of more extreme tropical cyclones (for a given central pressure). So the issue has been partly solved.
Winds		
W1. Under-estimation of strong gusts in convective situations	Although there is a helpful convective contribution in the computation of maximum gusts (as used in direct model output and the EFI), experience has shown that extreme gusts are generally under-represented, particularly when vigorous convection is involved, such as one might see with MCSs or squall lines - e.g. 60kt gusts or more might be observed when 30-40kt gusts are predicted. This relates to (i) an inability, at current model resolution, to represent the 3-d circulation around convective systems, and (ii) the fact that it is impossible to design an adjustment in the gust computation that will work in all cases.	Although a revised gust parametrisation was introduced in October 2021, which provides some general improvement, underestimation of extreme convective gusts will continue. New EFI parameters relating to severe convection were introduced in summer 2015; these provide useful pointers to when errors of this type are possible, as many cases have now shown. As part of a long-term initiative to improve CAPE representation ECMWF has progressively improved these parameters over the years, culminating in a final change in June 2023 (cycle 48r1). We now use a more continuous assessment of CAPE, based on hourly values, and a more physically correct computation method (previously a short-cut was used for computational reasons).
W2. Spurious short-fetch reduction in wind gusts	When a body of water lies downwind of land the 10m wind gust output parameter can exhibit an unrealistic localised reduction over that water, that spans a relatively short distance (perhaps 2 gridlengths). In one case the gust parameter changed from 23m/s over land, to 15m/s over a resolved lake. This issue probably relates to the fact that gusts are computed in part by adding to the mean wind a momentum transport component that depends on the vertical wind shear. It seems that if the shear is small, as can happen in the area just downwind (where roughness is small, sometimes because waves are also small), the reduction in the added component can be less than the concurrent increase in the mean wind, and so the net effect is that the diagnosed gust reduces, unrealistically. So users need to be aware that strong gusts could be underestimated just offshore (lake or sea). Lake Balaton is one area where this effect seems to have been noted.	Issue currently under investigation.
W3. Winds over mountains under-estimated	As some users have reported, wind speeds forecast over mountains (e.g. Norway, Iceland) tend to be too light compared to available observations. Errors can be particularly large when the geostrophic wind is large. 100m wind, available to customers as a standard model parameter, offers a viable and often more accurate alternative for predicting the mean 10m wind speed over mountains. However users must also consider the influence of observation site exposure.	This may relate to ways in which roughness is handled in the IFS. In recent tests of a new advection scheme wind representation over mountains has improved. Work on this complex topic continues.
Miscellaneous		

M1. Jumpiness in EFI and SOT, especially at short lead times	A consequence of the re-forecast strategy is that extreme events are sometimes not well sampled. Especially at short lead times, say 1 or 2 days, the 11 members that go up to make the re-forecast can be very similar, and so if the re-forecast dates (twice per week since May 2015) happen to be just before certain extreme events there may be some over-sampling, whilst if extreme events fall in-between the re-forecast dates, there may be some under-sampling. Thus the tails of the model climate (M-Climate) distribution can be jumpy as we move from one lead time to another, and as EFI and SOT depend heavily on these tails, much more than they depend on solutions around the median, they can be jumpy too.	The increase from 500 to 1980 re-forecast realisations effected in May 2015 has reduced the magnitude of this problem, but issues are still being identified from time to time (early 2023 comment). Adjustments to the re-forecast strategy are being closely considered in 2023; these may or may not alleviate this issue, depending on which of the related competing priorities is deemed most important.
M2. Sunshine duration irregularities	The integrity of this post-processed output parameter is strongly compromised by the radiation timestep in the model (3 hours in ENS, 1 hour in HRES), which because of computational cost is longer than the basic model timestep. This manifests itself in the sunshine duration parameter being (a) an undesirable function of longitude and (b) more generally unreliable.	Radiation code changes were introduced in cycle 41R2 (March 2016) which markedly reduced the dependance on longitude. Reliability also improved, providing a better match to the WMO sunshine definition. But see also more recent item M12 below.
M3. Sea ice evolution and associated weather	Sea ice cover does not change in any interactive way in the forecasts as we do not have a sea ice model. So none of the following are represented: sea ice formation due to low air temperatures, break up due to wind effects or melting, and advection by currents and winds. In turn this affects weather that relates, such as 2m temperatures over and downwind of, and convection triggered over water but not over ice. Wave model output will naturally also be affected. In the twice daily forecasts to day 15 sea ice cover is fixed. At longer ranges, to capture the seasonal cycle, there is relaxation towards the ice cover of the last five years; for monthly forecasts all ENS members use the average value, for seasonal forecasts members are divided into five sets (of 10) each one of which uses the pattern for one of those five years.	Sea ice model introduced operationally in Nov 2016 into ENS with cycle 43r1 (though note that HRES has no sea ice model as yet).
M4. Very poor SST evolution near New York	Due to the lack of resolution in the ocean component of the semi-coupled ENS system we are now running (introduced in Nov 2013 with 40R1), and an associated poor handling of the gulf stream wall, there is a major anomalous upward drift in SSTs over and S and E of the New York Bight (which itself lies just SE of New York city), in the first 10 days of the ENS forecasts. The area affected is about the size of England, and the size of the error that develops in 10 days can exceed 10C.	Problem alleviated slightly when ocean model in ENS went from 1.0 to 0.25 deg resolution in Nov 2016.
M5. 'Hot spots' near to glaciers	When cycle 41R1 was introduced on 12 May 2015 an error began to appear, over certain glaciated/partly glaciated regions (e.g. Iceland, the fringes of Greenland), on 2m temperature products that represent, directly or indirectly, anomalies. Affected fields include EFI/SOT (large positive values), Meteograms with climate (M-Climate too cold) and monthly forecasts (positive anomalies regularly forecast). This error is not a reflection of the absolute forecast values themselves - those should generally be OK - but is instead indicative of an error that was inadvertently introduced into the re-forecast suite. This error makes the 2m temperature forecasts in those re-forecasts, close to glaciers, much colder than they should be, causing the actual forecasts to look like they are indicating strong positive anomalies. Initially, in May, the error was less or non-existent because it was masked by residual seasonal snow cover.	The error has now been corrected. Affected re-forecasts will not themselves be rerun, although newly created re-forecasts should, from data times in early July 2015 onwards, be correct. This means that the mentioned issues slowly went away between early July and mid August 2015, as the fraction of the re-forecasts that were contaminated steadily reduced.
M6. 'Cold ring' around sea ice after day 10	Since the introduction of 41R1 on 12 May 2015 a ring of cold SST values (= -1.8C), about one gridbox wide, has appeared at the day 10/11 resolution change, along the edges of areas of sea ice. Locally SSTs may suddenly drop by more than 5C. This is a complex issue but relates to a change in the threshold at which sea ice cover is accepted (it is now 2%, it used to be 20%), the fact that we have to interpolate SST values onto a different grid when resolution changes, and the fact that SST is now set to -1.8C in gridboxes that include some sea ice.. A related complication is the fact that the two-way coupled ocean model runs on a different grid (1 degree). The ice cover threshold change was made to improve the handling of ocean waves in ice-margin zones, and to pave the way for introduction of a full sea ice model in the future. These advantages are considered to outweigh the 'cold ring' disadvantage. The day 10-15 EFI 2m temperature field can also be affected (showing a band of low values in the ice margin zones).	In March 2016 (cycle 41R2) the resolution change was moved to be at day 15, and so this issue no longer affects the twice daily ENS forecasts.
M7. Missing islands	A new land-sea mask was introduced with cycle 41R1 on 12 May 2015, in order to improve representativeness. Unfortunately there were some deficiencies in the source dataset, which has meant that a few islands, that should really be there, are no longer present in the IFS. The problems are mainly in HRES. Samoa is the island group where the greatest impact is seen. These issues will clearly reduce the utility of some IFS output, such as meteograms for some islands. For more details, including illustrations, go here. Note that the ocean wave model is not affected.	Corrected in March 2016, with the introduction of cycle 41R2.

M8. Seasonal lakes	In some locations lakes can undergo large changes in areal coverage, seasonally and with further modulation due to anomalous weather types. In the IFS lake areas are fixed, and are configured to match an ESA "GlobCover" dataset, which itself was derived by blending satellite images. So there are two potential sources of lake-cover-related errors in the IFS; one is inaccuracies in GlobCover, the other is temporal variations in lake size. In parts of Australia lake cover issues have had an adverse impact on forecasts of dew point, temperature, cloud cover and convection, not only in the immediate vicinity but also, via advection, in regions well beyond. Similar problems may also exist in other areas.	Through acquisition of more accurate lake-cover datasets - e.g. from forecast users - it may be possible to alleviate some of these problems. Non-interactive seasonal variation of lake extent is also being considered.
M9. Visibility biases	Visibility is a very difficult parameter to predict with a global model, a problem exacerbated somewhat by there being no aerosol emissions/transport. Since introduction of this diagnostic variable in May 2015 characteristics have been investigated, **over Europe**. Impressions of overall performance were positive, though the following general issues have been noted a) visibility in radiation fog tends to drop a bit too low, on average (e.g. 50m when 100m would be better) b) radiation fog formation tends to occur bit too late, and fog clearance a bit too early (e.g. 1-3 h typical bias in each case) c) background visibility (when no fog or precipitation) seems to be a bit too high overall d) hill fog seems to be under-represented (though this may relate to model orography limitations) e) visibility tends to be too low during rainfall f) visibility tends to be too high during snowfall	The model upgrade introduced in October 2021 specifically targeted items (a), (e) and (f). Performance is much better in these scenarios, with biases substantially reduced. Getting the details correct in radiation fog (a) remains challenging however. See here for more information.
M10. Sea ice in the Baltic	Sea ice cover analyses in the Baltic are 'discretized'. In other words certain ice configurations can appear often, others not at all. For example much the gulf of Finland tends to be relatively uniformly covered with sea ice of a set concentration, or just open water. In turn this is because satellite data, which feeds the Met Office OSTIA product used in the IFS, has particular difficulties in sensing ice where the water salinity is reduced by incursions of freshwater, as in the Baltic. This results in there being very few satellite sea ice pixels in the Baltic - the ones there are a long way from land - and in turn the information from these pixels gets copied into data-free areas nearby.	This problem has been alleviated somewhat via the use of a sea ice model, and via more recent coupling (June 2018) of HRES to dynamic ocean and ice models (see Fig 4 here).
M11. Stratospheric biases	IFS representations of the stratosphere suffer from various biases, which may be negatively impacting upon the skill of some longer range tropospheric predictions. These include different temperature biases at different levels, excess moisture 'leakage' from the troposphere into the lower stratosphere and unrealistic springtime breakdown of the polar vortex in the seasonal forecasting system (SEAS5). The representations of ozone and (large amplitude) gravity waves are also problem areas.	Model changes in cycle 47r1 (June 2020), in data assimilation and in the vertical interpolation methodology, have together substantially reduced stratospheric biases in the IFS. For example temperature biases in the analysis, above the 100hPa level, were reduced by up to 50%.
M12. Systematic overestimates of "sunshine duration"	The IFS "sunshine duration" parameter can, in some regions and some scenarios, provide large overestimates of sunshine duration relative to what is measured by conventional instrumentation. This can be in spite of there being, at the same time, a correct representation of the downward solar radiation in the IFS. Examples of this behaviour have been noted in summer time in Switzerland.	The way in which cloud optical properties are handled in the IFS is suspected to be the main cause of this discrepancy. ECMWF may be able to resolve this in future. A consequence of the new moist physics scheme introduced in October 2021 was that larger amounts of cloud were forecast than hitherto. Sunshine amounts reduce as a result, which may alleviate this problem. See here for more information.
M13. Convective Inhibition (CIN)	A weakness in ECMWF's CIN diagnostic has been identified. Sometimes values are much too high. The current advice is to not use this field for operational purposes.	ECMWF re-formulated and thoroughly tested new code for computing CIN, and this revision was introduced operationally with cycle 47r1 in June 2020.
M14. In thick fog relative humidity drops and temperature rises.	A bug introduced in cycle 47r3 in October 2021 has resulted in a tendency for 2m dewpoint and relative humidity (RH) values to sometimes drop to unrealistically low values once fog has been predicted/diagnosed in model output. This process is most apparent when dense fog is forecast (e.g. visibility < 100m). This is clearly unphysical and indeed wrong - RH should be close to 100% in such a situation. As these errors occur we sometimes also see a spurious increase in 2m temperature, with near surface superadiabats on model vertical profiles. In one NW Europe HRES case in Jan 2021, at one site, where visibility was forecast to be ~40m, the 2m temperature and dewpoint were respectively 0.5C and -10C. Both should have been ~ -2C.	This issue was investigated as a matter of priority. The cause was found to be positive feedback between two interacting and imperfectly represented mixing processes in the new moist physics scheme. A fix was implemented operationally with the 06UTC cycle on 22 Feb 2022.