TECHNICAL MEMORANDUM

769

Diagnosis of northern hemispheric regime behaviour during winter 2013/14

Mark J. Rodwell[1], Laura Ferranti[1], Linus Magnusson[1], Antje Weisheimer[1,2], Florence Rabier[1], David Richardson[1]

1) ECMWF, Reading, UK

2) Oxford University, UK

Forecast Department

November 2015

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

Series: ECMWF Technical Memoranda

A full list of ECMWF Publications can be found on our web site under:
http://www.ecmwf.int/en/research/publications

Contact: library@ecmwf.int

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

Abstract

During the northern winter of 2013/14, temperatures over North America were much colder than normal. On the other hand, Eastern Europe was anomalously warm, and strong storms and heavy precipitation affected Western Europe. The temperature signals are evident in the ECMWF ensemble (with smaller amplitude than in reality) at lead-times out to week four. The precipitation anomaly was predicted out to about 10 days. The study suggests that these anomalies were embedded within a hemispheric "regime" that was partly forced by tropical and underlying sea-surface temperatures via "Rossby Wave Source" forcing (associated with convection and its divergent out-flow). It is possible that this forcing may have been responsible for the moderate predictive skill found. Over the North Atlantic, we argue that enhanced baroclinicity associated with the planetary wave regime led to increased storminess, and suggest that the corresponding enhanced precipitation could also have acted to re-enforce the up-stream wave. Finally, we demonstrate the application of some new diagnostics of the data assimilation system (and of the ensemble of data assimilations) that are beginning to shed light on flow-dependent systematic model error and our representation of model uncertainty.

1 Introduction

During the 2014, the activity of the Diagnostics team was largely centralised within the new Forecast Department. The overall mission for the Diagnostics team is to organise activities aimed at understanding forecast deficiencies and to improve forecasts in collaboration with Research. An on-going activity is the development of diagnostic tools that help us investigate all aspects of the forecasting system - including accuracy and reliability of data assimilation and forecasts, and the model's mean climate and representation of key phenomena. Recently, there has been a growing emphasis on diagnostics of severe weather, regimes and regime transitions, and flow-dependent predictability (giving rise to situations where the outcome is highly certain and others that can lead to forecast "busts").

This technical memorandum is based on the first contribution to the Forecast Department Director's report, presented to the ECMWF Technical and Scientific Advisory Committees in October 2014. It aims to highlight the variety of diagnostic tools being used routinely, and in development, through their application to an important recent case-study - that of the highly anomalous Northern Winter period December 2013 – February 2014 (DJF 2014).

2 The observed anomalies

Figure 1(a) shows that, averaged over this period, Central Europe was mild, with mean temperatures up to 3K warmer than normal for this time of year and North America was very cold, with mean temperatures up to 5K below normal. By considering all the temperature anomalies over the Northern Hemisphere as a whole, it is possible to discern an apparent "wavetrain" over the North Pacific and North America. Notice also a strongly enhanced thermal gradient (baroclinicity) along the eastern coast of North America. An alternative view of the planetary-scale anomaly pattern can be seen in the 500 hPa height field - shown for the Northern Hemisphere in Figure 1(b). With this projection, it seems possible that the anomaly over the subtropical North Atlantic is also part of the same wavetrain. Notice also the strong low centre to the west of Europe. The mean vector wind anomalies associated with this feature are typically 10 ms^{-1} or more. Sensitivity experiments discussed later will be compared with the observed circulation using this 500 hPa height field. Figure 1(c) shows the observed precipitation anomaly field. Amongst other anomalies over the globe, Western Europe and the eastern North Atlantic

experienced heavy precipitation associated with very stormy conditions. There was also enhanced precipitation along the eastern coast of North America.

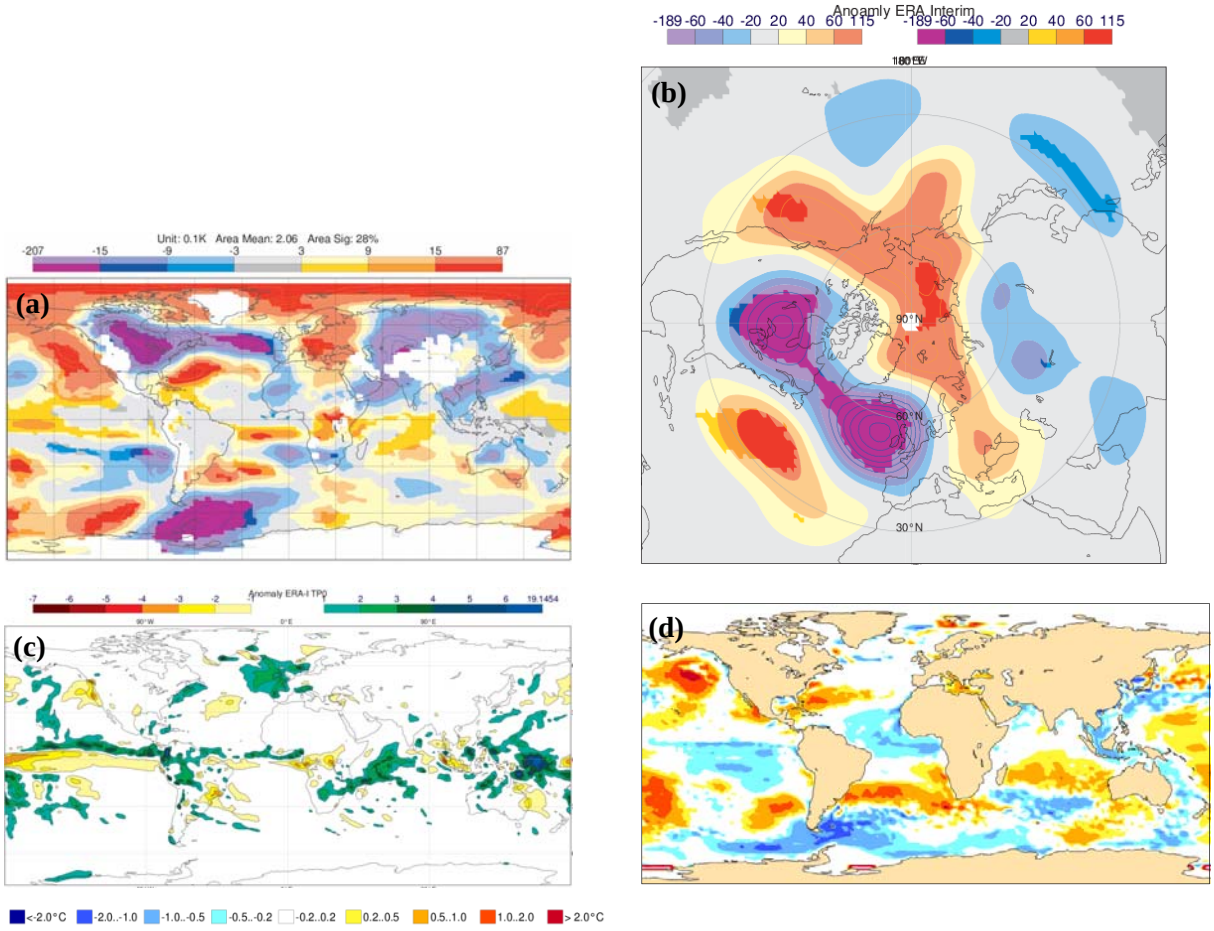

Figure 1 Mean observed anomalies, averaged between December 2013 and February 2014 (DJF 2014). (a) 850 hPa temperature anomaly. (b) 500 hPa geopotential height anomaly in m. (c) Precipitation anomaly in mmday^{-1} based on short (12-hour) forecasts. (d) Sea-surface temperature (SST) anomaly in K. In all panels, anomalies are relative to the ERA-Interim climate. Saturated colours in panels (a) and (b) indicate statistically significant mean anomalies at the 5% level allowing for auto-correlation in the daily data.

Despite the storminess over Europe, the medium-range forecasts for Europe had more skill (for 500 hPa height) than they had for the previous winter. A key question for forecasting at extended and seasonal ranges concerns the nature of the seasonal-mean anomalies. For example, was the European winter weather essentially a series of random storms that simply accumulated to a seasonal-mean? If this was the case, then predictability of the seasonal-mean would reduce quickly with lead-time. Alternatively, the storms might have been embedded within a larger-scale "regime" that was, itself, predictable to some degree. If this was the case then, although the individual storms may not have been predictable at longer lead-times, their aggregate effect on the seasonal-mean may have been predictable. Some support for this regime hypothesis can be found in the statistical significance test applied in Figure 1(a) and (b). The saturated colours in these panels suggest that we can be 95% confident that there is a seasonal-mean signal over-and-above the day-to-day "noise" for the North American and Central European

temperatures and the planetary wave pattern in the height field. This significant mean pattern could imply a teleconnection that provides predictability for the seasonal-mean anomaly in a remote region – such as over North America or Europe – either directly, or via its influence on transients (such as storms). Here we assess this regime hypothesis, and discuss what it says about potential predictability and future improvements to the forecast system.

A key potential source for predictability at extended and longer lead-times is likely to be from the oceans through the effect on the atmosphere of sea-surface temperature (SST) anomalies. There are other potential sources of predictability, such as from knowledge of sea-ice anomalies or the circulation in the stratosphere - including the quasi-biennial oscillation. Here, we focus on the SST aspect. Figure 1(d) shows the observed SST anomalies averaged over DJF 2014. Notice, in particular, warm anomalies over the eastern North Pacific, the subtropical North Atlantic, and the western tropical Pacific. There is some correspondence (but not perfect agreement) between these anomalies and the observed precipitation anomalies shown in Figure 1(c), and it is perhaps tempting to hypothesise a causal link whereby increased SSTs lead to enhanced evaporation, and thus more precipitation.

3 Forecast performance

Having examined the observed seasonal-mean anomalies, we now consider the performance of the ECMWF extended and seasonal forecasting systems. Figure 2(a) shows the ensemble-mean forecast of the 500 hPa height anomaly for DJF 2014 from the operational seasonal forecast "System 4". The signal is much weaker than that observed (notice that the contour interval has been halved relative to Figure 1b) and pattern agreement is by no means perfect. Nevertheless, it does show a statistically significant high anomaly over the eastern North Pacific and some increase in meridional gradient over the North Atlantic. The weakness of the signal could be because the observed anomaly was only weakly predictable, and/or it could reflect deficiencies in the model. Some evidence for the weakness of predictability comes from the fact that some individual members of System 4 were able to reproduce the observed seasonal-mean planetary wave pattern with the correct magnitude, together with the cold anomaly over North America and the enhanced precipitation over the Eastern North Atlantic.

We can obtain further understanding of the relative roles of predictability and model deficiency by running the atmospheric component of System 4 with prescribed (observed) SSTs. The ensemble-mean results (Figure 2b; now with the same contour interval as Figure 1b) show much better agreement with the observed anomaly. In particular, the high anomaly over the eastern North Pacific is reproduced with the observed magnitude, and there is an enhanced meridional gradient over the North Atlantic (although the trough over North America is a lot weaker than observed). Although the unrealistic nature of prescribed SST experiments can make them difficult to interpret in the context of the real-world, these results tend to point to weak inherent predictability of the coupled system, or ocean model deficiencies.

At sub-seasonal lead-times the ensemble (ENS) appears to indicate more skill in predicting features of the planetary wave pattern. For example, Figure 3(a) shows that the operational ENS (blue symbols) can predict the observed variations in weakly-mean North American 2m temperature (red dots) at a lead-time of 12-18 days. In addition to the predominance of correct cold forecasts, the relatively warm spell in the middle of January was also well predicted although its duration was perhaps overestimated.

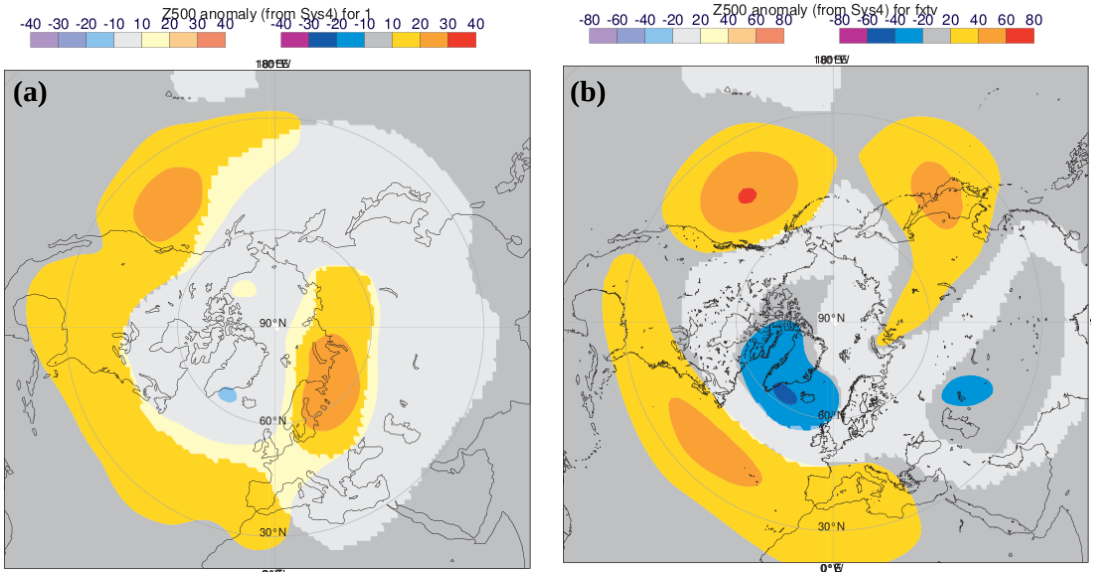

Figure 2 Ensemble-mean 500 hPa height anomaly for DJF 2014 based on (a) the operational coupled seasonal forecast System 4 (CY36R4) started on 1 November 2013 and (b) Integrations of the atmospheric model component of System 4 started on 1 November 2013 and forced with observed SST anomalies (run by Tim Stockdale). Both sets of integrations have 51 members and anomalies are relative to the same set of 450 re-forecasts started at the beginning of November (15 members for the 30 years 1981-2010). Statistical significance is indicated by saturated colours and based on the distribution of the 51-members (rather than of the 92-daily fields used in the test in Figure 1).

As a diagnostic of the ENS, it is interesting to know how well the ENS can reproduce the seasonal-mean anomaly for a range of sub-seasonal lead-times. Figure 3(b) shows the distribution of the average of the ENS forecasts shown in Figure 3a (and equivalent forecasts at different lead-times) over the DJF 2014 period. The observed seasonal-mean, North-American-mean 2m temperature anomaly of about -3K (red dot) is reproduced quite accurately, and with small uncertainty, by forecasts with lead-times 1-7 days (d1-7). This is associated with the atmospheric initialisation of the model - which generally includes a pre-existing anomaly - and the high degree of predictability at short ranges. The small discrepancy between observed and forecast values at this lead-time can be important for users, and is a current topic of investigation at ECMWF. It may be associated with interpolation and/or orography issues. As expected, as the lead-time increases, the ENS starts to lose the cold signal. However, it is interesting to observe that for lead-times between d12-18 and d22-28, the signal appears to partially "plateau" at anomalies of about -1K. This suggests the existence of extended-range "regime-predictability". The leadtimes showing monthly means (1 to 30 days, 31 to 60 days and 61 to 90 days) are based on the seasonal forecasting System 4. By aggregating the weekly mean forecast ranges, we can see that the monthly forecast at month 1 is better able to reproduce the seasonal-mean anomaly than is the seasonal forecast system at the same range. This may be due to the fact that the seasonal forecast is based on an older version of the model (but see later) and is run at coarser horizontal resolution. Figure 3(c) shows how well the ENS can reproduce the Eastern European warm 2m temperature anomaly. Again this is well reproduced with small uncertainty at leadtimes 1-7 days and a "plateau" at +1K is evident throughout the first month. Figure 3(d) shows the how well the ENS can reproduce the Eastern North Atlantic / Western European precipitation anomaly. There is a clear wet signal at leadtimes 1-7d and 5-

11d, but this is then lost more quickly; indicating that such extratropical precipitation anomalies may be relatively more difficult to predict than temperature anomalies.

In summary, therefore, the ENS showed skill in predicting the sub-seasonal variations in temperature out to at least 12-18 days, and was able to reproduce about half the seasonal-mean anomaly with lead-times up to 1 month. Beyond this lead-time, much of the signal is lost (unless we have much better knowledge of the true SSTs).

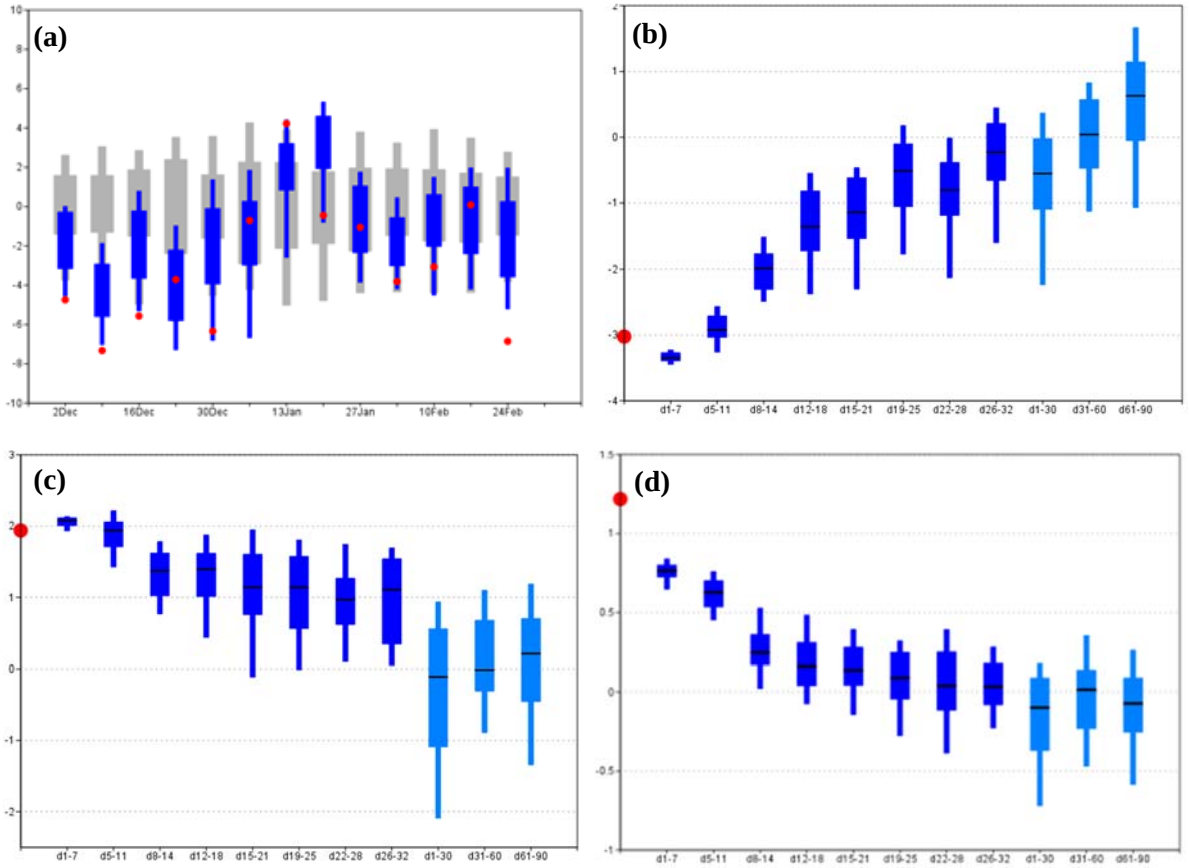

Figure 3 (a) Weekly-mean ensemble forecasts of North American (70-40N, 120-60W) 2m temperature anomalies in K from the monthly forecast system (CY40R1) at a lead-time of 12-18 days (blue symbols, indicating the 5, 25, 50, 75 and 95 percentiles of the 51-member ENS distribution). The corresponding observed weekly-mean temperatures are shown with red dots and the model climatological distribution (based on 5-member reforecasts from the corresponding calendar days for the years 1993-2012) is shown with grey symbols. Observed anomalies are relative to ERA-Interim re-analyses (for the corresponding calendar days for the years 1993-2012) and forecast anomalies are relative to mean model climatology. (b) Diagram showing the ability of the ECMWF ENS to reproduce the DJF 2014 seasonal-mean North American 2m temperature anomaly (red dot) for a range of lead-times. For each lead-time, symbols indicate the 5, 25, 50, 75 and 95 percentiles of the distribution derived by first averaging each nominal ensemble member forecast over the DJF 2014 period. Dark blue symbols are based on the monthly forecast system (CY40R1) and light blue symbols are based on the seasonal forecast System 4 (CY36R4). For consistency, the climatological distribution of System 4 is based here on the same years 1993-2012. It is important to note that the distributions shown are diagnostics of the ENS, and do not represent useable seasonal-forecasts since they are not available until the season has (nearly) ended. (c) As (b) but for Eastern European (70-40N, 0 -30E) 2m temperature. (d) As (b) but for Western European / Eastern North Atlantic (60-40N, 35W-10E) precipitation in mmday^{-1}.

4 The nature of the planetary wave pattern

Since the atmospheric model can reproduce some of the large-scale circulation anomalies when forced with observed SSTs this suggests that, in the mean, the planetary wave pattern may have "latched-on" to the observed SST anomaly pattern. We now examine further the nature of the mean planetary wave pattern and which aspects of the SST anomaly field might be playing a role in its existence. One way to gain a better understanding of the role of the (tropical) SST in forcing an extratropical circulation response is through the "Rossby-wave-source" (RWS) diagnostic. This quantifies the effect of the divergent circulation – e.g. associated with SST-forced atmospheric convection – on the large-scale Rossby waves (vorticity anomalies) along the Jet Stream. The shading in Figure 4 shows the observed (i.e. analysed) anomalous RWS for DJF 2014. The arrows show the anomalous divergent flow. What appears to be happening is that enhanced precipitation over the western tropical Pacific (at least partly forced by the SST anomalies there) leads to divergent flow that transports low planetary vorticity from the tropics into the subtropical western North Pacific (the blue region centred at 150^{o}E, 30^{o}N – "RWS1"). At the same time, negative extratropical RWS centres can be seen at 150^{o}W, 35^{o}N ("RWS2") and 75^{o}W, 35^{o}N ("RWS3") and these coincide with the enhanced precipitation and warm SST anomalies highlighted in Fig. 1. Since Figure 4 suggests these RWS features are associated with local divergence - the "Ballerina effect" - it is tempting to extend our hypothesis to suggest that outflow above SST-forced precipitation may be playing a role in "pinning-down" the mean planetary waves along the Jet Stream; as highlighted by the anomalous streamfunction contours, and the wavetrain previously highlighted in Figure 1b. A forth negative RWS centre is seen over the eastern North Atlantic (centred at 15^{o}W, 55^{o}N – "RWS4"). This is associated with the heavy precipitation experienced during DJF 2014, although not directly linked to any under-lying SST anomaly.

To investigate the relationship between the RWS and the Jet-stream wave anomalies, seasonal coupled hindcast experiments have been made where the atmosphere is strongly relaxed towards the observed (ERA Interim reanalysis) state within each of the highlighted negative RWS anomaly centres RWS1-4. Since we use the latest model version (CY40R1), we first produce a 28-member ensemble of unforced control runs starting on 1 November 2013. Figure 5(a) shows the ensemble-mean DJF 2014 anomaly field of 500 hPa heights relative to the model climatology. It appears that this model cycle is no better able to predict the observed planetary wave anomaly pattern (Figure 1b) than was System 4 (Figure 2a).

The DJF 2014 anomaly fields of 500 hPa heights relative to the model climatology for the experiments relaxed within the regions RWS1-4 are shown in Figure 5(b)-(e), respectively. For each experiment, aspects of the observed planetary waves are greatly improved – including upstream (west) of each RWS centre. Since each RWS centre is at least partly due to divergent outflow above a precipitation anomaly, it would appear that each of the highlighted precipitation anomalies plays a role in maintaining the entire Jet-stream wave. Further experimentation is required to (1) quantify the extent to which the precipitation anomalies are caused by the underlying SST anomalies (and how much they are consistent with moisture advection by the wave itself), and (2) quantify how much of the total RWS anomalies is explained by moist (rather than adiabatic) processes.

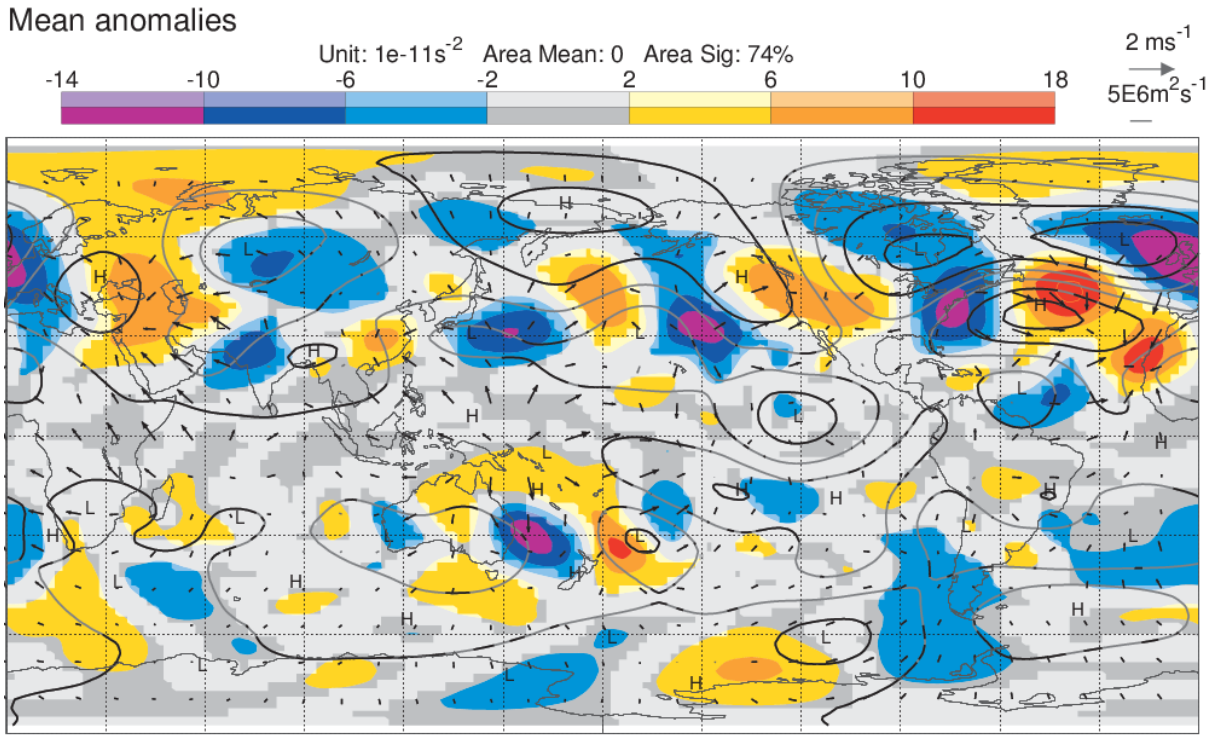

Figure.4 Shading shows the upper-tropospheric Rossby wave source (RWS) anomaly for DJF 2014. This has been spectrally filtered to highlight anomalies at the scale of the streamfunction anomalies (shown with contours). Arrows show the anomalous divergent winds. All fields shown have been integrated over the layer 300-100 hPa. Anomalies are relative to the 1981-2013 climatology and saturated colours, black contours and black arrows are statistically significant at the 5% level.

Figure 1(a) indicated enhanced temperature gradients between the warm western subtropical North Atlantic and the cold North American continent. These gradients are clearly associated with the SST anomalies and the Jet Stream wave, and act to increase the baroclinicity in this region – which is key to the developing storms that lashed the western coasts of Europe. It is interesting that the associated precipitation over the eastern North Atlantic and Europe during this time acts to produce a RWS anomaly that feeds-back positively on the mean Jet Stream wave (Figure 5e) – thereby helping to perpetuate the observed regime (along with the other RWS anomalies that are more directly linked to SST anomalies). This upstream effect could be purely tropospheric, or it could involve the modulation of the (stratospheric) Arctic vortex by the applied relaxation.

There is also some confirmation about our interpretation of the observed wavetrain pattern, and global precipitation and SST anomalies during DJF 2014 in the historical data. For example, cluster analysis applied to 500 hPa height shows that one of the dominant hemispheric regimes looks remarkably similar to Figure 1b (personal communication Franco Molteni, 2014), as does a composite of periods during 1981-2010 with cold anomalies over North America. Precipitation and SST anomalies coincident with either the hemispheric cluster (regime) or cold North American temperatures also have similarities with those observed during DJF 2014.

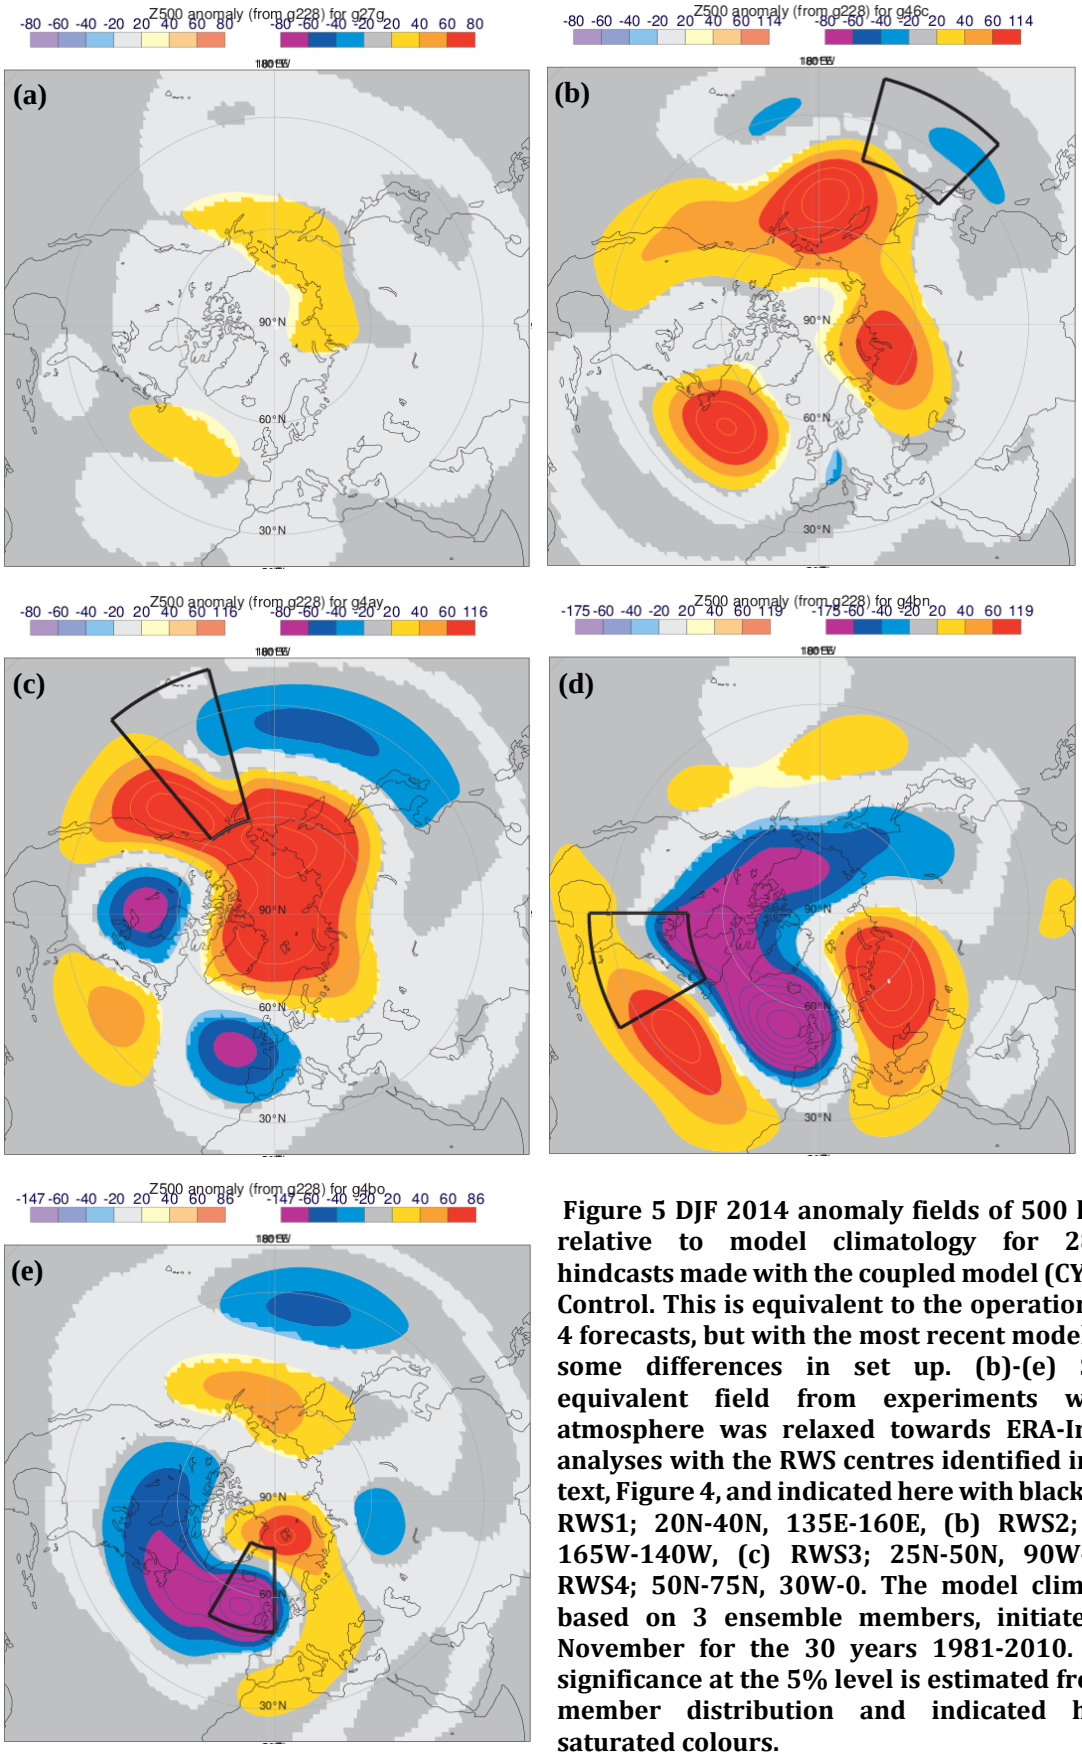

Figure 5 DJF 2014 anomaly fields of 500 hPa height relative to model climatology for 28-member hindcasts made with the coupled model (CY40R1). (a) Control. This is equivalent to the operational System 4 forecasts, but with the most recent model cycle and some differences in set up. (b)-(e) Show the equivalent field from experiments where the atmosphere was relaxed towards ERA-Interim re-analyses with the RWS centres identified in the main text, Figure 4, and indicated here with black boxes: (a) RWS1; 20N-40N, 135E-160E, (b) RWS2; 20N-60N, 165W-140W, (c) RWS3; 25N-50N, 90W-60W, (d) RWS4; 50N-75N, 30W-0. The model climatology is based on 3 ensemble members, initiated from 1 November for the 30 years 1981-2010. Statistical significance at the 5% level is estimated from the 28-member distribution and indicated here with saturated colours.

In summary, these results provide some support to the hypothesis that the DJF 2014 northern extratropical circulation anomaly represents a "regime" of the coupled system rather than being simply the statistical aggregate of unrelated synoptic events. It appears to have involved (1) forcing by tropical western Pacific SSTs and precipitation, (2) forcing by extratropical SST anomalies (via precipitation) that happened to be evenly separated by the wavelength for stationary Rossby waves, (3) increased baroclinicity, storminess and precipitation over the North Atlantic, and (4) some self-reinforcement of the planetary waves through the diabatic forcing of vorticity.

5 The synoptic scale

So far, we have concentrated on the existence and predictability of the hemispheric regime that appears to have been associated with the mean flow anomalies of DJF 2014. We now consider the synoptic-scale weather that was embedded in this larger-scale flow. How well does the model represent the salient aspects of the physical processes within the observed synoptic systems, and how well does the ensemble of data assimilations (EDA) "serve" the ensemble (ENS) forecast by accurately representing uncertainties in our knowledge of the true state?

The first of these questions is investigated through use of the "initial tendency / analysis increment" budget of the data assimilation system. Data assimilation is the process whereby new observations are combined with a previous short ("background") forecast to produce an estimate (analysis) of the new state of the atmosphere that is consistent with estimated errors in the background and observations. The "analysis increment" - the difference between the new analysis and background - can be considered as a correction to the background forecast. This is particularly true when averaged over many data assimilation cycles since chaotic uncertainties and random observation errors are averaged-away. This correction to the background can sometimes be related to issues in a particular parametrized process in the model, or to the model's "dynamical core". This is because the evolution of the background can be decomposed into the sum of the accumulated tendencies of each of the individual processes – each acting on a state that is as close to the truth as is currently possible to estimate. Figure 6 shows the overall budget. The sum of the process tendencies (Figure 6a-e), together with the analysis increment (Figure 6g) should match the analyses evolution (Figure 6h). Any discrepancy (or "residual", Figure 6f) is due to other process or numerics in the model that have been ignored here. As long as this residual is smaller than the increment, we can be confident that the increment is correcting one (or more) of the displayed processes. Figure 6 shows this budget for 500 hPa temperatures and winds averaged over DJF 2014 for the control (unperturbed) member of the EDA (184 data assimilation cycles = 92 days x 2 cycles per day). Because we have averaged over a whole season, the evolution term is very small. The residual term is also much smaller than the increment so we can be reasonably confident that the increment does reflect mean errors in one or more of the displayed processes. These processes are displayed with a larger contour interval. Latent heating associated with the precipitation is clearly evident in the "cloud" and "convection" processes. To some extent, this heating is offset by radiative and dynamical cooling. However, the cooling increment over Western Europe suggests that the moist physics may be over-doing the precipitation – in essence predicting more precipitation and extending this further into Europe than that observed (which, as we have seen, was already a lot wetter than normal). An alternative hypothesis - that the observations are biased cold - is less likely since different observation types (microwave radiance, and direct measurements from aircraft) tend to agree that the background is too

warm over Western Europe. This issue does not seem to be specific to the DJF 2014 season, and can also be seen in day 1 and day 2 forecasts that tend to be biased warm over Western Europe.

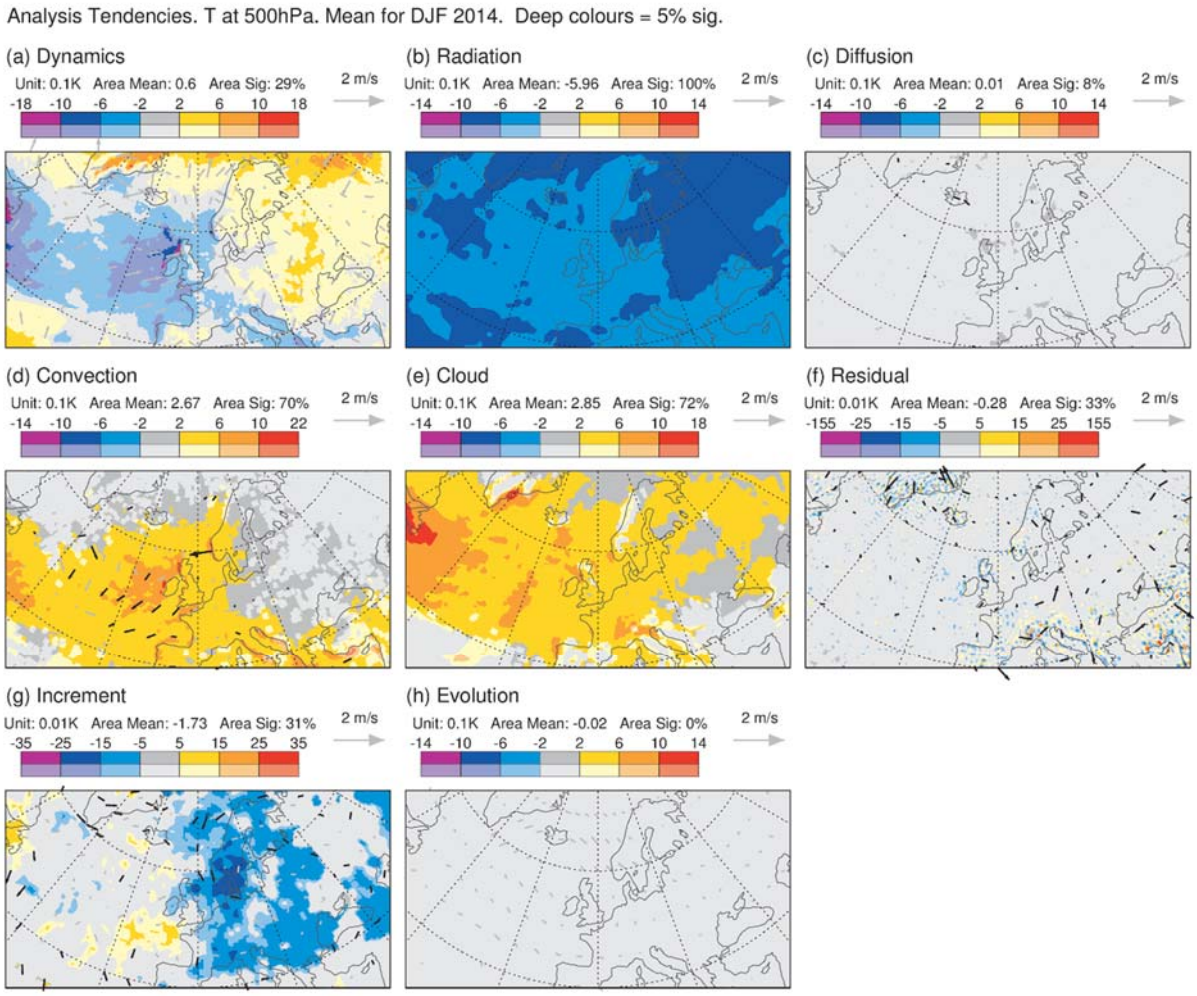

Figure.6 Initial tendency / analysis increment budget averaged over all 184 EDA control data assimilation cycles during the period DJF 2014. The budget is for 500 hPa temperature and wind. (a)-(e) Accumulated tendencies for the modelled processes Dynamics, Radiation, Vertical diffusion (and gravity wave drag), Convection and Cloud, respectively. (g) The mean analysis increment. (h) The mean analysed evolution (the mean tendency over the entire period). (f) The residual = (h)-(a)-(b)-(c)-(d)-(e)-(g). Statistical significance at the 5% level is indicated with saturated colours and black arrows.

In addition to improving mean forecast errors, it is increasingly important to improve the prediction of uncertainty in the forecast. This is particularly the case for short-range forecasts of severe weather – such as the intense storms experienced during DJF 2014 over Western Europe. A new diagnostic has been developed recently to assess how well the EDA is able to represent uncertainty in the initial conditions. This diagnostic is based on a variance budget of the EDA in "observation-space". In brief, when averaged over sufficient assimilation cycles, the predicted squared error in the ensemble-mean of the EDA background forecasts (relative to the observations – i.e. the squared "departure" in the language of data assimilation) should equal the sum of the squared estimated bias of the ensemble-mean, the

ensemble variance of the background forecasts (the squared "spread"), and the squared estimated observation error (hence the reason for calculating this budget in observation-space). Any discrepancy in this budget - in particular any large and statistically significant residual - would suggest that the EDA background has the wrong amount of uncertainty. Any ENS initialised from this EDA is likely to give poorer estimates of the probability of weather events – including extreme weather – than would an EDA with a well-balanced variance budget. (Note that for the current ECMWF ENS, initial spread is effectively inflated relative to the EDA spread using a tuning parameter based on past performance). In theory, all budgets for all observation types and at all grid resolutions should be balanced simultaneously, and this should provide a lot of power to the diagnostic.

Figure 7 shows such an EDA variance budget averaged over DJF 2014 based on AMSUA channel 5 microwave brightness temperatures (satellite measurements). This observation set is chosen because of its global coverage and the fact that is reflects mid-tropospheric temperatures – so that it is comparable with Figure 6, although note that the budget is based on squared quantities here. This new diagnostic highlights many interesting features but we will simply note here that the residual tends to be strongly negative and statistically significance in regions where the observation uncertainty is large – i.e. where observation density is lowest (over Western and Southern Europe for example). Such an imbalance may reflect issues with attempts to represent correlated observation errors – an important topic of current research. We also see a more direct issue with EDA spread over the western North Atlantic, where the residual may be indicating too much spread – with implications for the predicted uncertainty of down-stream development over the first few days of the ENS. Indeed, quantification of the traditional ENS spread-error relationship shows a ~20% excess of spread developing over the UK and Scandinavia between days 2 and 4.

In summary, new diagnostics of the data assimilation system are beginning to shed light on potential issues in the model's representation of the dynamics and physics and also on the representation of uncertainty in the EDA/ENS.

6 Discussion

The anomalous northern winter of 2013/14 was characterised by a planetary wave that brought cold weather to North America, and embedded synoptic systems that brought storminess, rain and milder temperatures to Europe. The ENS showed some skill in predicting the sub-seasonal variations in temperature out to at least 12-18 days, and was able to reproduce about half the seasonal-mean anomaly with lead-times up to 1 month. Results suggest that the northern extratropical circulation anomaly represented a "regime" of the coupled system that may have partly been forced by tropical western Pacific SSTs and corresponding precipitation (via the "Rossby Wave Source" associated with the convective out-flow), and by extratropical SST anomalies (again via precipitation) that happened to be evenly separated by the wavelength for stationary Rossby waves. We suggest that such links to SST may have been the reason for the (moderately enhanced) predictive skill found over the season. Increased baroclinicity, storminess and precipitation over the North Atlantic may also have provided some self-reinforcement of the planetary wave through the diabatic forcing of vorticity. New diagnostics of the data assimilation system are beginning to shed light on flow-dependent systematic errors in our representation of such physical processes, and diagnostics of the EDA are starting to highlight deficiencies in our flow-dependent representation of model uncertainty.

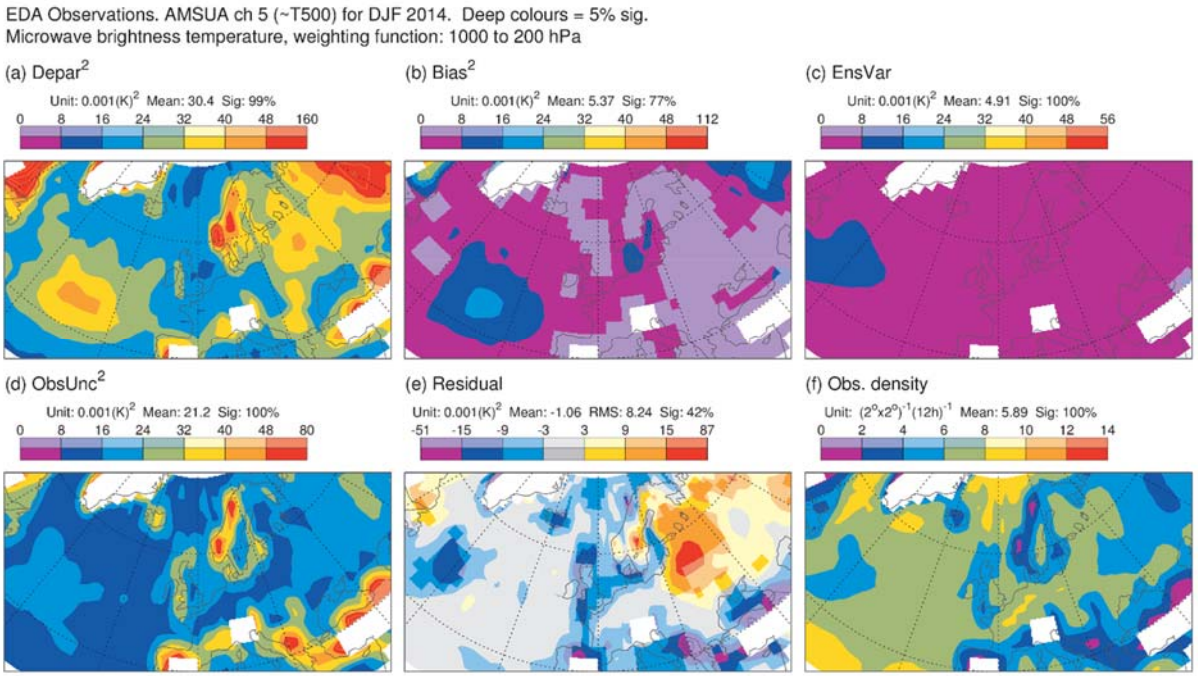

Figure 7 Variance budget of the Ensemble of Data Assimilations (EDA) averaged over all 184 EDA perturbed data assimilation cycles during the period DJF 2014. The budget is for AMSUA channel 5 microwave brightness temperatures (which measure mid-tropospheric temperature) binned onto a 2°x2° grid. (a) Mean-squared departure of the ensemble-mean of the EDA background forecasts (relative to the observations). (b) Mean-squared bias of the ensemble-mean. (c) Mean ensemble variance of the background forecasts (i.e. squared "spread"). (d) Mean-squared estimated observation error. (e) Residual = (a)-(b)-(c)-(d). (f) Observation density. Statistical significance at the 5% level is indicated with saturated colours.

The ultimate aim of Diagnostics is to have a beneficial impact on the development of ECMWF's forecast systems. Not all weather events get the same treatment as the winter of 2013/14 received but, through the routine monitoring of our analyses and forecasts using the diagnostic tools discussed above (and others), we hope to achieve this overall aim.