

# The TIGGE experience: Does the multi-model concept work for medium-range weather forecasts?

**Renate Hagedorn**

*ECMWF, Shinfield Park, Reading,  
RG2 9AX, United Kingdom*

## **ABSTRACT**

The validity of the multi-model concept for medium-range weather forecasts was assessed by comparing the relative benefits of multi-model forecasts provided by the THORPEX Interactive Grand Global Ensemble (TIGGE) project with reforecast-calibrated ensemble predictions from the European Centre for Medium-Range Weather Forecasts (ECMWF). Considering the statistical performance of global probabilistic forecasts of 850-hPa temperature and 2-m temperatures, a multi-model ensemble containing nine ensemble prediction systems (EPS) from the TIGGE archive did not improve on the performance of the best single-model, the ECMWF EPS. However, a reduced multi-model system, consisting of only the four best ensemble systems, provided by Canada, the US, the UK and ECMWF, showed an improved performance. This multi-model ensemble provided a new benchmark for the single-model systems contributing to the multi-model. However, reforecast-calibrated ECMWF EPS forecasts were of comparable or superior quality to the multi-model predictions, when verified against ERA-interim reanalyses. This improved performance was achieved by using the ECMWF reforecast dataset to correct for systematic errors and spread deficiencies. Further experimentation revealed that applying different weights to the individual model components of the TIGGE multi-model ensemble improved only marginally its performance. Therefore it does not seem worthwhile to employ computationally expensive tools for determining optimized weights when combining single-model ensembles of relatively similar quality..

## **1. Introduction**

The main motivation for investing into research activities on Numerical Weather Prediction (NWP) lies in the expectation that improved weather forecasts lead to enhanced socio-economic benefits. As such, the ultimate goal of all research related to NWP is to improve the quality and utility of weather forecasts. There are of course many ways to achieve this goal, ranging from work on the model system per se to research on the provision of user-optimized forecast products. All of these activities are valuable contributions to the general objective, and therefore none of the single efforts can be judged as more important than another. On the contrary, only through the diversity of approaches the overall goal can be achieved.

Post-processing of Direct Model Output (DMO) from NWP models is one of the many ways to improve weather forecasts. The term “post-processing” encompasses any means of manipulating the DMO in order to provide improved predictions. However, here we will concentrate on two specific methods: (i) combining single-model forecasts into a multi-model forecast, and (ii) calibrating single-model forecasts with the help of specific training datasets. Both of these approaches have been proven in the past to be successful in improving forecast quality. For example, the concept of multi-model forecasting has been extensively studied in the DEMETER project (Palmer et al., 2004), leading to a number of publications on the potential superiority of multi-model predictions on the seasonal timescale (see e.g. the special issue on the DEMETER project in *Tellus-A* 57(3)). Studying the rationale behind the success of multi-model ensembles, Hagedorn et al. (2005) concluded that “the key to the success of the multi-model concept lies in combining independent and skilful models, each with its own strengths and weaknesses.” In particular the fact that the performance of the single-

model ensembles varies, and thus in an operational environment the “best” model cannot be easily identified, makes the multi-model ensemble overall the most reliable choice. However, based on systematic toy model simulations, Weigel et al. (2008) and Weigel and Bowler (2009) demonstrated that even under the assumption that there is a clearly identifiable best single-model system, a multi-model ensemble can still improve the performance of this best model.

This result is particularly relevant in the context of applying the multi-model concept to medium-range weather forecasts, which has been at the heart of the THORPEX Interactive Grand Global Ensemble (TIGGE) project (Bougeault et al., 2010). First results from comparisons of the performance of individual TIGGE models indicated that in contrast to the seasonal timescale, where it can be difficult to define a “best” single-model which outperforms all other models on virtually all aspects, on the medium-range timescale it is much easier to identify a single-model which is clearly superior to all other models (Park et al., 2008). Therefore, the initial research question posed in the context of seasonal forecasting: “Does the combination of single-model ensembles with overall similar levels of skill lead to a more skilful multi-model ensemble?” changes in the context of medium-range forecasting to: “Does adding information from less skilful models to the best model lead to a more skilful multi-model ensemble?” As Weigel and Bowler (2009) pointed out in their theoretical study “it is possible to construct and combine reliable forecasts such that the multi-model has indeed higher skill than the best component forecast alone”, and early diagnosis of the TIGGE dataset confirms this theoretical result with real forecast data (Park et al. 2008, Johnson and Swinbank, 2009).

In order to assess the validity and potential of the multi-model concept also in the broader context of other post-processing methods, we will also discuss the calibration of single forecast systems with the help of specific training datasets. A number of different calibration methods have been proposed for operational and research applications, and a recent comparison of several methods can be found in Wilks and Hamill (2007). As most calibration methods are based on the idea of correcting the current forecast by using past forecast errors, they require some sort of training dataset. With this set of past forecast-observation pairs, correction coefficients for a regression-based calibration scheme can be determined. It has been shown that such calibration techniques are particularly successful when a “reforecast” training dataset is available (Hamill et al. 2004, 2006, 2008; Hamill and Whitaker 2006, 2007; Hagedorn et al. 2008). A reforecast dataset is a collection of forecasts from the past, usually going back for a considerable number of years or decades. In order to ensure consistency between reforecasts and actual forecasts, ideally the reforecasts are produced specifically with the same model and data assimilation system that is used to produce the actual forecasts. The availability of a large number of past forecast-observation pairs consistent with the current model system is a major factor of the success of the calibration technique used in this study.

One can expect that both these post-processing methods, the multi-model concept and the reforecast calibration, have their own strengths and weaknesses. Hence it is only natural to compare the potential benefits of both approaches, which is one of the aims of this publication. However, it is not our intent to come up with a final judgement on which is the better method, but instead to provide some indication for potential users to decide which approach might be the more appropriate choice for their specific circumstances. In contrast to Weigel et al. (2009) who have investigated a similar question on the seasonal timescale, this study concentrates on the medium-range timescale of forecasts up to 15 days.

A description of the datasets used can be found in section 2. The post-processing methods are presented in section 3. The results are discussed in section 4 with a summary and conclusions following in section 5.

## 2. Datasets

### 2.1. Forecast datasets

In this study forecasts from nine global Ensemble Prediction Systems archived in the TIGGE database at ECMWF are used. The main features of the model systems can be found in Table 1, together with a list of the model centres operationally running the forecast systems and providing the data for the TIGGE archive. Further detailed information on the model systems can be found in Park et al. (2008) or on the TIGGE website at ECMWF (<http://tigge.ecmwf.int/models.html>).

The investigations will focus mainly on the winter season December 2008 to February 2009 (DJF-2008/09), with some additional results also shown for the spring season March to May 2010 (MAM-2010). Results for the upper air field of 850-hPa temperature (T850), and the near surface variable 2-m temperature (T2m) will be discussed. The main comparisons will be done using forecasts starting at 00 UTC, the start time for which also ECMWF reforecasts are available. The comparisons involving all nine single-model forecasts from the TIGGE database are done for 12 UTC forecasts since some of the contributing centres produce forecasts only for 12 UTC and not at 00 UTC.

All forecasts have been interpolated to a common  $2.5^\circ \times 2.5^\circ$  grid using the interpolation routines provided by the ECMWF TIGGE data portal (<http://tigge.ecmwf.int>). Since the resolutions of the models are finer than  $2.5^\circ \times 2.5^\circ$  to varying degrees, the values at this lower resolution verification grid can be regarded as representing the average forecast over the  $2.5^\circ \times 2.5^\circ$  areas. One might expect that this sort of smoothing of the forecasts could improve the scores. However, sensitivity studies on the impact of the verification resolution have demonstrated that performing the verification on a higher  $1.5^\circ \times 1.5^\circ$  grid essentially did not change the result.

Centre	Horizontal resolution in archive	No. of vertical levels	No. of perturbed members	Forecast length (days)
<b>BOM</b>	$1.5^\circ \times 1.5^\circ$	19	32	10
<b>CMA</b>	$0.56^\circ \times 0.56^\circ$	31	14	16
<b>CMC</b>	$1.0^\circ \times 1.0^\circ$	28	20	16
<b>CPTEC</b>	N96( $\sim 1.0^\circ \times 1.0^\circ$ )	28	14	15
<b>ECMWF</b>	N200 ( $\sim 0.5^\circ \times 0.5^\circ$ ) N128 ( $\sim 0.7^\circ \times 0.7^\circ$ )	62	50	15
<b>JMA</b>	$1.25^\circ \times 1.25^\circ$	40	50	9
<b>KMA</b>	$1.25^\circ \times 1.25^\circ$	40	16	10
<b>NCEP</b>	$1.0^\circ \times 1.0^\circ$	28	20	16
<b>MetOffice</b>	$1.25^\circ \times 0.83^\circ$	38	23	15

*Table 1: Main features of the nine TIGGE model systems used in this study, at the time of producing DJF 2008/09 forecasts. BOM: Bureau of Meteorology (Australia), CMA: China Meteorological Administration (China), CMC: Canadian Meteorological Centre (Canada), CPTEC: Centro de Previsão de Tempo e Estudos Climáticos (Brazil), ECMWF: European Centre for Medium-Range Weather Forecasts (International), JMA: Japan Meteorological Agency (Japan), KMA: Korea Meteorological Administration (Korea), NCEP: National Centres for Environmental Prediction (USA), MetOffice: The UK Met Office (United Kingdom).*

The forecasts are mainly evaluated by calculating the Continuous Ranked Probability Score (CRPS) or its skill score (CRPSS), a diagnostic focussing on the entire permissible range of a certain variable (Hersbach, 2000). The CRPS is very suitable for systems issuing continuous probability forecasts, however, here we evaluate all forecasts in the classical way of probabilities retrieved from discrete ensemble members.

## 2.2. Training datasets

Every calibration needs a set of past forecast-observation pairs, also called training dataset. Usually it is beneficial to have a training dataset as large as possible to achieve a robust calibration. However, the existence of seasonally varying systematic errors suggests that it also can be beneficial to restrict the use of the available training data to only that data for a similar time of year. Using the ECMWF reforecast dataset (Hagedorn, 2008) enables us to satisfy both of these requirements. A set of reforecasts are operationally produced at ECMWF once per week, with start dates from the past 18 years. That is, on every Thursday the operational EPS is not only run for the actual date, but also for the same calendar day of the past 18 years. The only difference of these reforecasts to the actual EPS forecasts is the reduced number of perturbed members (4 instead of 50) and that ECMWF reanalyses instead of operational analyses are used in the initialization. Before 12 March 2009 a combination of ERA-40 reanalyses (Uppala et al., 2005) and operational analyses were used for the initialization, and from that date onwards only ERA-interim analyses (Simmons et al., 2007) have been used to provide a more consistent initialization dataset. The training dataset used in this study consists of the reforecasts produced for the five calendar days closest to the target date to be calibrated. In this way both the seasonally varying aspects are preserved and the number of forecast-observation pairs (18 years  $\times$  5 start dates = 90) should be sufficient for a robust estimation of the calibration coefficients, at least for the quasi-Gaussian variables studied here, 850-hPa and 2-m temperature.

At present, this type of reforecast dataset is only available for the ECMWF EPS and not for the remaining single-models of the TIGGE archive. That is, the reforecast based calibration can only be applied to ECMWF forecasts. However, a simple bias-correction procedure has been applied to all single-model systems and the multi-model ensemble. This calibration is based on a training dataset consisting of the last 30 days before the start date of the forecasts (Hagedorn et al. 2008).

## 2.3. Verification datasets

A number of considerations have to be taken into account when choosing the verification dataset to assess the performance of different single- and multi-models. On one hand, using model independent verification data like station observations ensures a fair treatment of all models. On the other hand, comparisons of the model performance over larger areas or for variables not directly available in observational datasets require the use of analyses, which commonly exhibit some of the bias of the forecast model used. There are a number of possibilities for the choice of analysis product in the context of comparing single- and multi-model predictions. The first option is to use each model's own analysis as verification dataset. However, this has the disadvantage that (i) the multi-model ensemble has no own analysis, and (ii) it would be difficult to draw conclusions from the resulting scores and skill scores when their calculation is based on different reference datasets. Another possibility is to use the average of all analyses of the participating models or some weighted average, also called multi-model analysis. Such an average analysis would fulfil the condition of being as fair as possible to all models participating in the comparison. On the other hand, averaging all analyses, including less accurate ones, might not necessarily lead to an analysis closest to reality. Additionally, such a multi-

model analysis cannot be used as verification dataset in this reforecast-comparison study because it is only available for the TIGGE forecast period, i.e. from 2007 onwards. This is not sufficient because the calibration of ECMWF forecasts based on the reforecast training dataset requires a consistent verification dataset for the entire training and test period, i.e. the verification dataset has to be available from 1991 onwards.

There are several candidate reanalysis data sets available, including the ECMWF ERA-interim (Simmons et al. 2007), ERA-40 (Uppala et al. 2005), and the NCEP-NCAR reanalysis (Kanamitsu et al. 2002, Kalnay et al. 1996). Considering the competing requirements of being as fair as possible to all models involved and being as accurate as possible, it was decided to place greater emphasis on the accuracy of the results and therefore to choose the ECMWF ERA-interim re-analysis (Simmons et al. 2007) as main verification dataset. The two important advantages of this choice are the acknowledged high quality of this analysis product (it used 4D-Var, a recent version of the ECMWF forecast model and a T255L40 resolution) and the availability of this dataset for the entire training and test period (1991 up to near realtime). The obvious drawback of this option is that the ERA-interim re-analyses are certainly not entirely independent of one of the models in the comparison, the ECMWF model. However, before discussing the relative performance of the forecasts itself, the sensitivity using different analyses as the verification is demonstrated. This will enable the reader to evaluate the potential impact of the choice of verification data set on the subsequent results.

The relative impact of using ERA-interim as verification instead of the multi-model analysis can be estimated by calculating the Continuous Ranked Probability Skill Score (CRPSS) from the CRPS that uses ERA-interim as verification ( $CRPS_{ERA-I}$ ) and the CRPS that uses the multi-model analysis as verification ( $CRPS_{MMA}$ ).

$$CRPSS = 1 - \frac{CRPS_{ERA-I}}{CRPS_{MMA}}$$

If the results were insensitive to the choice of verification dataset,  $CRPS_{ERA-I}$  and  $CRPS_{MMA}$  would be equal, and consequently the CRPSS would be zero. However, the negative values of the CRPSS in Figure 1a and 1b indicate a higher  $CRPS_{ERA-I}$ , i.e. a worse performance of the models when they are verified against ERA-interim instead of the multi-model analysis.

Generally, the skill scores worsen for all models, though the level of changes depends strongly on the lead time, variables or areas under investigation. For example, in tropical areas the analyses of the single-models vary extremely between different model systems (Park et al., 2008), and consequently using ERA-interim as verifying analysis would have a strong impact on the results. Therefore results are not discussed for the tropics in this study. The choice of verification dataset has least impact for longer lead times and extra-tropical upper air fields, with noticeable differences occurring mainly during the first three to four days (Fig. 1a). Assuming that we would like the impact of changing the verification to ERA-interim to be similar for each individual model, the CRPSS curves in Fig. 1a and 1b should be as close as possible. As such, it is important to note that *all models* follow the same general pattern of greater sensitivity to the choice of analysis for early lead times and less sensitivity at longer lead time. However, the performance of the CMC model worsens the most in the given examples, and one should keep in mind this fact when interpreting later results on the relative performance of the forecasts. There are larger differences between the models when looking at near surface variables like 2-m temperature, where each analysis may have its own systematic bias, and therefore the overall sensitivity to the choice of analysis is much more pronounced than for upper-air

fields (Fig. 1b, note the different scale to Fig. 1a). Additionally, it can be seen that up to a lead time of three days, the ECMWF T2m forecasts likely benefit from the fact that their initializing analyses are closer to the ERA-interim verification than the initialization of the other models.

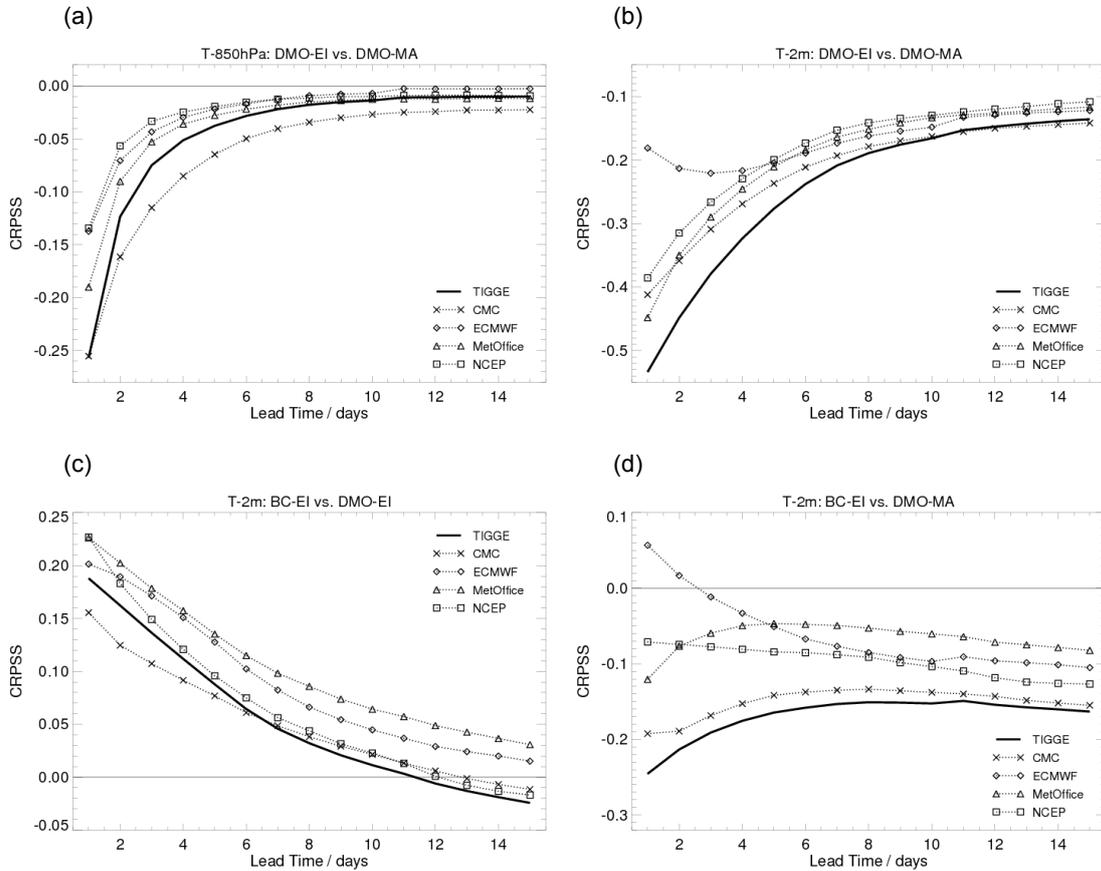


Figure 1: Illustration of the impact of the verification dataset and bias-correction on the relative skill of the predictions. The CRPSS is defined as  $CRPSS = 1 - CRPS(exp) / CRPS(ref)$ . Scores are calculated for forecast from the TIGGE multi-model (solid line) and the single-models (dotted lines with symbols; CMC: crosses, ECMWF: diamonds, Met Office: triangles, NCEP: squares) starting in DJF 2008/09 and averaged over the Northern Hemisphere ( $20^{\circ}N - 90^{\circ}N$ ). (a): 850-hPa temperature DMO forecast using ERA-interim as verification (exp) and the multi-model analysis as verification (ref); (b): 2-m temperature DMO forecast using ERA-interim as verification (exp) and the multi-model analysis as verification (ref); (c): 2-m temperature BC forecast using ERA-interim as verification (exp) and DMO forecast using ERA-interim as verification (ref); d: 2-m temperature BC forecast using ERA-interim as verification (exp) and DMO forecast using the multi-model analysis as verification (ref).

In order to reduce this stronger impact close to the surface, a bias correction (BC) is applied to all forecasts using the last 30 days of forecasts and ERA-interim analyses (for details see section 3.2). The positive values of the CRPSS shown in Fig. 1c indicate the improvements achieved by applying this bias correction. The calibration corrects the forecasts of all models most effectively at short lead times, with the MetOffice and NCEP achieving the largest skill improvement at a lead time of one day. For the remaining lead times, the MetOffice forecasts continue to gain the most from the bias correction, followed by ECMWF, NCEP, CMC and the multi-model forecast. The multi-model system generally gains the least from this explicit bias correction, because it has overall a bias of smaller magnitude than the contributing models. This implicit bias correction in the multi-model

concept arises as different models may have different biases that can compensate each other leading to a lower net bias (Pavan and Doblas-Reyes, 2000).

The final comparison between the bias-corrected forecasts verified against ERA-interim and the uncorrected forecasts verified against the multi-model analysis demonstrates that the bias correction reduces the strong lead-time dependence of the impact of using ERA-interim (Fig. 1d compared to Fig. 1b). Apart from the very early lead times of up to three days, the impact is now relatively similar for all models. All these results, together with the fact that a consistent dataset is indispensable for the reforecast calibration, justify the choice of using ERA-interim as verification dataset. However, for the time being we ask the reader to keep in mind these initial findings on the impact of using ERA-interim as verification dataset, and to relate all further conclusions to this initial discussion.

### 3. Post-processing methods

#### 3.1. Multi-model combination

The most basic way of constructing a multi-model ensemble is to simply combine the individual ensemble members from the contributing models with the same weight. This approach is not only an easy and robust way of combining different models, it also has been proven to be quite successful in improving on single-model predictions (Park et al 2008; Hagedorn et al., 2005; Shin et al., 2003). Since the main goal of this study is to assess the general validity of the multi-model concept for medium-range weather forecasts, we will first and foremost investigate the standard multi-model ensemble constructed by giving equal weights to all contributing members, noting that through the different number of members in the individual EPSs an implicit weighting will be applied. That is, model systems with a higher number of ensemble members will have a greater impact in the final multi-model prediction than model systems with fewer members. The performance of two equally weighted TIGGE multi-model ensembles will be compared with single-model forecasts. The first includes all ensemble members from the nine model systems listed in Table 1 (i.e. 248 members, 239 perturbed plus 9 control runs), and the second, called TIGGE-4, consists of only the 117 (113 plus 4) members of the CMC, ECMWF, MetOffice and NCEP ensembles.

There have been also many attempts to improve on the equal-weight method, with some of these studies claiming to be able to improve on this method (e.g. Krishnamurti et al., 1999; Robertson et al., 2004), others concluding that it is very difficult to achieve significant improvements (Peng et al., 2002, Doblas-Reyes et al., 2005; Johnson and Swinbank, 2009). As a contribution to this discussion, four different weighting schemes have been tested for the TIGGE-4 multi-model. The schemes range from more sophisticated minimization procedures to simple “intuitively” fixed weights:

- WGT-1: weights for single models determined proportional to their respective inverse mean square error in the training dataset
- WGT-2: weights for single models determined by “downhill simplex method” (Nelder and Mead, 1965), minimizing the CRPS in the training dataset
- WGT-3: weights determined from a random distribution
- WGT-4: weights fixed to:

$$W_{\text{CMC}} = 0.1, W_{\text{ECMWF}} = 0.6, W_{\text{MetOffice}} = 0.2, W_{\text{NCEP}} = 0.1$$

For all four schemes it was ensured that the individual single-model weights summed up to one.

### 3.2. Bias correction

A bias-correction procedure is applied to near-surface fields. As mentioned in section 2.3, this procedure aims to reduce the impact of using ERA-interim as verification dataset (see also Fig. 1c and 1d) and makes the comparison with the reforecast-calibrated ECMWF forecasts fairer. In fact, since the reforecasts are only available for the ECMWF EPS, the calibration procedure applied to the remaining models can only be based on a training dataset consisting of a limited number of previous forecasts. Taking into account this restriction, we apply a bias-correction procedure based on the past 30 days of forecasts (BC-30). The procedure itself calculates at every grid point  $x$  and for every lead time  $t$  a correction  $c(x,t)$

$$c(x,t) = \frac{1}{N} \sum_{i=1}^N e_i(x,t) - v_i(x,t)$$

as the average difference between the ensemble mean  $e(x,t)$  and the verification  $v(x,t)$  for all  $N$  cases in the training dataset. This correction is applied to all ensemble members of the forecast to be corrected, i.e. the ensemble distribution itself is not altered in itself but is shifted as a whole. However, because this type of 30-day bias correction has proven to have a significant impact only on near-surface variables like 2-m temperature, the results shown for the upper-air variable 850-hPa temperature are not based on BC-30 forecasts but compare DMO forecasts with the reforecast-calibrated ECMWF EPS described in the next section.

### 3.3. ECMWF reforecast calibration

The availability of the ECMWF reforecast dataset enables the application of more sophisticated calibration techniques than the simple bias-correction described above. Here we use a combination technique “EC-CAL” based on the Non-homogeneous Gaussian Regression (NGR) and results from the pure bias-correction (BC). The NGR technique itself is described in detail in Gneiting et al. (2005) and has been already previously applied to ECMWF EPS forecasts (Hagedorn et al., 2008). Essentially, NGR is an extension to conventional linear regression by taking into account information contained in the existing spread-skill relationship of the raw forecast. Using the ensemble mean and the spread as predictors, it fits a Gaussian distribution around the bias-corrected ensemble mean. The spread of this Gaussian is on the one hand linearly adjusted according to the errors of the regression model using the training data, and on the other hand depends on the actual spread according to the diagnosed spread-error relationship in the training dataset. Thus, one important feature of this methodology is being able to not only correct the first moment of the ensemble distribution but also correct spread deficiencies.

After applying the NGR calibration, the forecast Probability Density Function (PDF) consists of a continuous Gaussian distribution, not an ensemble of realizations. However, in order to be able to compare the performance of the calibrated probabilities with the frequentist probabilities based on individual ensemble members, a synthetic ensemble is created from the calibrated Gaussian by drawing 51 equally likely ensemble members from the calibrated PDF. That is, the synthetic ensemble is realized by sampling the members at the 51 equally spaced quantiles of the regressed Cumulative Distribution Function (CDF).

Experimenting with the choice of training dataset and calibration method revealed that combining the simple bias-correction based on the 30-day training data (BC-30) and the NGR calibration based on reforecasts (NGR-RF) is superior to the pure NGR-RF calibration, in particular for early lead times.

As already seen in Fig. 1c, the 30-day bias correction can improve the CRPS of the DMO by about 20% for early lead times. The reforecast based NGR calibration is even more effective, with improvements of more than 25% (Fig. 2).

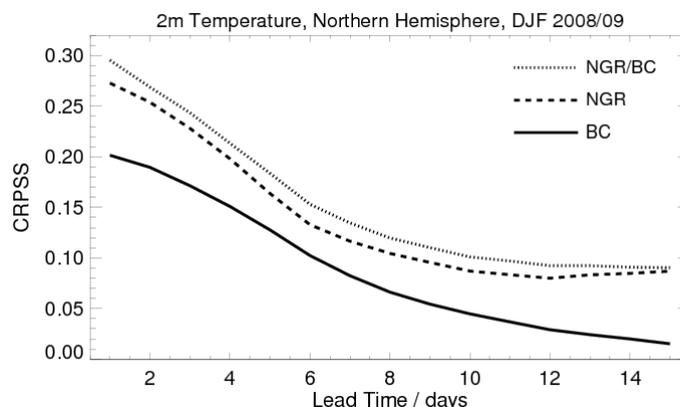


Figure 2: Illustration of gain in skill depending on the calibration method applied to ECMWF direct model output. The CRPSS is defined as  $CRPSS = 1 - CRPS(exp) / CRPS(ref)$ , with  $CRPS(exp)$  being the CRPS of the bias corrected forecasts (solid), the NGR calibrated forecast (dashed), and the NGR/BC model combination (dotted).  $CRPS(ref)$  is in all cases the CRPS of the DMO forecasts. All scores are calculated for 2-m temperature forecasts starting in DJF 2008/09 and averaged over the Northern Hemisphere ( $20^{\circ}N - 90^{\circ}N$ ).

However, combining the NGR-RF and BC-30 ensembles can lead to further slight improvements. The two ensembles are not combined by taking all members from both ensembles to form a new ensemble with twice the number of members, but by first ordering both the bias-corrected and NGR-calibrated ensembles and then averaging the corresponding members. In this way the final combined calibrated system still contains only 51 members. Some experimentation with different weights for the NGR-RF and BC-30 ensembles revealed that applying equal weights at all lead times leads to overall best results. For the current version, the slightly improved performance might be caused by the fact that the BC-30 calibration contains information on the bias more relevant to the current weather regime than the overall bias diagnosed from the reforecast dataset. However, using a refined version of the NGR-RF calibration by, for example, including soil moisture as an additional predictor might diminish the positive impact the BC-30 contribution can have. A further advantage of adding the BC-30 calibrated ensemble to the Gaussian NGR-RF ensemble is the fact that through this procedure any non-Gaussian characteristics of the original ensemble may be retained to some degree.

## 4. Results

### 4.1. TIGGE multi-model ensembles versus single-model systems

A first impression on the level of skill of the single-model systems is given by comparing the CRPSS of the 850-hPa temperature over the Northern Hemisphere for forecasts of the winter season DJF 2008/09 (Fig. 3a). The scores are based on uncalibrated DMO ensembles since (i) applying the bias-correction has no significant impact for upper-air variables like T850 and (ii) we do not need to account for the choice of verification dataset in the case of T850, as demonstrated in section 2.3. The performance of the T850 forecasts varies significantly for the different models, with the CRPSS dropping to zero for the worst models at a lead time of five days and for the best models around day 15. That is, the time range up to which the model predictions are more useful than the reference forecast, which is in this case the climatological distribution, changes considerably from one model to

another. The climatological distribution is estimated from ERA-40 reanalyses in the period 1979–2001 (Uppala et al., 2005; Jung and Leutbecher, 2008). Since not all forecasting centres integrate their models to 15 days lead time, the performance of the multi-model ensemble combining all nine single-model systems can only be assessed up to the maximum forecast range covered by all individual models, which is nine days. Except for the first two forecast days, this multi-model prediction does not significantly improve over the best single-model, i.e. the ECMWF EPS. Note that the significance levels of the difference between the single-model systems and the multi-model ensemble have been assessed using a paired block bootstrap algorithm following Hamill (1999). Similar results can be observed for other variables like e.g. the bias-corrected 2-m temperature (Fig. 3b). Again, the best model is only for the first two to three days significantly worse than the multi-model ensemble, and their performance cannot be distinguished later on.

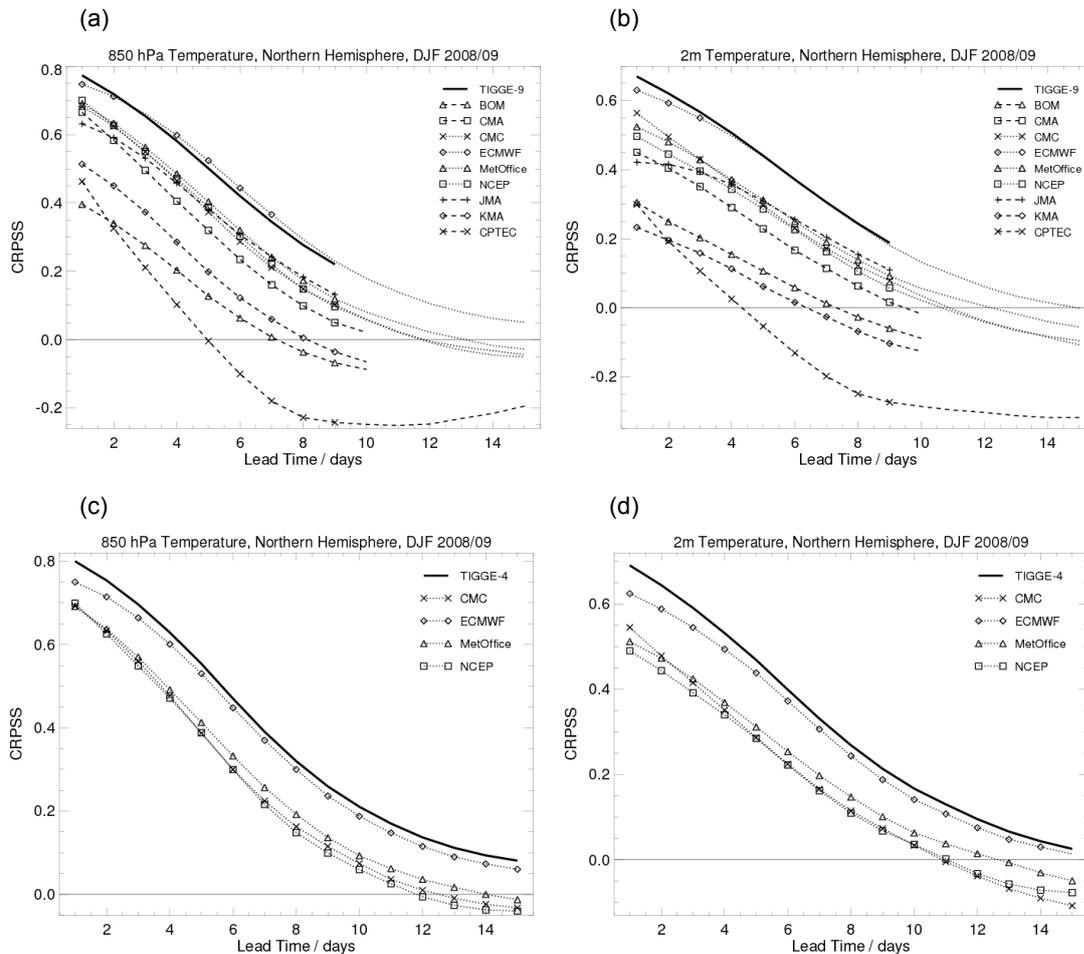


Figure 3: Continuous Ranked Probability Skill Score versus lead time for 850-hPa temperature DMO forecasts (left column, subpanel a and c) and 2-m temperature BC-30 forecasts (right column, subpanel b and d) in DJF 2008/09, averaged over the Northern Hemisphere ( $20^{\circ}\text{N} - 90^{\circ}\text{N}$ ). Figures (a) and (b) in the upper row show results for the TIGGE-9 multi-model (solid line) composed of nine single-models and the scores of all nine contributing single-models (dashed and dotted lines with symbols). Figures (c) and (d) in the lower row show results for the TIGGE-4 multi-model (solid line) composed of the four best single-models with lead time up to 15 days and the scores of the four contributing single-models (dotted lines with symbols). Symbols are only plotted for cases in which the single-model score significantly differs from the multi-model score on a 1% significance level.

The inability of the multi-model ensemble to significantly improve over the best single-model system might be caused by the fact that it consists of all nine single-models, i.e. it includes also the models with rather poor performance. In order to eliminate these possibly detrimental contributions, a new multi-model (TIGGE-4) containing only the four best single-model systems with lead time up to 15 days was constructed and compared to the four contributing single-models: CMC, ECMWF, MetOffice, and NCEP (Fig 3c and 3d). In fact, this reduced version of the full multi-model ensemble gives now significantly improved scores over the whole forecast period and for both upper-air and surface variables. This result indicates that a careful selection of the contributing models seems to be important for medium-range multi-model predictions.

#### 4.2. TIGGE multi-model ensemble versus reforecast-calibrated ECMWF

After having established a new benchmark for the best single-model, the ECMWF EPS, the question is now whether it might be possible to achieve similar improvements by calibrating the ECMWF EPS based on its reforecast dataset. Comparing the CRPSS of the reforecast-calibrated ECMWF EPS (EC-CAL) with the TIGGE-4 multi-model scores previously shown in Figure 3 reveals that indeed the calibration procedure significantly improves ECMWF's scores (Fig. 4).

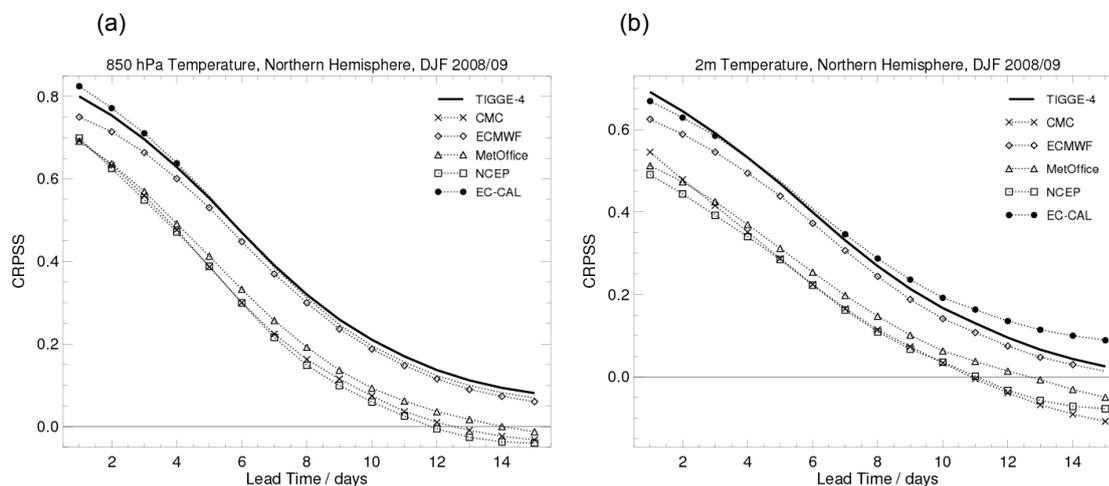


Figure 4: Continuous Ranked Probability Skill Score versus lead time for the TIGGE-4 multi-model (solid line), for the contributing single-models itself (dotted lines with symbols, CMC: cross, ECMWF: diamond, MetOffice: triangle, NCEP: square), and for the reforecast calibrated ECMWF forecasts (dotted lines with bullets). Symbols are only plotted for cases in which the single-model score significantly differs from the multi-model score on a 1% significance level. (a) 850-hPa temperature DMO and EC-CAL forecast scores averaged over the Northern Hemisphere (20°N - 90°N) for DJF 2008/09, (b) as in (a) but for 2-m temperature BC-30 and EC-CAL forecast scores.

Overall the performance of the EC-CAL predictions is as good as the TIGGE-4 multi-model ensemble, and for longer lead times it can be even better. For 850-hPa temperature predictions the EC-CAL's CRPSS lies above the multi-model CRPSS for early lead times, and for longer lead times the skill scores are slightly lower than for the multi-model ensemble, though not statistically significant (Fig 4.a). Considering the slight advantage in the early lead times for ECMWF forecasts when using ERA-interim as verification and the lack of statistical significance of the difference in the CRPSS for longer lead times, it can be concluded that for T850 the reforecast-calibrated ECMWF forecasts are of comparable quality as the TIGGE-4 multi-model forecasts.

This result is confirmed also for 2-m temperature forecasts, for which the calibration is even more effective for longer lead times (Fig. 4b). This indicates that the systematic component of the error is more dominant in the case of 2-m temperature, and thus the calibration procedure is able to further reduce the Root Mean Square Error (RMSE) of the ensemble mean. However, the general level of skill at those long lead times is very low. Therefore, these improvements - as relevant as they might look in terms of overall scores - might not add very much in terms of improving the usefulness of the predictions in a real forecast situation.

In order to indeed further investigate the mechanisms behind the improvements, Figure 5 shows the RMSE of the ensemble mean and the spread of the different ensembles.

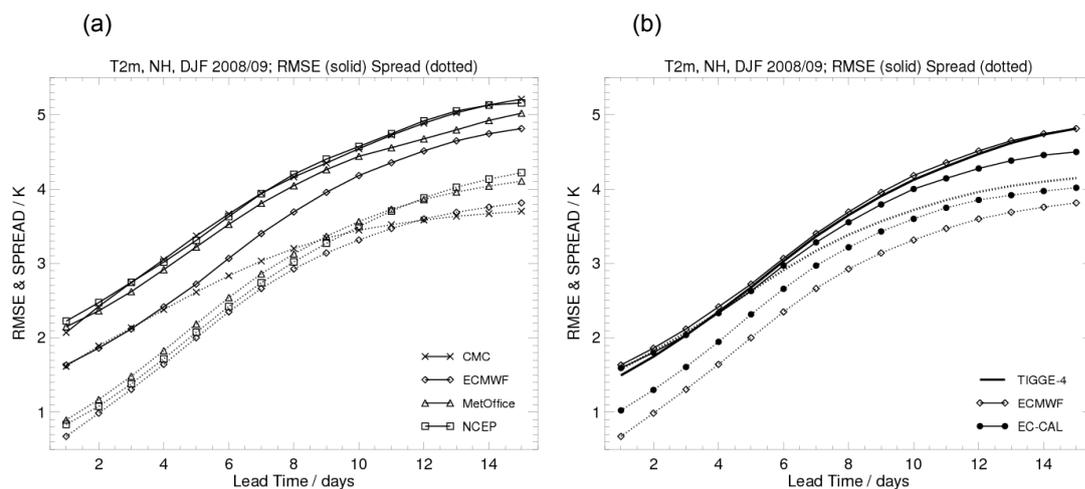


Figure 5: Root Mean Square Error of the ensemble mean (solid lines) and ensemble standard deviation (“spread”, dotted lines) versus lead time. a: results for the single-model BC-30 forecasts depicted using lines with symbols (CMC: cross, ECMWF: diamond, MetOffice: triangle, NCEP: square). b: as in (a) but without the CMC, MetOffice and NCEP results, including instead the results for the reforecast calibrated ECMWF (lines with bullets as symbol) and TIGGE-4 multi-model results (curves without symbols). All scores are calculated for 2-m temperature forecasts starting in DJF 2008/09 and averaged over the Northern Hemisphere (20°N - 90°N).

Ensemble forecasting aims to construct uncertainty information so that the observations can be considered as indistinguishable from the ensemble members of the forecast. Since both the ensemble mean and the analysis have an error, a necessary (but not sufficient) condition for a reliable ensemble is for the sum of the squared spread of the ensemble and the variance of the analysis error to be close to the squared difference of the ensemble mean and the analysis (Saetra et al., 2004, Candille et al., 2007). It is difficult to quantitatively estimate the true analysis error variance, however, it is planned to extend the current diagnostic to incorporate this aspect in future work (e.g., multi-model estimates of Langland et al. 2008, or direct estimates via the ensemble Kalman filter, e.g., Houtekamer and Mitchell 1998). Until this more quantitative assessment, we just qualitatively state that the spread of the ensemble should be somewhat smaller than the standard deviation of the ensemble mean, especially at short forecast leads when analysis error is of similar magnitude to spread and error. However, for 2-m temperature all single-model systems are seriously under-dispersive (Fig. 5a), i.e. much more than could be explained by the analysis error variance currently not accounted for. CMC starts with the smallest spread deficiency at the beginning of the forecast, but due to a serious mismatch in the growth of spread and error it has the worst spread-error relation for longer lead times. The remaining three models have a similar level of spread, however, the significantly lower RMSE of the ECMWF EPS implies not only a slightly better spread error relationship compared to the

MetOffice and NCEP ensembles, it is also one of the main reasons for its significantly better probabilistic scores discussed before. The effect of combining the single-model systems or calibrating the ECMWF EPS can be seen in Figure 5b. The RMSE of the multi-model ensemble is slightly reduced for early lead times, but the most noticeable change is the very much improved spread-error relation in particular up to a forecast range of day-6. In contrast to that, the reforecast-calibrated ECMWF EPS has not such a perfect spread-error relation, though it is improved compared to the original EPS spread. The reason for this is the above discussed methodology of combining the BC-30 and NGR-RF calibration. Applying the pure NGR calibration should lead to a near perfect spread-error relation, but as discussed above, the advantages of possible reductions in the systematic error provided by the 30-day bias-corrected ensemble may outweigh the slight disadvantage of a supposedly poorer 2nd-moment calibration. Since the under-dispersion is not fully corrected in the reforecast-calibrated ensemble, the main improvement of its probabilistic scores comes from the reduction in the RMSE, in particular for longer lead times.

### 4.3. TIGGE weighted multi-model ensemble

After having established that overall the performance of the TIGGE multi-model and the reforecast-calibrated ECMWF EPS are at a comparable level, the question remains whether the multi-model performance can be further improved by applying different weights to its individual components. Generally, only very moderate improvements are achieved by the two tested weighting methods which determine the weights based on a training dataset of forecast-observation pairs from the previous 30 days, WGT-1 and WGT-2 (Fig. 6). The weighting method based on the inverse mean square error (WGT-1) starts with improvements of about 3% for early lead times, dropping to near zero for longer lead times. The slightly more sophisticated method of finding optimal weights by minimizing the CRPS in the training dataset (WGT-2) achieves slightly higher improvements in the

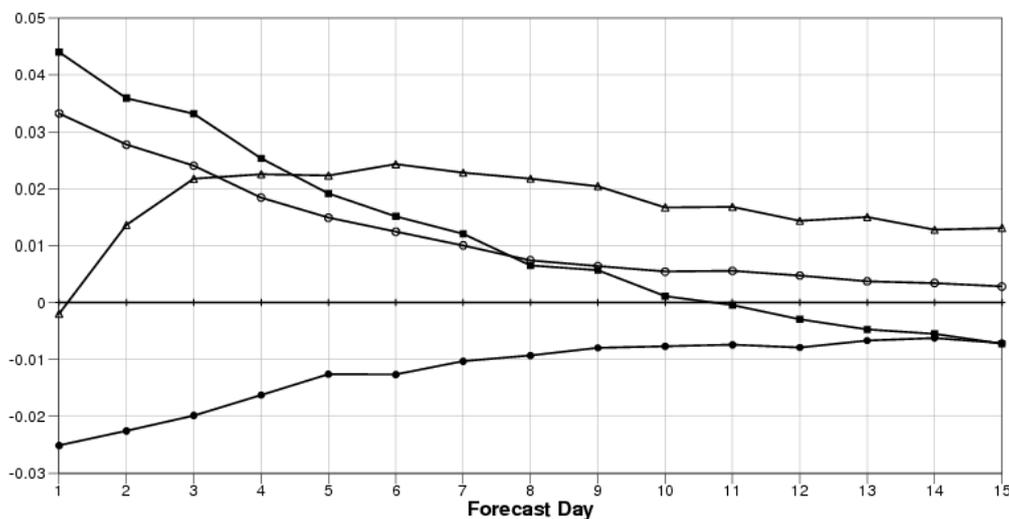


Figure 6: Illustration of impact on skill depending on the weighting method applied to the TIGGE-4 multi-model ensemble. The CRPSS is defined as  $CRPSS = 1 - CRPS(exp) / CRPS(ref)$ , with  $CRPS(ref)$  being in all cases the CRPS of the equally weighted TIGGE-4 multi-model ensemble, and  $CRPS(exp)$  being the CRPS of the different weighting methods. The four lines with different symbols represent the percentage increase/decrease of the skill for the weighting method based on the inverse MSE (WGT-1, open circles), the CRPS minimization method (WGT-2, squares), random weights (WGT-3, solid circles) and fixed weights (WGT-4, open triangles). All scores are calculated for 2-m temperature forecasts starting in MAM 2010 and averaged over the Northern Hemisphere ( $20^{\circ}N - 90^{\circ}N$ ).

early forecast range, but falls back to even negative values for longer forecast ranges. These results indicate that applying different weights to the individual single-model ensembles does not necessarily lead to significant improvements compared to the equally weighted multi-model ensemble.

In order to investigate more generally what level of impact any weighting scheme could have on the performance of the multi-model, a deliberately “bad” weighting scheme, i.e. giving random weights (WGT-3), was applied. As expected, applying random weights leads to a worse performance of the multi-model. More interestingly, the level of deterioration is in the same order of magnitude as could be achieved with the first two weighting schemes. This confirms that generally any weighting of the four individual model components has only a limited effect on the overall performance of the predictions. The reason for this is the overall similarity of the four single-model ensembles, resulting in only minor changes of the multi-model ensemble when giving different weights to the single-model components. Finally, we tested the impact of applying a combination of weights based on more intuitive reasoning (WGT-4). Interestingly, the results are in the same order as the improvements from the more objectively derived weights. This demonstrates that employing more sophisticated weighting schemes, which can be computationally expensive, might not be worthwhile in this context.

## 5. Summary and conclusions

The main aim of this study was to investigate whether the multi-model concept is similarly successful for medium-range forecasting as it has been proven in the seasonal forecast range. This has been done by comparing the relative benefits of TIGGE multi-model forecasts versus the reforecast-calibrated ECMWF EPS. A major issue in such a verification study was the choice of verification dataset. ERA-Interim was chosen here because: (i) one intent of this study was to evaluate reforecast-calibrated ECMWF products (requiring forecasts and observations or analyses from many years past); (ii) use of analyses rather than observations facilitated a verification over a larger region, including oceans, and (iii) the highest-quality analyses were deemed preferable. Since the choice of ERA-Interim as verification data set may favour ECMWF forecasts, we demonstrated the relative impact of using different reanalyses as verification. It was shown that the level of advantage the ECMWF EPS gained when using ERA-interim as verification for surface variables was reduced by applying a simple bias-correction to all models using the past 30 days of forecasts and analyses.

The performance of nine single-model systems from the TIGGE archive was compared with the performance of the full TIGGE multi-model, consisting of all these nine models. This full multi-model version did not improve on the best single-model, the ECMWF EPS. However, when combining only the four best single-model ensembles (CMC, ECMWF, MetOffice, and NCEP), the multi-model ensemble outperformed the ECMWF-only EPS forecasts, though we note that this result did not apply to all model variables and all lead times. However, by taking advantage of the reforecast dataset which was available for the ECMWF EPS and using it as training dataset to produce reforecast-calibrated forecasts, the ECMWF EPS scores were improved to such an extent that its overall performance was as good as the TIGGE multi-model system, and often better.

The reforecast calibration procedure was particularly helpful at locations with clearly detectable systematic errors like areas with complex orography or coastal grid points. In such areas the calibration procedure essentially applied a statistical downscaling to the forecasts. The multi-model approach, in contrast, might be advantageous in situations where it is able to suggest alternative solutions not predicted by the single-model of choice. Further investigations on the mechanisms behind the improvements achieved by the post-processing methods led to the conclusion that both

approaches tend to correct similar deficiencies. That is, systematic error and spread deficiencies were improved to a similar extent by both approaches. Experiments assessing the impact of giving different weights to the individual single-model ensembles demonstrated that only minor improvements could be achieved. This limited potential of applying sophisticated weighting schemes to the multi-model ensemble is caused by (i) the relative similarity of the single-model ensembles and (ii) the fact that the relative errors are not very stable in time and thus it is difficult to determine robust and stable weights.

Finally, considering the performance improvements made possible by the availability of the ECMWF reforecast dataset, other modelling centres might start providing reforecasts for their model systems in the not so distant future. In that case it would be interesting to study the relative benefits achievable for reforecast-calibrated multi-model or single-model systems, respectively. Furthermore, we suggest exploring the relative merits of multi-model versus reforecast-calibrated predictions for other user-relevant variables like precipitation and wind speed, in particular in the context of extreme events.

### Acknowledgments

I would like to thank all members of the Probabilistic Forecasting Division at ECMWF for the intense and fruitful discussions in the context of this work, not forgetting the kind and efficient support I received in all scientific and technical aspects. My special thanks goes to Tom Hamill for the excellent collaboration on the reforecast-related work; without Tom we might not have operational reforecasts at ECMWF!

### References

- Bougeault, P. and Coauthors, 2010: The THORPEX Interactive Grand Global Ensemble (TIGGE). *Bull. Amer. Meteor. Soc.*, in press.
- Candille, G., C. Côté, P.L. Houtekamer, and G. Pellerin, 2007: Verification of an Ensemble Prediction System against Observations. *Mon. Wea. Rev.*, **135**, 2688-2699.
- Doblas-Reyes, F. J., R. Hagedorn and T.N. Palmer, 2005: The rationale behind the success of multi-model ensembles in seasonal forecasting. Part II: Calibration and combination. *Tellus-A*, **57**, 234-252.
- Gneiting, T., A. E. Raftery, A. H. Westveld III, and T. Goldman, 2005: Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Mon. Wea. Rev.*, **133**, 1098-1118.
- Hagedorn, R., 2008: Using the ECMWF reforecast dataset to calibrate EPS forecasts. *ECMWF Newsletter*, **117**, 8-13.
- Hagedorn, R., F. J. Doblas-Reyes and T. N. Palmer, 2005: The rationale behind the success of multi-model ensembles in seasonal forecasting. Part I: Basic concept. *Tellus-A*, **57**, 219-233.
- Hagedorn, R., T. M. Hamill, and J. S. Whitaker, 2008: Probabilistic forecast calibration using ECMWF and GFS ensemble reforecasts. Part I: two-meter temperatures. *Mon. Wea. Rev.*, **136**, 2608–2619.
- Hamill, T. M., 1999: Hypothesis tests for evaluating numerical precipitation forecasts. *Wea. Forecasting*, **14**, 155-167.
- Hamill, T. M., and J. S. Whitaker, 2006: Probabilistic quantitative precipitation forecasts based on reforecast analogs: theory and application. *Mon. Wea. Rev.*, **134**, 3209-3229.

- Hamill, T. M., and J. S. Whitaker, 2007: Ensemble calibration of 500-hPa geopotential height and 850-hPa and 2-m temperatures using reforecasts. *Mon. Wea. Rev.*, **135**, 3273-3280.
- Hamill, T. M., J. S. Whitaker, and X. Wei, 2004: Ensemble reforecasting: improving medium-range forecast skill using retrospective forecasts. *Mon. Wea. Rev.*, **132**, 1434-1447.
- Hamill, T. M., J. S. Whitaker, and S. L. Mullen, 2006: Reforecasts, an important dataset for improving weather predictions. *Bull. Amer. Meteor. Soc.*, **87**, 33-46.
- Hamill, T. M., R. Hagedorn, and J. S. Whitaker, 2008: Probabilistic forecast calibration using ECMWF and GFS ensemble forecasts. Part II: Precipitation. *Mon. Wea. Rev.*, **136**, 2620–2632.
- Hersbach, H., 2000: Decomposition of the Continuous Ranked Probability Score for Ensemble Prediction Systems. *Wea. Forecasting*, **15**, 559–570.
- Houtekamer, P. L., and H. L. Mitchell, 1998: Data assimilation using an ensemble Kalman filter technique. *Mon. Wea. Rev.*, **126**, 796-811.
- Johnson, C., and R. Swinbank, 2009: Medium-Range multi-model ensemble combination and calibration. *Quart. J. Roy. Meteor. Soc.*, **135**, 777-794.
- Jung, T., and M. Leutbecher, 2008: Scale-dependent verification of ensemble forecasts. *Quart. J. Roy. Meteor. Soc.*, **134**, 973-984.
- Kalnay, E., and coauthors, 1996: The NCEP/NCAR 40-year reanalysis project. *Bull. Amer. Meteor. Soc.*, **77**, 437-471.
- Kanamitsu, M., W. Ebisuzaki, J. Woollen, S.-K. Yang, J. J. Hnilo, M. Fiorino, and G. L. Potter, 2002: NCEP-DOE AMIP-II Reanalysis (R-2), *Bull. Amer. Meteor. Soc.*, **83**, 1631-1643.
- Krishnamurti, T. N., C. M. Kishtawal, T. E. LaRow, D. R. Bachiochi, Z. Zhang, C. E. Williford, S. Gadgil, S. Surendran, 1999: Improved Weather and Seasonal Climate Forecasts from Multimodel Superensemble. *Science*, **285**, 1548-1550.
- Langland, R. H., Maue., R. N., and C. H. Bishop, 2008: Uncertainty in atmospheric temperature analyses. *Tellus A*, **60**, 598-603.
- Nelder, J. A., and R. Mead, 1965: A simplex method for function minimization. *Computer Journal*, **7**, 308-313.
- Palmer, T.N. and Coauthors, 2004: Development of a European Multi-Model Ensemble System for Seasonal to Inter-Annual Prediction (DEMETER). *Bull. Amer. Meteor. Soc.*, **85**, 853–872.
- Park, Y.-Y., R. Buizza, and M. Leutbecher, 2008: TIGGE: preliminary results on comparing and combining ensembles. *Quart. J. Roy. Meteor. Soc.*, **134**, 2029-2050.
- Pavan, V., and F. J. Doblas-Reyes, 2000: Multi-model seasonal hindcasts over the Euro-Atlantic: skill scores and dynamic features. *Climate Dyn.*, **16**, 611-625.
- Peng, P., A. Kumar, H. van den Dool, and A. G. Barnston, 2002: An analysis of multimodel ensemble predictions for seasonal climate anomalies, *J. Geophys. Res.*, **107(D23)**, 4710, doi:10.1029/2002JD002712.
- Robertson, A.W., U. Lall, S. E. Zebiak, L. Goddard, 2004: Improved combination of multiple atmospheric GCM ensembles for seasonal prediction. *Mon. Wea. Rev.*, **132**, 2732–2744.

- Saetra, Ø., H. Hersbach, J.-R. Bidlot, and D.S. Richardson, 2004: Effects of Observation Errors on the Statistics for Ensemble Spread and Reliability. *Mon. Wea. Rev.*, **132**, 1487-1501.
- Shin, D. W., S. Cocks, and T. E. Larow, 2003: Ensemble Configurations for Typhoon Precipitation Forecasts. *J. Meteor. Soc. Japan.*, **81** (4), 679-696.
- Simmons, A., Uppala, S., Dee, D., and Kobayashi, S., 2007: ERA-Interim: New ECMWF reanalysis products from 1989 onwards. *ECMWF Newsletter*, **110**, 25-35.
- Uppala, S. M. and Coauthors, 2005: The ERA-40 re-analysis. *Quart. J. Roy. Meteor. Soc.*, **131**, 2961-3012.
- Weigel, A. P., and N. E. Bowler, 2009: Comment on ‘Can multi-model combination really enhance the prediction skill of probabilistic ensemble forecasts?’. *Quart. J. Roy. Meteor. Soc.*, **135**, 535-539.
- Weigel, A. P., M. A. Liniger, C. Appenzeller, 2008: Can multi-model combination really enhance the prediction skill of probabilistic ensemble forecasts? *Quart. J. Roy. Meteor. Soc.*, **134**, 241-260.
- Weigel, A.P., M.A. Liniger, and C. Appenzeller, 2009: Seasonal Ensemble Forecasts: Are Recalibrated Single Models Better than Multimodels? *Mon. Wea. Rev.*, **137**, 1460–1479.
- Wilks, D. S., and T. M. Hamill, 2007: Comparison of ensemble-MOS methods using GFS reforecasts. *Mon. Wea. Rev.*, **135**, 2379-2390.

