J1.8 DOES INCREASED MODEL RESOLUTION ENHANCE PREDICTABILITY?

Zoltan Toth[1], Yuejian Zhu[1], Istvan Szunyogh[2], Mark Iredell, and Richard Wobus[1]
Environmental Modeling Center, NCEP, NWS/NOAA
Washington DC 20233

1. INTRODUCTION

In this paper we explore the effect of horizontal resolution on forecast skill and predictability in general as manifested in forecasts started from the best available analyses (control forecasts) and from slightly perturbed states around them (ensemble forecasts) at the National Center for Environmental Prediction (NCEP). Confirming many earlier studies (e. g., Simmons and Hollingworth 2001), Szunyogh and Toth (2001) found that increased horizontal resolution had a clear positive impact on the quality of short range Numerical Weather Predictions (NWP). They also found that truncating the T126 high resolution control and ensemble forecasts after *one day* of integration to a lower T62 resolution had a clear negative effect on forecast performance. The high resolution ensemble forecasts truncated after 3 days also showed a clear degradation. Based on a 30–day experimental period, however, Szunyogh and Toth (2001) reported that the high resolution control forecast truncated at *three day* lead time performed better than forecasts run with the more realistic high resolution (T126) model version all the way to 7 days.

After reviewing the related results of Szunyogh and Toth (2001) the impact of truncating control forecasts after a few days of integration will be studied on an expanded set of forecasts (section 2). The effect of horizontal truncation on the quality of ensemble forecasts (both in terms of ensemble mean and probabilistic forecast performance) will be evaluated in the next section (section 3). In search for an explanation of the results, in section 4 predictability properties of the different model versions will be explored in a perfect model environment. A tentative explanation for the effect of model truncation on control and ensemble forecast performance will be put forward in section 5 while section 6 offers some preliminary conclusions.

2. PERFORMANCE OF TRUNCATED CONTROL FORECASTS

In this section we use three different data sets to explore the effect of horizontal truncation introduced in the course of control forecast integrations on forecast performance.

2.1 *Results from an earlier study*

First we review related results presented by Szunyogh and Toth (2001) for the period January 13 through February 11 1999. In this study the January 2000 operational version of the NCEP MRF forecast model (Derber et al. 1998) was used to generate T126 control forecasts. Here we consider two other forecast setups studied by Szunyogh et al: one where the forecasts are truncated to T62 after 3 days of integration, and another where the forecasts start at the lower T62 resolution. Root mean square (rms) error results presented in Fig. 2 of Szunyogh and Toth (2001) are reproduced here for convenience in Fig. 1.

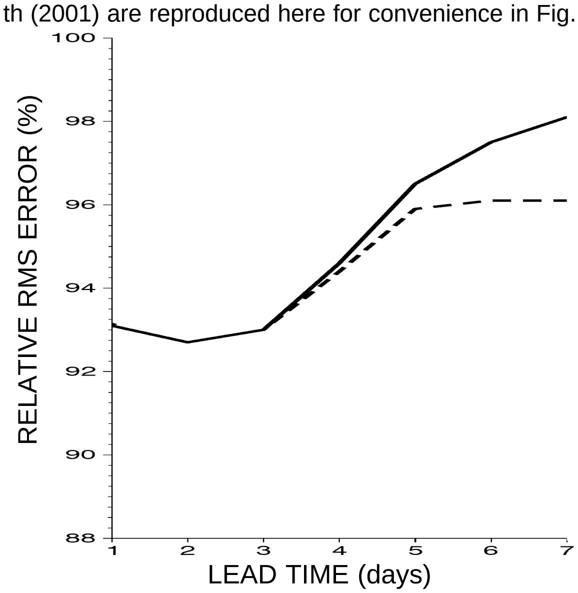

Fig. 1. 500 hPa geopotential height RMS error for the T126 control forecast run at full resolution (solid line) and truncated to T62 resolution after 3 day lead time (dashed), expressed as a percentage of error in the low resolution T62 forecast for the NH extatropics, 13 Jan. – 11 Febr. 1999. Adapted from Szunyogh and Toth (2001).

The figure shows that increased horizontal resolution during the first 3 days of integration reduces Northern Hemisphere extratropical (30N–70N) 500 hPa geopotential height rms forecast error. Interestingly, maintaining the high resolution beyond three days lead time, however, slightly degrades forecast performance. These and all other results reported later are confirmed using pattern anomaly correlation (not shown) as another measure of forecast performance. Note also that all forecasts in this paper are verified on a common regular 2.5 degree latitude/longitude grid.

2.2 *Follow up experiments*

Since Szunyogh and Toth (2001) limited their study to a 30–day period, and the truncation of the control forecasts after three days resulted only in a modest improvement in skill, follow up experiments were carried out to confirm the results. In these tests the NCEP MRF T170 horizontal, 42 level (L42) vertical resolution operational forecasts were truncated at 84 hour lead time and continued at T62 horizontal, and 28 level (L28) vertical resolution with the September 2001 operational version of the NCEP MRF model. In Fig. 2 500 hPa geopotential height forecast error results (solid lines) are averaged for January (only for days 12–31

1 SAIC at NCEP. Corresponding author address: Z. Toth, NCEP/EMC, 5200 Auth Rd., Room 207, Camp Springs, MD 20746. Zoltan.Toth@noaa.gov

2 University of Maryland, Institute for Physical Science and Technology

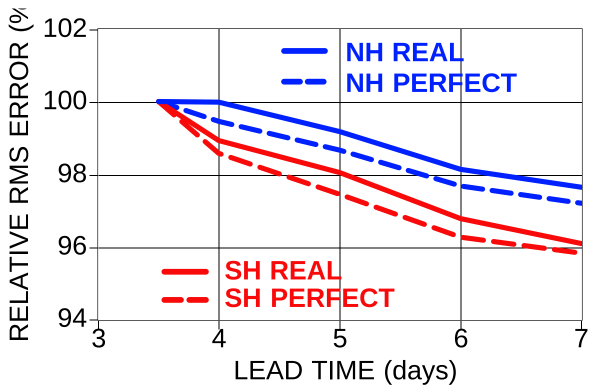

Fig. 2. 500 hPa geopotential height RMS error for the experimental T62/L28 control forecast initiated after truncating the high resolution operational T170/L42 forecast at 84 hour lead time, expressed as a percentage of error in the high resolution forecast, for NH and SH extatropics (blue and red continuous lines), averaged for January, February, and July of 2000. Dashed lines correspond to perfect model evaluation results where the forecasts are verified against the following day's T170/L42 forecast.

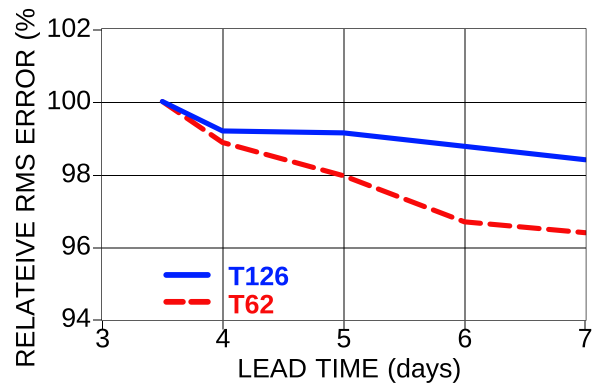

Fig. 3. 500 hPa geopotential height RMS error for experimental T62/L28 (red) and T126/L28 (blue)control forecasts initiated after truncating the high resolution operational T170/L42 forecasts at 84 hour lead time, expressed as a percentage of error in the high resolution forecast, averaged for the NH and SH extatropics, for February 2001 (continuous lines). Dashed lines correspond to perfect model evaluation results where the forecasts are verified against the following day's T170/L42 forecast.

since the operational configuration was changed on 12 January 2001), February, and July of 2001, for the NH and SH extratropics (30S – 70S). The results indicate that the forecasts truncated at 84 hour lead time performed better than the forecasts that maintained high resolution for the entire forecast period. The truncated forecast outperformed the high resolution model on 57/66 days out of the 79 cases considered for the NH/SH extratropics respectively (not shown). Assuming that data from only every fifth day is independent these results are statistically significant at the 5%/1% levels respectively. When data from the two hemispheres are combined and considered independent the results are statistically significant at the 0.1% level.

For the month of February 2001 another set of experiments was carried out in which the forecasts were truncated to an intermediate T126/L28 resolution (instead of T62/L28). To reduce sampling fluctuations, combined monthly average scores are displayed for the NH and SH extratropics in Fig. 3 (solid lines). As seen from Fig. 3 the modestly truncated forecast still outperforms the T170/L42 high resolution forecast but to a lesser degree than the forecasts truncated to T62/L28. When the 7–day lead time daily scores for the NH and SH extratropics are combined and only every fifth day is considered independent the T126/L28 model is found to yield poorer forecasts than the T62/L28 version at the 10% statistical significance level (43 wins out of 56 cases). The advantage of the T126/L28 against the T170/L42 model version (34 wins) is statistically not significant.

2.3 *Evaluation of operational forecasts*

Results in Figs. 1–3 pertain to experiments carried out over 4 months in 3 separate seasons. They indicate that forecast error growth can be reduced if an initially high resolution forecast is truncated at 72–84 hour lead time and integrated forward at a lower horizontal/vertical resolution. This phenomenon can be further studied by verifying operational forecasts. The operational forecast suite at NCEP includes a T170 horizontal and 42 vertical level resolution (out to 7 days) Medium Range Forecast (MRF, high reolution control) and a lower resolution forecast that starts from the same initial condition truncated to T126 horizontal and 28 level vertical resolution (ensemble control forecast). The latter forecast is further truncated at 60/84 hour lead time (before/after 12 January 2001), after which it is integrated at T62 horizontal resolution.

Based on the results presented in the previous sections one can expect that the higher resolution T170/L42 model version produces superior forecasts during the first few days while the error growth in the truncated forecasts is reduced for the period after the truncation when compared to that in the high resolution forecasts. This is confirmed by Fig. 4 where the 500 hPa geopotential height

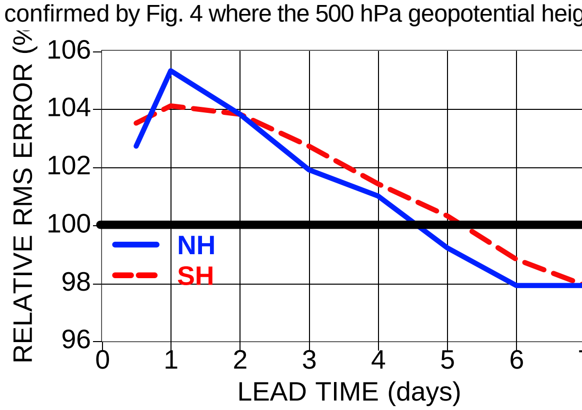

Fig. 4. 500 hPa geopotential height RMS error for the operational ensemble control forecast (T126/L28 until 60/84 hours, then truncated to T62/L28), expressed as a percentage of error in the operational T170/L28 resolution MRF forecast, for NH (solid) and SH extatropics (dashed), averaged for July 2000 through August 2001.

high resolution and truncated ensemble control forecast errors averaged for the available 14 month period between July 2000 and August 2001 are compared.

Note first in Fig. 4 that at 2.5 day (and shorter lead time, before the truncation takes effect) the high resolution model performs better on both hemispheres. Monthly mean scores for the T170/L42 forecasts are better than the slilghly lower resolution T126/L28 forecasts in all 14 months for both hemispheres up to 2.5–day lead time (not shown). These results are significant at the 0.01% level. Second, at longer lead times the truncated version of the model performs better. At 7–day lead time, for example, in 13 out of the 14 months the T62/L28 truncated forecast verifies better than the high resolution T170/L42 forecast over both hemispheres, despite the fact that the low resolution forecasts have 4–5% larger errors earlier on. When the difference between errors at 7 and 3.5 day lead times is considered, the truncated model exhibits a lower error growth in all 14 months investigated, both over the NH and SH extratropics. The results are statistically significant at the 0.01% level for both hemispheres and when they are combined and considered independent, the significance level is 0.0000005%.

3. PERFORMANCE OF TRUNCATED ENSEMBLE FORECASTS

In this section we explore the effect of horizontal truncation on the quality of ensemble forecasts. We will use a 10–member ensemble generated by Szunyogh and Toth (2001) at T126/L28 and T62/L28 resolution using the breeding method operational at NCEP (Toth and Kalnay 1997) for the 30–day period described in section 2.1. In addition, just as in the case of the control forecasts, a third set of ensemble forecasts were also generated where the horizontal resolution is truncated from T126 to T62 at 3–day lead time. First, the performance of the three different sets of ensemble mean forecasts are compared in a fashion similar to that followed in section 2.1 above with respect to the control forecasts (see Fig. 1). Fig. 5, which is

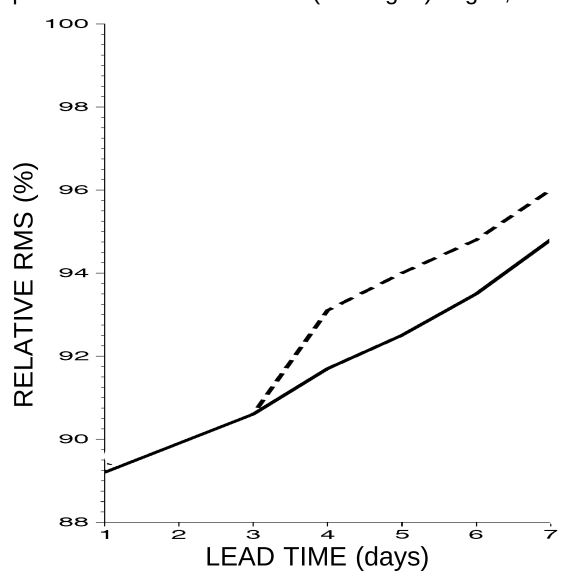

Fig. 5. Same as Fig. 1 except for ensemble mean forecasts. Adapted from Szunyogh and Toth (2001).

a counterpart of Fig. 1, contains the ensemble mean verification results. As expected, the ensemble mean – unlike the control forecast – displays a clear indication of loss of skill when the horizontal resolution is truncated to T62.

Next we investigate a series of measures that are used routinely at NCEP to evaluate the performance of probabilistic forecasts. Each measure is evaluated grid point by grid point and statistics are accumulated over the NH extratropics (except where noted). For a fuller description of the measures considered the reader is referred to Zhu et al. 1996, Toth et al. 1998, and Zhu et al. 2001. Fig. 6 shows how often each set of ensemble forecasts fails to

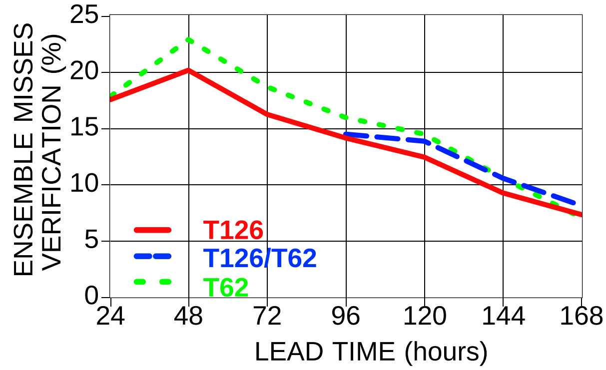

Fig. 6. Percentage of cases when 10–member T126, T62, and initially T126 (that is truncated to T62 resolution at 3–day lead time) ensemble forecasts fail to encompass the 1000 hPa geopotential height verifying analysis over the NH extratropics for January 13 – February 11 1999 (in excess of 18.2% that is expected due to the limited size of the ensemble).

encompass the 1000 hPa geopotential height verifying analysis, above that expected from a perfect ensemble of the same size (10 members). Generally the truncated model results lie in between the high and low resolution curves. Note that at days 5 and 6 the number of outliers for the truncated version fall close to the T62 curve, indicating that in terms of this measure the truncated ensemble suffers a clear loss of skill compared to the high resolution, T126 ensemble.

Fig. 7 displays the Brier Skill Score (BSS) results.

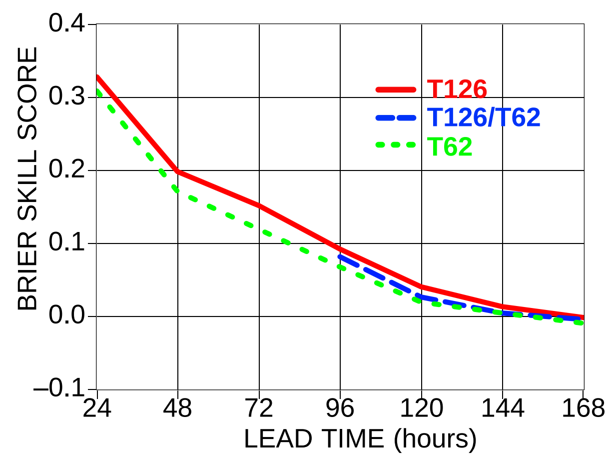

Fig. 7. Brier Skill Score for T126, T62, and initially T126 (that is truncated to T62 resolution at 3–day lead time) 10–member ensemble forecasts for the 500 hPa height, NH extratropics, January 13 – February 11 1999, for 10 climatologically equally likely bins; results shown here are the average for all bins.

Again, the results for the truncated ensemble, as expected, lie in between the T126 and T62 resolution ensembles, with scores close to those of the T62 resolution ensemble at days 5 and 6. The Ranked Probability Skill Score (RPSS) results for the truncated ensemble lie closer to the results of the T62 ensemble than to the T126 ensemble in the 4–6 day lead time range (Fig. 8). The two

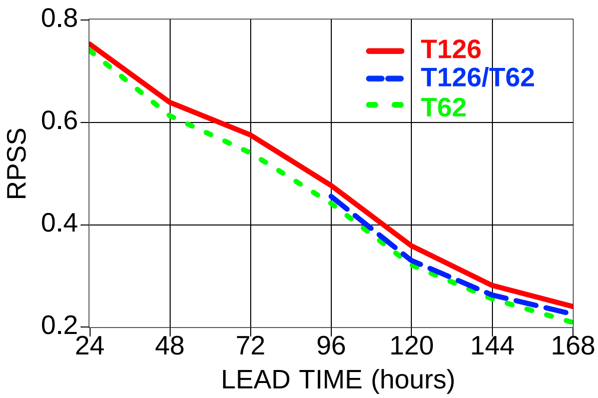

Fig. 8. Same as Fig. 7 except for Ranked Probability Skill Score (RPSS).

measures that are most closely related to the actual use of ensemble based probabilistic forecasts, the Relative Operating Characteristics area (ROC–area) scores (not shown), and the Economic Value results (Fig. 9) both indi-

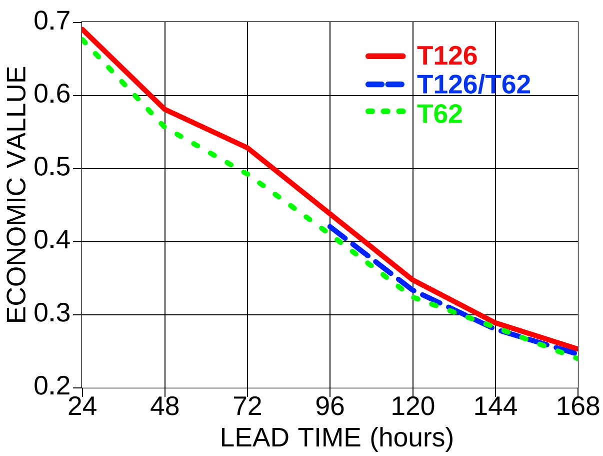

Fig. 9. Same as Fig. 7 except for economic value for users with 0.1 cost–loss ratio.

cate that the performance of the truncated ensemble is closer to that of the low resolution ensemble at all lead times considered (days 4–7) after the truncation is introduced. These results confirm the findings of Szunyogh and Toth (2001) and demonstrate that the performance of the truncated ensemble is poorer than its high resolution counterpart in all respects investigated.

4. CONTROL FORECASTS IN A PERFECT MODEL ENVIRONMENT

The ensemble results presented in the previous section require no special explanation. When a higher resolution model, which is presumably more realistic, is used throughout the whole integration the ensemble performance is superior to an ensemble that is generated, after 3 days lead time, with a reduced resolution version of the model. The improved performance of the truncated control forecast discussed in section 2, however, is somewhat unexpected and calls for an explanation. The fact that the truncated ensemble forecasts exhibit reduced skill is an indication that the peculiar behavior of the control forecasts may not be related to model imperfection. In this section we investigate this issue by evaluating the same set of forecasts studied in section 2.2 but now in a perfect model environment.

Following Lorenz (1982) we will consider a numerical forecast started from a standard analysis as truth, and consider a numerical integration with the same (or a different) model started from an analysis a day earlier as a forecast for truth. The 1–day forecast error, defined in the traditional sense as the difference between the numerical forecast and its verifying analysis takes the role of “initial error” in the perfect model setup. “Verification” in this setup involves the comparison of the day old forecast with the following day forecast (truth). In the following we will consider the forecast data set for January, February and July 2001, introduced in section 2.2. The T170/L42 forecasts will play the role of “truth” and we will verify integrations started a day earlier with both the same T170/L42 model version (perfect model) and the same model but truncated at 84 hour lead time to T62/L28 (less realistic, imperfect model).

The verification results for the perfect and truncated models in the perfect model scenario are presented as dashed lines in Fig. 2. Interestingly, the dashed perfect model verification lines, to a good approximation, run parallel to the solid verification lines based on a regular comparison of forecasts against corresponding analysis fields (section 2.2). Forecasts made by the imperfect (less realistic), truncated model version verify better than those made with the perfect (more realistic) high resolution model in all three months (not shown). A statistical analysis of the daily results (not shown), where data only from every fifth day is considered independent, indicate that the results for the NH/SH extratropics (truncated T62 forecasts better in 54/60 out of 79 total number of cases) are statistically significant at the 5%/10% level.

Similar results were obtained with the operational data set discussed in section 2.3. These results are consistent with those presented by Simmons and Hollingworth (2001) who found that error growth over the NH extratropics is enhanced in more recent (and more realistic) versions of the ECMWF operational forecast model, especially during summer (see their figure 7).

5. DISCUSSION

5.1 *Explanation of results*

As noted earlier, the fact that truncation did not improve but rather degraded ensemble performance (section 3) indicates that the cause of the improved control forecast performance after truncation cannot be explained by model imperfections. The reproduction of the results in controled perfect and truncated model experiments (section 4) gives yet another clear indication that the unexpected behavior of the truncated control forecast is not due to model problems. How can the results then be explained?

It has long been recognized (Lorenz 1969) that smaller scale atmospheric phenomena have progressively shortened predictability time periods. Though some argue that orographic or other external forcing may lend some extra predictability to small scale features such enhanced predictability must be rather confined and limited (see, e. g., Nuss 2001). We argue that it is reasonable to assume that all predictability on scales smaller than those represented in the T62 low resolution version of the truncated model used in this study is lost by 72 hour lead time. In other words at and beyond 3 days lead time features in the T170 or T126 resolution forecasts on scales smaller than T62 do not directly correspond with the observed (or analyzed) features on those scales.

The following question can then be raised. Is it better to carry features in a model that, on their own, have no predictability at all; or is it better to truncate those scales on which predictability is lost out from the model? Potentially, both approaches have advantages. The full resolution model (T170 and T126 in this study) is presumably more realistic. This is indicated by more accurate short range forecasts (see Figs. 1 and 4). But it has the disadvantage of carrying unpredictable features which act as random noise in the model. These features can interact with features on scales that retain some predictability and can potentially negatively influence forecast skill (see, e. g., Tracton 1990).

The exclusion of these unpredictable features (that are beyond its resolution) from the truncated (T62) model has a potential advantage. The predictable scales are not perturbed by noise and may retain their skill for a longer period of time. The obvious negative effect of this approach is that a less realistic model is used. Small scales and their interaction with larger scales are not represented in the truncated model. Instead the effect of the small scales on the larger scales are parameterized by diffusion. Therefore, the truncated model is less realistic in the sense that observed features can be simulated (i.e., forecast at short lead times) with less fidelity than with the higher resolution model.

The noise that a high resolution forecast carries does not affect the skill of the forecast directly. This is because the variance on these small scales is negligable compared to that on larger scales (and because all forecasts are verified on a coarse 2.5 degree grid). The effect of the smaller scales manifests itself only through interactions with the larger, predictable scales. Let us now consider the scales of motion resolved by the truncated T62 model. The noise in a higher resolution model will introduce perturbations onto these larger scales – random perturbations that create alternative forecast trajectories for the real system that is being modeled. The larger scales in the real system (either the actual atmosphere or the high resolution model that represents it in the perfect model environment of section 4) in fact have their own random forcing due to the small scales not represented in the truncated model. When a single high resolution model forecast is verified two trajectories (the forecast and reality) that are both perturbed randomly by the small scales are compared. In the truncated forecast, however, the large scales are *not* perturbed randomly by the small scales. It means that on the predictable scales only one of the two trajectories (the reality) is perturbed, leading to an error smaller than that for the high resolution forecast.

Let us consider now the climatological distribution of the small scales not represented in the T62 model and compare a state in this climatological distribution (small scale motions in ”reality”) to either another randomly chosen state (small scale motions in a high resolution forecast) or to the climatological mean (assumed to be zero in the truncated forecast). It is well known that the expected distance of a randomly chosen point from the mean of a normally distributed sample is $\sqrt{2}$ times smaller than that from another randomly chosen point. One can argue then that the difference in the effect of the small scales on the larger, predictable scales will also be larger when two random realizations of the small scales are present than when only one realization is present. It follows from these arguments that the exclusion of the smaller scales and their associated random forcing may provide an advantage in case of the truncated forecasts.

In the context of the global control forecast configurations considered in this study (T170 or T126 model versions truncated after 3–3.5 day lead time to T62) the positive effect of noise reduction clearly overwhelms the negative effect of less realism/fidelity, leading to improved forecast skill for the truncated forecasts. To what extent the results can be generalized to other (larger or smaller) scales in the atmosphere or to other physical systems is an open question.

5.2 *Implications for NWP forecast evaluation*

The results presented in this study may have some implications for Numerical Weather Prediction (NWP) verification and model development activities. Looking at the control verification results alone a modeler may come to the conclusion that the model truncated after a few days is a better model of the atmosphere. Considering the ensemble results it is obvious that this is not the case. The improved scores for the high resolution ensemble demonstrate that the high resolution model is indeed more realistic, even in the medium range. Each of the alternate scenarios induced by the random perturbations provided by the small scales that are left unresolved in the truncated forecasts but are retained in the high resolution ensemble leads to a somewhat poorer individual forecast (including the control forecast) in terms of its skill, but to a somewhat more realistic trajectory, leading to a more skillful ensemble (Fig. 6). Realism here is considered from a climatological point of view, in terms of simulated features that better resemble those observed over a long period of time. Or more precisely, from a conditional climatological point of view, when only those NWP forecast scenarios are considered that are possible given the observed initial condition with its associated uncertainty.

Note that when even longer, climatological forecasts are considered atmospheric initial conditions already lost their influence and the forecast errors are in a fully nonlinearly saturated mode, given a particular set of projected boundary conditions. A proper evaluation of a model, again, requires a comparison of a “climatological” ensemble generated by one (or more) long integration(s), and an ensemble of observed or analyzed states, both corresponding to a given set of boundary forcing values. Higher resolution models typically demonstrate an advantage in such settings as well (see, e. g., Tibaldi et al. 1990).

It is only in the short range where the small scales in the model directly correspond with those in reality that the advantages of a high resolution model are clearly demonstrated in single control forecasts. At the other extreme of the time scale, in climate forecasting and evaluation, the ensemble approach is widely used (e. g., Kumar and Hoerling 2000). As shown above the use of ensembles is also important over the intermediate time scales where initial value related predictability is limited but, unlike in climate forecasting, not completely lost. As shown, the evaluation of single control forecasts alone in this environment may lead to misleading conclusions.

The true value of higher fidelity in the high resolution model in this case becomes apparent only when a number of perturbed forecast members are considered collectively as part of an ensemble. It is only in the context of an ensemble that improved predictability in terms of better ensemble mean (Fig. 5) and probablistic forecasts (Figs. 7–9) is realized. The results presented in this paper provide an example for the advantage the evaluation of ensemble (instead of single control) forecasts may bring to the NWP community. Other examples include cases of low predictability associated with forecasts of mesoscale features in the short range (Fritch et al. 1998).

5.3 *Practical considerations*

The results presented in sections 2 and 4 indicate that truncated control forecasts can reduce errors by as much as 2–4% (see Figs. 2 and 4), corresponding, for example, to the extension of predictability at 7 day lead time by 6 hours. To place this in a broader perspective, this improvement is comarable to that due to the use of the T170/L42 model in place of the T126/L28 model version (2–5% rms error reduction at 4 times higher cpu cost, see Fig. 4), and is equivalent to advances in data quality and coverage that have been achieved during the past 20 years; and is similar to advances due to the combined effect of improved data coverage, data assimilation, and NWP modeling recently achieved at NCEP in the course of approximately two years of efforts (cf. Fig. 7 of Kistler et al. 2001). These findings can have important practical implications for NWP operations.

Many NWP centers (e. g., NCEP, European Centre for Medium Range Forecasts, Canadian Meteorological Centre, Fleet Numerical Meteorology and Oceanography Center, Japan Meteorological Agency) currently run a high resolution control forecast along with a low resolution ensemble. Except at NCEP where the high resolution control forecast is truncated at 7 day lead time to save computer resources (Tracton and Kalnay 1993), the high resolution of the control is maintained to the end of the integration, typically through 10 days. Sometimes it is argued that a higher resolution model can alert the user of features that a lower resolution ensemble cannot detect. Such an arrangement of high resolution control and lower resolution ensemble forecasts, however, may be suboptimal, at least beyond a few days lead time, considering the results shown in earlier sections.

Note that the generation of the truncated forecasts (T62/L28) studied in this paper require 8–31 times less computational resources than that of the high resolution (T126/L28 or T170/L42) forecasts. So the reported forecast improvements can be achieved through the use of significantly less resources. If these computational resources can be freed up, a further increase in resolution at the short range (where this may be beneficial), more sophisticated physical parametrization or numerical algorithms, or enhanced ensemble forecasting may become possible. Based on the results presented in the study of Buizza et al. (1998) and in section 4 above, further increases in the resolution used in ensemble forecasting, for example, can be quite beneficial.

6. CONCLUSIONS

In this paper the effect of horizontal resolution on single control and ensemble forecasts was studied in various experimental and operational forecast data sets. The benefits of increasing the horizontal resolution during the initial few days of integrations was demonstrated both for the control and ensemble forecasts. As expected, it was also found that truncating ensemble forecasts around 3 day lead time in the course of the integration had a clear negative impact on all performance measures, including probabilistic scores. Curiously, however, control forecasts exhibited *improved performance* in terms of reduced error growth as compared to forecasts that were continued to be integrated at the high resolution.

The control forecast results were reproduced when high resolution and truncated forecasts were evaluated in a controled perfect model environment against following day high resolution forecats. This suggests that the inability of the high resolution control integration to produce forecasts of higher quality than the truncated version is most likely not related to problems in model formulation.

According to the tentative explanation offered in the previous section the small scale features in the high resolution forecasts that are eliminated from the truncated integrations have no skill at all. These motions act as random noise that once interact with the motions on the larger and more predictable scales can reduce the skill of *individual* forecasts. Such noise, however, since it realistically represents natural processes, can enhance performance when an ensemble of forecasts is considered.

Earlier studies (e. g., Toth et al. 1998, Mylne 1999, Richardson 2000, Zhu et al. 2001) suggested that beyond 3–4 days lead time ensemble forecasts offer more value to all potential users than control forecasts do, even if the latter are run at a higher resolution. The results of the present study can be interpreted as yet another indication for the limited utility of single forecasts.

The results suggest that for optimal short range performance the most realistic, high resolution models should be used. If an ensemble approach is followed, this is true for longer lead times as well. If only a single control forecast is used, however, best performance is achieved when the model is truncated as the smallest resolved scales lose their skill. We conclude that in the quest for ever improved numerical weather predictions increased resolution, beyond short lead times, may be an asset only if applied along with an ensemble approach.

7. ACKNOWLEDGEMENTS

The interest and support of Drs. Stephen Lord and Hua–Lu Pan of NCEP in the course of the research is greatly appreciated. We benefited from discussions with Drs. Steve Tracton of NCEP and Adrian Simmons of ECMWF.

8. REFERENCES

A complete list of references is available in the web version of this document at: http://sgi62.wwb.noaa.gov:8080/ens/ens_info.html