Monthly and Seasonal Forecasting at ECMWF

Tim Stockdale

European Centre for Medium-range Weather Forecasts

ECMWF has been developing seasonal forecast systems for several years, and although the technology is not yet mature, the point has been reached where use is starting to be made of the forecasts. ECMWF is also about to start experimental running of a monthly forecast system. Sensible interpretation and use of numerical seasonal and monthly forecasts is not easy, and the present status of key issues will be discussed.

1. ECMWF's new seasonal forecast system.

The first experimental seasonal forecast system at ECMWF was introduced in 1997, and was based on the version of the atmosphere model then available (Cycle 15r8, operational for medium range forecasting in early 1997). At that time, only limited work had been done to create a seasonal forecasting infrastructure, and the model was introduced into regular use because performance seemed to be good enough to be of potential use, and because there was a large El Niño event in progress. Since then, work has continued both to assess the performance of seasonal forecast systems, and to develop more comprehensive systems suitable for sustained operational production of seasonable forecasts. Although this turned out to be a larger task than expected, a new seasonal forecast system based on the new framework is presently being introduced.

Limited types of seasonal forecasts can be made with statistical techniques or simplified models, but the strategy at ECMWF is to use comprehensive numerical models of the ocean and atmosphere, initialised to as realistic a state as possible. The new forecast system consists of Cycle 23r4 of the IFS (operational for medium range forecasting in 2001) at T_L95 resolution with 40 vertical levels. The ocean wave model is now included, and the new physics includes improvements to the land surface and radiation, and changes to the convection and other physics. The model behaviour in the tropics differs from that in Cycle 15r8; whether it is judged "better" or "worse" will depend on which features are considered most important - i.e., the improvement is ambiguous, and not guaranteed to give better results for the specific application of seasonal forecasting.

The ocean model is HOPE, with an equivalent 1° resolution, increased to 0.3° near the equator. The model physics has also been retuned. The ocean assimilation scheme has undergone many changes. These include better quality control, a proper treatment of salinity variation, and the introduction of balanced velocity increments. These improvements appear to reduce the amount of noise in the ocean analysis. It remains true, however, that the ocean model does not fit the data very well in the vicinity of the equator. This appears to be a universal problem, and is not properly understood.

A new ensemble strategy has been introduced. The forecasts are based on a 40-member ensemble, and all ensemble members start at 00 UTC on the first day of a month. Compared to the previous method of a "lagged average" ensemble, this allows proper control over the uncertainties in the initial conditions, and also results in a more timely production of the forecast. Most of the predictability in the coupled system on seasonal timescales comes from the ocean initial conditions. These in turn are strongly influenced by the wind fields used to drive the ocean. To allow for uncertainties here, five separate ocean analyses are run, four of which use perturbed wind fields, with the perturbations based on differences between different wind analysis products. A larger ensemble of ocean analyses would be better, but would be more expensive. Accounting for other uncertainties (e.g. errors in data, data sampling issues, even model error) is also desirable, but will need more research and development. To create 40 forecasts from 5 analyses, additional perturbations are made to the SST field at the start of each forecast. This is because experience has shown that the coupled model is sensitive to the initial SST, and because the nominal SST field used to prepare the ocean analyses is known to have moderate errors. As with the wind perturbations, the SST perturbations are based on random sampling of differences between differently produced analyses of the observations.

A new data archive has been built. All output from the new seasonal forecasting system is stored in MARS, and available for retrieval and use by ECMWF Member States and other authorized users. A new stream is used, which allows proper reference to the operational system used, the method of production and the ensemble member number. The layout of data on tape allows efficient access to both the real time forecasts and the necessary calibration runs, and pre-computed forecast monthly means of all archived quantities are also available. This is a big improvement over the old data archiving, and is designed to allow many people to use and access the data, and construct tailored products for specific user needs.

2. Issues in interpretation

It is important to remember that with seasonal forecasts, we are always working in a probabilistic framework. This can be somewhat unfamiliar, both to the scientists and meteorologists developing the forecast systems, and to the users of forecast products. It complicates both the production and verification of forecasts. Only a few issues can be discussed here.

Because we use ensemble techniques, it is conceptually straightforward to distinguish when the forecast for a particular date differs from a climatological forecast. In practice, the finite ensemble size will introduce sampling errors. For our ensemble sizes of 40 for the forecast and 70 for the climate, a large or moderate shift in the pdf of a climate variable can be detected fairly reliably. Small shifts, and moderate changes in the shape of the pdf, are not reliably determined. A key point to bear in mind is that although we may be very confident that the model pdf is shifted, this is not the same as being confident that the true pdf is shifted. And the other key point is that even if the true pdf is shifted, this does not mean that the observed outcome will be shifted in a similar way. In the future we hope to be able to provide forecasts which are calibrated against actual outcomes, i.e. a calibration which accounts for all forms of forecast error, not just the mean model drift. Such a calibration is inherently probabilistic, and needs large datasets if it is to be done properly.

The potential value of a seasonal forecast depends very much on the application being considered. Theoretical models such as the Cost/Loss model have been developed for medium range ensemble forecasting, and can be used to assess the utility of forecast systems with different skill characteristics. A key point is that different users are likely to need different products if optimal use is to be made of the often limited forecast information.

A frequently asked question is "How good are these seasonal forecasts?" It would be nice to have a few simple statistics that could give a reasonable and representative answer to anyone who asked such a question. Unfortunately, given the range of potential uses of seasonal forecasts, and the variety of geographical areas and meteorological variables of interest, it is difficult to develop meaningful general statistics. The sampling error inherent in the limited number of past cases is also large, especially in areas where predictability is low for most of the time. We use some basic statistics, especially for the skill of forecasting SST anomalies, which help us assess the relative performance of different prototype systems. For specific users of forecasts, it is not clear what sort of representation of past skill is most appropriate. One possibility is simply to provide a graphical representation of past forecasts and actual outcomes for a specific quantity of interest. Formal scores come with very large error bars, given the present number of cases, and are in any case less intuitive.

Seasonal forecasts obtained directly from numerical models are not reliable, i.e., the model pdf cannot be assumed to be a close approximation to the true pdf. Some of this unreliability comes from our inability to specify correctly the uncertainty in the initial conditions. As discussed earlier, our ensemble generation techniques need further development. However, the major source of error in our forecasts is believed to be errors in the numerical models themselves. Model error seems to be a much tougher problem for seasonal forecasting than, for example, for medium range NWP. This is probably because of the much longer timescales involved (i.e., model errors have plenty of time for their impact to accumulate), the small signals being sought (i.e., the forecast changes in pdf are typically much smaller than the amplitude of an individual synoptic system), and the crucial role played by tropical convection (which is notoriously difficult to handle realistically in numerical models). The impact of model errors can be ameliorated to a modest extent by using multi-model ensembles. That is forecast products are based on the output of several different models, each of which has different model error characteristics. Such a multi-model strategy is being pursued at ECMWF, initially by working together with the UKMO. Note that although this strategy is worthwhile, and can help reduce the number of severe embarrassments, most of the degradation caused by model error remains, and it is not an alternative to improving the models themselves.

From all of the above, it is clear that communication with users is a key issue. Simply issuing "forecasts" to a disparate community of users will result in a loss of much of the potential value of seasonal forecasts. This is especially true in regions of modest to low predictability, and at times when there are no strong perturbations to the global climate. Adapting forecast products to specific user needs, and ensuring that users fully understand the forecast strengths and weaknesses, will be a vital part of future operational seasonal forecasting.

3. Scope for future development.

Numerical seasonal forecasting is a young science, and there is much scope for future development and improvement. Firstly, we need better numerical models. Today's models are often little or no better than empirical techniques when it comes down to aspects relevant to seasonal forecasting. Yet the indications are that the potential predictability is much higher than the models are yet able to achieve, i.e. that there is scope for real and substantial improvement. It also seems to be the case that the model error is dominated by systematic and regime-dependent problems, again a sign that we should be able to do better. The speed of future model improvement is not known. At the moment there are many other issues for us to address, but at some time in the future it could be that model error becomes the single dominant outstanding issue.

It is also important that seasonal forecast systems are better characterized. A system may have quite a large number of faults and limitations, but if they are known they can be allowed for, and the system can be very useful for those things it gets right. In order to characterize systems properly, we need a significantly larger number of past cases (hence more computer resources), and a large amount of work (hence human resources). Over time, however, progress will be made.

There is also the issue of calibration of probabilistic forecasts, as has already been mentioned. This again requires a lot of past cases, and also suitably large ensemble sizes. A Bayesian framework seems to be most appropriate for this sort of work. It may be that different applications will want to make different prior assumptions about the probability distributions, to ensure an appropriate degree of robustness for each application.

Most applications want forecasts valid locally, or at least for smaller regions than are properly resolved by our present models. Higher resolution and/or appropriate downscaling techniques will need to be developed and applied.

At the moment ECMWF creates a very limited set of graphical products from its seasonal forecast system. Although the new data archiving allows for much more activity in constructing custom applications, it is still true that default graphical products will very often be used at least to explore the possibility of a worthwhile seasonal signal being present. To this end, the range and appropriateness of products needs to be improved.

Finally, the actual integration of seasonal forecasts into applications has barely, if at all, begun. Ideally the full probabilistic information of a seasonal forecast should be passed through to the user application, so decision making can be made based on appropriate probability distributions of the relevant outcomes.

4. Monthly forecasting

Seasonal forecasting at ECMWF looks at the coming six months, and forecasts are calculated at monthly intervals. Medium range forecasts look to the coming 10 days, and are calculated daily. In between these we will shortly introduce monthly forecasting. The aim here is to produce an outlook for the coming month. Scientifically this is challenging, because it goes beyond the deterministic range (so the atmosphere initial conditions have only a limited impact on what will happen), and yet it is not really long enough for the boundary conditions (principally SST) to have a strong impact either. Nonetheless, ensemble techniques should be able to tell us what predictable signals are present, and where they are.

The specific configuration planned at ECMWF is to run a 51-member ensemble at $T_L159L40$ resolution. The structure of the forecasts will be the same as in the seasonal system, except that they will include a real-time ocean analysis component and also the same Singular Vector computations as are used in the medium-range EPS. Ideally the forecasts would be made weekly, but initially they will be run only every second week. As with the seasonal forecasts, a set of back integrations will also be run, consisting of a 5-member ensemble for the 1991-2000 period. This will allow the output of the monthly forecasts to include all of the seasonal-type products, as well as the products from the 10-day EPS system. This monthly forecasting is experimental, and is designed to tell us more about what can be expected from such a system.

5. Summary

Numerical seasonal forecasting is now emerging operationally. This is true not just at ECMWF, but at various organizations worldwide, and within the structures and protocols of the WMO.

Seasonal forecasting, as practised today, has useful skill in some areas at some times. However, proper use of the forecasts is still at any early stage. It is an important challenge for both the scientific community and national meteorological services to encourage appropriate (and discourage inappropriate) use of seasonal forecasts.

Importantly, there is scope for real improvements in forecasts. Today's levels of reliability do not set the limit of what to expect in the future. However, the actual level of predictability will always be limited by the physics of the world in which we live. Some users will be disappointed that the future cannot be foretold more exactly. Others will be very pleased with what they can do with the seasonal forecasts we will one day be able to produce.