



Met Office Challenges on Convection*

Mike Cullen, Alison Stirling, Michael Whitall

August 28, 2015

Below three different challenges are defined on convection, which are followed by concluding remarks.

1 Validating Convective-Scale Rainfall Forecasts and Estimating their Uncertainty

By Michael Whitall

1.1 Introduction

Weather forecasting centres are increasingly making use of high-resolution NWP models, which cover a local domain of interest (e.g. the UK) with a mesh-size of the order ~ 1 km. These models are capable of explicitly simulating convective showers and thunderstorms (whose dynamics occur on relatively small scales), so they are a vital tool for forecasting high-impact, heavy-rainfall events.

The model configurations used for these simulations are under ongoing development, and it is important to validate and compare new experimental versions of the models in terms of how well they predict the location and timing of convective rainfall.

Since uncertainties in the initial conditions and model physics grow more rapidly at the smaller scales, it is important to take a probabilistic approach to forecasting rainfall patterns during convectively unstable conditions. Ensembles of model simulations are performed to sample the likely range of rainfall patterns, but limited computational resources restrict the ensemble size so that it is not sufficient to estimate robust probabilities for the chance of rain at each location.

Forecasters and model-developers currently lack the statistical tools needed to validate convective-scale rainfall forecasts in a purely objective way. Also, current methods for estimating the “chance of rain” from the limited ensemble data are ad-hoc and sensitive to arbitrary parameter choices. In section 1.2, we describe the requirements for a “displacement metric”, measuring the positional error of rain-showers in a forecast relative to observations. In section 1.3, we

*Environmental Modelling in Industry Study Group, Isaac Newton Institute of Mathematical Sciences, Cambridge, Sept. 21-24, 2015 (Maths Foresees & PURE Networks)

describe the need for a robust statistical method to estimate the plausible range of rain-shower locations (and hence the chance of rain at each location) from a forecast ensemble with a limited size.

1.2 Validation of Convective-Scale Rainfall Forecasts

When developing high-resolution weather forecast models, it is vital to be able to assess the relative performance of different model configurations. A crucial element of model performance to assess is the skill of the model in predicting the location and timing of intense convective showers and thunderstorms. At a given time, the true spatial pattern of rainfall intensity can be estimated from radar observations, and this pattern can be compared with the 2-D field of surface rainfall predicted by the forecast model, which has been initialised using data from an earlier time. To illustrate this, figure 1 shows made-up examples of an observed rainfall pattern, a forecast rainfall pattern, and the difference between the two.

Traditionally, the error of a model relative to the observations is assessed by calculating the point-wise root-mean-square difference between the two:

$$\text{r.m.s.e.}_P = \sqrt{\frac{1}{XY} \int_{y=0}^Y \int_{x=0}^X (P_{\text{forecast}} - P_{\text{observed}})^2 dx dy}$$

However, in the case illustrated by figure 1, this is not necessarily a meaningful measure of the error of the model. Unless the forecast areas of rain actually overlap spatially with the observed areas, the r.m.s.e. will be maximal and will always be the same. If an improved forecast model B predicts the rain-showers closer to their observed locations than forecast model A, but they still don't overlap significantly with the observed locations, the r.m.s.e. will misleadingly suggest that forecasts A and B are equally poor.

Also worryingly, a forecast model that unrealistically smooths the predicted rainfall pattern (e.g. predicts light drizzle across the whole domain) will typically get a better r.m.s.e. score than one that models realistic rain-shower structures, and only slightly misplaces the shower locations.

Clearly, we need a bulk measure of the spatial displacement of the rain-showers between the forecast and observations. i.e. what is the integrated distance over-which you have to spatially move the forecast rainfall to reproduce the observed rainfall pattern? If the rain-rate is expressed in units of $kgm^{-2}s^{-1}$, the bulk displacement error will have units rain-rate \times area \times distance = $kg s^{-1}m$. To estimate the displacement error, we need to find the most efficient displacement mapping to move the forecast rainfall mass onto the observed pattern. A way of visualising the displacement error metric is as follows:

Imagine that both the observed and forecast rainfall patterns are divided up into equal, infinitesimal bits of rain-rate dP , where each bit has a precise coordinate x, y in space. Then imagine that every bit of the forecast rainfall pattern dP_{forecast} must be connected to an equal bit of the observed rainfall pattern dP_{observed} by a straight length of string. The displacement metric is

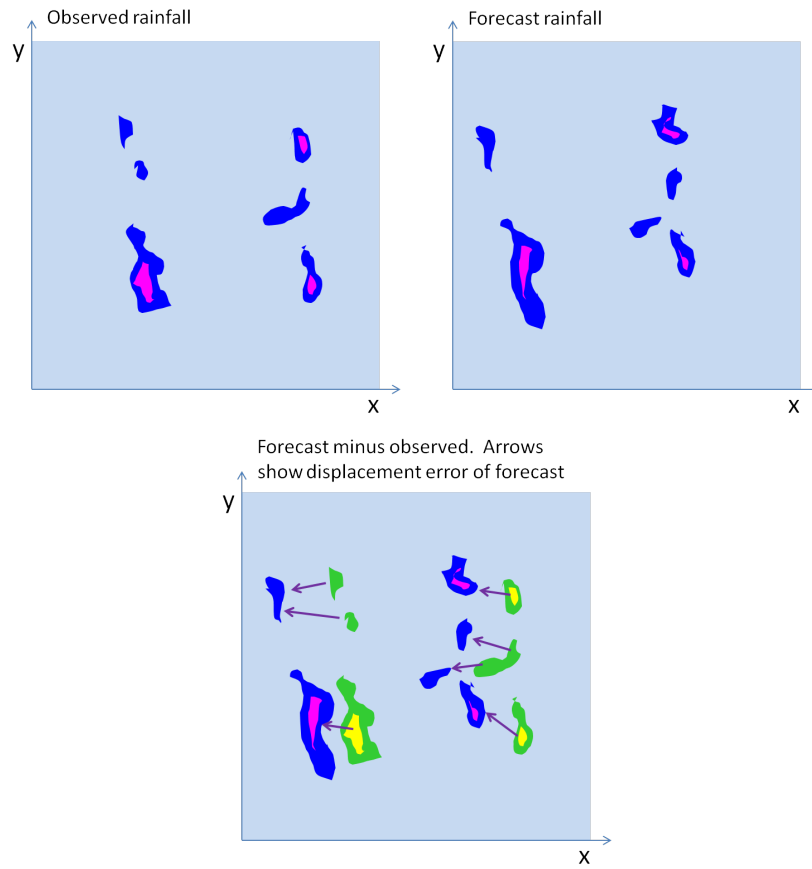


Figure 1: Schematic example maps of surface rain-rate; (top) observed rainfall, (middle) model-predicted forecast rainfall, and (bottom) forecast minus observed, with arrows indicating the displacement we wish to estimate.

then the shortest possible total length of string needed to connect all dP_{forecast} to all dP_{observed} .

The question I want to ask a mathematician is, how do you efficiently calculate this displacement metric? Some potentially complicating factors that need to be bourn in mind are:

- The observations and the forecast don't in general have equal domain-integrated rain-rates, so the error metric should have components both from spatial displacement and also the surplus / deficit of the forecast rainfall relative to the observed.
- The domain being considered will have boundaries, since we usually want to evaluate the model over a particular region of interest (e.g. the local UK domain simulated by our 1.5 km resolution forecast model). How should we handle rainfall on the domain boundaries when calculating the displacement error?

1.3 Estimating the chance of rain from sparsely-scattered showers using a limited-size forecast ensemble

When forecasting the rainfall pattern on days with sparsely-scattered convective showers, the uncertainties on the convective scale usually grow very quickly so that after a few hours, the forecast really has very little confidence about the exact locations of the showers. However, there is usually still some pattern to the distribution of rain-showers that is predictable (e.g. they sometimes have a tendency to form over hilly regions and then drift down-wind).

In order to quantify the uncertainty in the forecast rainfall and see a representative range of possible weather outcomes, forecasters use ensembles of model runs, where each ensemble member randomly samples the uncertainties in the initial state and the model physics. However, convective-scale model-runs are very computationally expensive, so we are limited to running quite a small ensemble.

If we had an infinitely large ensemble, the chance of rain at a particular location would be straightforward to calculate, simply as the fraction of the ensemble members in which rain is forecast to occur at that location. Unfortunately, the combination of sparsely-scattered rain-showers and a restricted ensemble size (of about 12 members) means that the fraction of ensemble members producing rain at an individual model grid-point utterly fails to give a robust probability. Many grid-points will fall between the locations of showers in the ensemble members, misleadingly implying a zero chance of rain.

Current efforts to get round this problem involve considering arbitrary square neighbourhoods that cover many grid-points, and assuming that if a rain-shower is forecast to occur somewhere in that neighbourhood, it is just as likely to occur anywhere in the neighbourhood. A probability of rain at a point can then be estimated by considering how many ensemble members forecast rain within the wider neighbourhood. However, this method has problems:

- The forecast probability of rain is highly sensitive to the chosen neighbourhood size.
- The model might simulate a particular tendency for showers to occur in one part of a neighbourhood rather than others (e.g. due to the influence of hills); the neighbourhood method would ignore this information and assign a uniform chance of rain over the whole neighbourhood.

The question I want to ask: is there a statistical method of deriving a more robust probability of rain from a limited ensemble size (which cannot on its own sample the full range of possible locations of the scattered showers)? And is there such a method which doesn't include any ad-hoc parameters or assumptions, and which can retain forecast signals for the most-likely locations of showers on all spatial scales?

I could envisage that such a method might involve building an emulator, which can take the different spatial patterns of rainfall predicted by the forecast ensemble members, and cheaply generate many more rainfall patterns deemed to have equal probability of occurrence to the model ensemble members.

A key piece of information it could use might be the rainfall location displacement error metric described in section 1.2 above. You could compare all possible pairs of forecast ensemble members and calculate the rainfall displacement error between each pair. During the calculation you could store the 2-D field of horizontal displacement flux of rain-rate needed to move the rainfall from the pattern predicted by one ensemble member to the pattern predicted by another (this corresponds roughly to the arrows shown in the forecast minus obs picture in figure 1). This 2-D field (its magnitude averaged over all pairs of ensemble members) would give an indication of regions that lie between forecast shower locations in the different ensemble members, and which might therefore also have an increased chance of rain. You might want to down-weight the influence from pairs of ensemble members in which the rain has to be moved over larger distances (e.g. the forecast might be undecided over whether the showers will occur over one range of hills or another range 100 km away, but there could be a good reason why the rain isn't forecast to occur over the low-lying area in-between). The emulator-generated rainfall patterns could be designed to have the same statistics of the displacement error metric between each pair of patterns as the metric generated between pairs of model ensemble members.

Such an emulator would be straightforward to validate; run it to generate hundreds of plausible rainfall patterns for a large number of forecast case-studies, use the resulting super-ensemble to estimate the chance of rain at every grid-point in the model for each forecast, stratify the forecasts / grid-points based on this chance of rain, and check the observations to see what fraction of the time rain actually occurred for each forecast chance of rain. E.g. if you consider all occasions on which you predict a 50% chance of rain, 50% of those occasions should have observed rain!

2 Convection-Wave coupling

By Alison Stirling

A long-standing problem in the representation of convection in global models is getting the coupling right between convective activity and tropical waves. At present we think parameterised convection disrupts the progress of tropical waves, and there is minimal feedback between the two. We believe getting this coupling right to be crucial for improving the longer term predictability of our forecasts.

a) The Met Office has been approaching this problem by first considering what the atmospheric response to convective heating should be (in the absence of wave disturbances). To do this we have been simulating plumes of warm air rising from the surface, and aiming to characterise the transient divergence / convergent response to these plumes. At present we have a somewhat ad-hoc representation of this divergence pattern that might capture some of the non-local behaviour, but could this be calculated mathematically?

So, given a localised heating term that grows in height as a function of time, $h(x, t)$. What is $\mathbf{v}(\mathbf{x}, t)$? (For a variety of environmental profiles and Coriolis parameters).

b) A second part to the problem that we have not considered, but is equally important: express mathematically how an area of convection interferes with tropical wave progression, where convection can be represented as a column of lower stability than its surroundings.

3 Adjustment of a column of convectively unstable moist air

By Mike Cullen and Dorian Goldman

When moist air parcels rise, they cool, eventually reaching a point where their water vapour content starts to exceed the saturation value for that height (or pressure level). The excess water vapour condenses, releasing latent heat which acts to increase the parcels temperature. This boosts its buoyancy relative to the environment and it rises further and, in the absence of mixing, tracing out a saturated adiabat.

Air parcels are observed to arrange themselves quickly so that the parcels with highest potential temperature rise to the top of an atmospheric column, ensuring that the potential temperature is monotonically increasing at all times. This monotonicity can be applied as a constraint to the problem of adjustment, as proposed in Cullen and Purser (1984). Without the effects of moisture, potential temperature is conserved in time and hence monotonicity is also conserved. But what happens when moisture effects are included?

M.J.P Cullen & R.J. Purser. An extended Lagrangian theory of semi-geostrophic frontogenesis. *J. Atmos. Sci.* **41**, 1477-1497 (1984).

4 Concluding Remarks: linking a convection model to the explicit large-scale dynamics

The convection problem can be viewed as the occurrence of systematic smaller scale motions in a an environment which otherwise varies fairly smoothly in space. A natural way of treating this situation is adaptive mesh refinement, where a much higher (convection permitting) resolution is used in the convection regions. While this is likely to work as well as a model with convection permitting resolution everywhere, the prevalence of atmospheric convection. particularly in the tropics, means that it is not affordable.

Convection schemes are usually designed as single column models which replicate the averaged fluxes obtained from cloud resolving models. An affordable substitute for adaptive mesh refinement would thus to be use the convection scheme in the place of the large-scale model in regions where the convection is diagnosed. In current practice, the convection scheme is added on top of the large-scale model.

The replacement strategy has several potential large advantages. Exact mass balance in the column would not be required, allowing remote subsidence as shown in diagnostics from high resolution convection permitting simulations where appropriate. The situation where explicit convection develops where the convection scheme does not remove instability fast enough would also be avoided. Currently the convection scheme has to be forced to work unrealistically fast in order to prevent this.

Though this looks very attractive, it will be extremely challenging to devise a way of making this work, as a formulation which is both analytically and numerically stable would have to be found. This would then need evaluation in a simple model, such as a vertical slice, to show that it was viable.