On the chance of rain: probability forecasts in an operational setting

Ed Wheatcroft

Dynamic Days Europe Conference, Exeter

7 September 2015

Abstract

In the forecasting of nonlinear dynamical systems, ensembles of model simulations are often used to account for initial condition uncertainty, For example, in weather forecasting, ensembles of around 50 members are launched every 6 or 12 hours, with each simulation extending a week or more from its launch time. It is often convenient to interpret the information contained in ensembles at any given time by using them to form predictive distribution functions called forecast densities. Applying traditional density formation techniques such as kernel density estimation would aim to recover the distribution of the ensemble. This, however is not satisfactory when the model is imperfect since our aim is to estimate the distribution of the outcome which is necessarily different to that of the ensemble. We show how an approach called kernel dressing can be used to account for model error and can yield more skillful forecast densities when measured using a logarithmic scoring rule. In the second part of the talk, we ask whether the information contained in forecasts of the same target time but launched sequentially in time can be combined in some way to yield improved forecast skill. We argue that, if the forecast densities could be considered to represent probabilities of the predictand, a Bayesian approach would be expected to yield optimal skill. We show, however, that this is not the case, casting doubt on the validity of treating forecast probabilities in this way. The ideas in this talk are demonstrated using both low dimensional chaotic systems such as Lorenz ’63 and using a high dimensional model of real-world weather observations.

LSE