Keith Runcorn Prize Lecture: Can randomness reduce uncertainty? The use of stochastic physics in weather and climate prediction

Hannah Christensen

November 2015

I would like to start by asking the following question: can you trust a weather forecast? For example, consider the case of a deterministic, or 'best guess', forecast for the next few days. Such a forecast indicates what the weather forecaster considers to be the *most likely* weather for the coming week, with no indication of how certain they are in their prediction. If you pose this question, 'can you trust a weather forecast', to a member of the public, the likely answer is: no, you can't trust a forecast. Or at the very least, that you should take such a forecast with a pinch of salt — they are often wrong. So when provided with a best-guess forecast for the coming week, even if no indication of the uncertainty in the forecast is given, the public will infer some uncertainty in the forecast given their past experience. Unfortunately for the weather forecaster, the public tends to remember those occasions when the forecast went horribly wrong rather than those occasions when the forecast was spot on. The perfect example of this is the most infamous UK weather forecast of all time: Michael Fish's mis-forecast of the Great Storm of 1987. On the evening of the 15th October, Michael Fish gave the following forecast: "Earlier on today, apparently, a woman rang the BBC and said she heard there was a hurricane on the way; well, if you're watching, don't worry, there isn't ...". The following morning, people woke to scenes of devastation: gusts of 90-122 mph had swept across the south of the UK overnight, uprooting 15 million trees, and leading to the deaths of 18 people. We can repeat the forecast for that night using a modern weather forecasting system. If we produce a 'best-guess' forecast, very much like the forecast which Michael Fish would have seen, it predicts a mild, low pressure system over the south of the UK — certainly nothing concerning, or out-of-the-ordinary. Even given thirty years of advances in the computer simulators used to make weather predictions, the forecast model still predicts a calm night, instead of the observed 'Great Storm'.

However, in addition to the 'best-guess' forecast, a modern weather forecasting system also predicts the uncertainty in the forecast. Fifty alternative, but equally likely forecasts are produced which represent the uncertainties involved in making a forecast: this is called an ensemble forecast. If we produce this probabilistic ensemble forecast for the night of the Great Storm, we see that while most ensemble members indicate a calm night, a third of the ensemble members indicate the possibility of a very deep low-pressure system, with very tight isobars, indicating a strong wind storm. Instead of leaving it up to the public to infer how certain they can be in the forecast, it is important to explicitly calculate this using our forecasting model. In the case of the Great Storm, it was as if the atmosphere was teetering on a knife edge, and it was impossible to tell which way it would fall. Presenting this uncertainty information to the public gives a much more accurate and useful forecast. An important property of these probabilistic forecasts is that they are *reliable*. This refers to the statistical consistency of the ensemble forecast with the observations. For example, it asks the question, if I collect together all the occasions when the forecast predicted rain with a 20% probability, and looked at what actually happened on those occasions, did it rain 20% of the time? If the forecast probabilities are consistent with those observed, the forecast is said to be reliable.

In order to achieve reliability, we must represent all sources of uncertainty in the forecast. There are two sources of uncertainty in particular which it is important to represent in a weather forecast. The first is initial condition uncertainty. We only have access to limited satellite and weather balloon data with which to estimate the current state of the atmosphere for initialising our models. We can represent the uncertainty in the initial conditions by perturbing them slightly between the fifty members in our ensemble forecast. As the atmosphere is a chaotic system, these very small errors can grow rapidly and, in some cases, can lead to very different forecasts after just a few days. The second key source of uncertainty in a weather forecast is model uncertainty. This stems from the fact that our computer simulator is merely a model of the atmosphere, and includes many approximations and simplifications. In particular, small scale processes are not resolved by the model, and must be approximated through parametrisation schemes. Errors introduced into the forecast in this way can also grow rapidly in time. While accurate representation of initial condition uncertainty is well understood in the atmospheric community, there is still much debate as to the optimal way to represent model uncertainty. The remainder of this talk will discuss this important question, and will focus in particular on two proposed techniques: stochastic parametrisation schemes, and perturbed parameter approaches.

We will begin by considering these two representations in an idealised model: the Lorenz '96 System. This can be thought of as a 'toy model' of the atmosphere. It consists of a set of coupled equations with two types of variables arranged in a ring: the large-scale, low frequency variables represent large scale atmospheric dynamics, and they are coupled to small-scale, high frequency variables, which can represent individual convective clouds. We use this system to perform a set of idealised experiments. We run the full set of equations with both large and small scale variables — this is what our "real atmosphere" is doing. We can also build a forecasting model of this system, where we assume that the small scale variables are unresolved. However, we must represent their influence on the large scale variables, so we develop a simplified representation, or parametrisation scheme, to use in the forecast model. This parametrisation scheme is an approximation — it gives the most likely impact of the sub-grid scales on the resolved scales, but does not represent the day-to-day variability. We can now explore what the two representations of model uncertainty look like in this system. In a perturbed parameter approach, we take the uncertain parameters (physical constants) in a parametrisation scheme and perturb them between ensemble members to explore the possible range in their values. In a stochastic parametrisation scheme we introduce random numbers into our equations of motion to represent errors in our parametrisation scheme. Instead of representing the most likely impact of the sub-grid scales on the resolved scales, a stochastic scheme represents one potential realisation of the sub-grid variables. Using the Lorenz '96 model, we perform a series of ensemble weather forecasts. All members are initialised from perfect initial conditions, removing initial condition uncertainty. This allows for a clean test of the two representations of model uncertainty in the system. We find that, while the perturbed parameter scheme improves on weather forecasts which do not represent model uncertainty, it is the stochastic approach which gives reliable forecasts. Some days, the ensemble spreads out indicating high uncertainty whereas on other days, the ensemble members stay close together indicating little uncertainty in the forecast. So, returning to the title of the talk, in some cases, randomness can reduce uncertainty.

Having shown the potential of stochastic parametrisation schemes in a simple model, we move to considering their impact in an operational weather forecasting model — the

Integrated Forecasting System (IFS) used at the European Centre for Medium Range Weather Forecasts (ECMWF). In particular, we are interested in the representation of uncertainty in the convection scheme, as this is the parametrisation scheme to which models are most sensitive. The IFS contains two operational representations of model uncertainty, both stochastic. For comparison, we develop a perturbed parameter scheme. A Bayesian parameter estimation approach is used to measure the uncertainty in four of the parameters in the convection scheme — the resultant joint distribution in the parameters is used to determine the degree of perturbation between ensemble members. We also consider a generalisation to one of the stochastic approaches to address the uncertainties in different atmospheric processes independently. This is in contrast to the current stochastic approaches, which are holistic in nature. We find that the perturbed parameter scheme improves on the operational stochastic schemes — the resultant forecasts are more reliable. However, the new 'independent' stochastic approach performs the best, and produces reliable ten-day forecasts in areas of the world where convection is an important atmospheric process.

So it appears that stochastic parametrisation schemes are a powerful tool for representing uncertainty in weather forecasts, but could they also be used to represent model uncertainty in climate prediction? There are good theoretical reasons for including a stochastic parametrisation in a climate model. We hope that such a scheme will improve the short-timescale 'weather' variability in the model. In turn this can improve the statistics of the modelled climate through noise induced drift, noise enhanced variability and noise activated regime transitions. In fact, in both the Lorenz '96 system and in coupled climate models, we do observe an improvement in the representation of the climate of the system when a stochastic scheme is included. In the Lorenz model, we find that including a stochastic parametrisation scheme leads to an improvement in the regime behaviour of the system. In the climate model developed at the US National Center for Atmospheric Research, we see an improvement in the simulation of the El Nino-Southern Oscillation (ENSO). This is the name given to the irregular oscillation in sea surface temperature observed in the Tropical Pacific, where it is the dominant mode of climate variability. In fact, we're currently experiencing a Large 'El Nino' episode, with sea surface temperatures approximately 3^{o}C above normal.

So to conclude, can randomness reduce uncertainty in forecasts? I would argue that in weather forecasts, the answer is 'yes'. Stochastic parametrisation schemes allow us to produce reliable probabilistic forecasts, which indicate how predictable the coming weather is, and therefore how certain we can be. In the case of climate models, I would say the answer is 'maybe'. While stochastic parametrisations certainly improve the ability of a climate model to represent the real atmosphere, it is yet to be tested whether this will reduce or increase uncertainty about future climate change.

The President. If you'd like to know what the weather's going to be tomorrow now is that time to ask. [Laughter]. We have time for a few questions.

Dr. G. Q. G. Stanley. Very fascinating. It would be lovely to know what the weather's going to be like for the weekend so if you can tell us, we'd appreciate that. [Laughter]. I know there is a technique where you can introduce noise into a signal you're processing and then what happens when you do the Fourier analysis is it basically gets rid of the other noise. You put it back together and you get the crisp signal out of it so I guess that is roughly what you're doing at this point?

Dr. Christensen. So that's very relevant for the discussion about ENSO — stochastic resonance is what you're referring to. The idea is that if you put a white noise forcing into your system then the system picks up the particular frequency which is characteristic of the system and then resonates with that frequency. In that way you can magnify a signal among the noise. So that would predict an enhancement of some of the variability in a system. If you put additive noise into a very simple model of ENSO, you see it enhancing ENSO, increasing the magnitude. But I suppose what we saw here was a reduction in power in ENSO when we put in a stochastic parameterization, and you can see this in terms of the standard deviation of the sea surface temperatures as a function of month. It's very interesting — it doesn't seem to be kind of classic stochastic resonance. I think in part that's because we're using the multiplicative noise term as opposed to an additive noise term and you tend to think of stochastic resonance for additive noise but I think it's definitely very relevant.

Dr. Stanley. So to follow on from that, is when you're results at the end, do you then post- process it and take out the randomness you've injected?

Dr. Christensen. Ah, no so for the randomness in a weather forecasting model, the timescale of the randomness is about six hours and the spatial scale is about 500 km. Because we're perturbing the physics tendencies themselves then that in turn affects the flow of the model. So there's no removing of the randomness at the end. I suppose what we do, however, is make is a range of forecasts for the future. So if you just want to work out your best guess forecast, I like to think of a stochastic parameterization scheme as somehow inverting the order of your averaging. So you can explore these possible flow situations and then average to find your best guess, or if you use a deterministic parameterization scheme you're kind of doing the averaging as you go along. You're putting into the model the best guess representation of the cloud within a grid box, and then seeing how the model responds to that most likely cloud. Whereas, in a stochastic scheme you explore possible clouds I suppose and then average at the end.

Dr. G. Q. G. Stanley. I suppose the bad news for astronomers is that they don't need good seeing, the randomness there would help their results I suppose.

Professor P. G. Murdin. That was a very interesting talk indeed. I wondered, but it contained extremely subtle attitudes about communicating about probability and forecasting. I wondered if you'd given any thought as to how it is one would go about explaining this to the public at large and the political politicians [Laughter] to the people who are in charge of the especially in relation to climate modelling, not to mention weather forecasting. Whether you've given any thought to how to get across the ideas that you've been talking about?

Dr. Christensen. Thank you, that's a really important question. I think in the realm of climate modelling, communicating uncertainty is a little bit ahead of where it is in weather forecasting in the UK. I think that's because there are other very obvious things that are uncertain about predicting the future climate, such what carbon dioxide emission track are we going to go along. So certainly I think that's communicated, or is attempted to be communicated, fairly well to the policy makers. Especially out of these big International Panel on Climate Change reports — so that's very explicit in terms of what is very likely or possible or less likely in terms of impacts on the climate. For weather forecasting, I think it's a real shame that we don't actually give out more of this probabilistic information to the public. The Met Office has made some efforts in recent years to move towards this but I think they're worried that if they just present a PDF to the public that they'll think that they're hedging their bets in some way. [Laughter].

Professor Murdin. They are, but that's the nature of the equations.

Dr. Christensen. Well, exactly but I think if people communicate to the public that some days we only know that it will rain with say forty percent certainly and other days we know it's not going to rain and other days we're much more certain, I think that would be better. There were rumours that the Met Office and the BBC fell out over the Met Office wanting to put more probabilistic information into their forecasts.

Dr. M. B.. So is that why the BBC gave up using the Met Office or are about to give up using the met office?

Dr. Christensen. Well that is the rumour, but it might be also a financial issue.