

Possible or probable?

Myles R. Allen

As a record-breaking hot summer in Europe draws to a close, the issue of global warming is once more in the news. Away from these headlines, those with the job of keeping our doorsteps dry over the coming decades, and our water, food and energy supplies secure, are now actively planning for climate change. This transition, from politics to practicalities, presents the climate-research community with new challenges. Today's coastal and water-supply engineers do not need old-style 'projections' of how the climate might respond to rising levels of the greenhouse gases, no matter how detailed. Projections of what might happen in the future are fine for lurid headlines, but practical planning needs exactly the opposite kind of information. The challenge of probabilistic — or risk-based — climate forecasting is to start saying what changes can be ruled out as unlikely, rather than simply ruled in as possible.

The Intergovernmental Panel on Climate Change (IPCC) recognized the need for a probabilistic dimension to climate forecasting in their 2001 Assessment Report, although they refused to assign an explicit probability to their 1.4–5.8K range for projected warming over the twenty-first century — despite intense pressure to do so. One of the main reasons for the IPCC's reluctance to interpret this range as a formal uncertainty estimate was because it depends on the spread of results from a small number of climate models, included primarily because they happened to be available at the time. Models taken 'off the shelf' tend to be too similar to each other to be considered the kind of random, mutually independent sample beloved of statisticians. Checking one model's results against another's, when both have been developed with reference to the same observations, does not quite deserve the philosopher Ludwig Wittgenstein's jibe about buying several copies of the morning newspaper to assure yourself that what it says is true — it's more like buying two newspapers that rely (to an unknown degree) on the same wire service.


So what can be done? At first glance, generating a probabilistic climate forecast seems straightforward. Any forecast beyond the next few years must allow for uncertainty in how our models represent the climate system. This uncertainty is found in the values of crucial parameters that are not well constrained by observations — such as the reflectivity of clouds — and more generally in how processes such as cloud formation should be modelled. So we begin with

a representative sample of possible models, weight them by some measure of their similarity to the real world, and use this weighted 'ensemble' of forecasts to infer future risks.

The problem is obtaining that representative sample of possible models defined, crucially, without reference to the observations used to weight the ensemble. If the same observations are used to select models initially as are subsequently used to weight them, then we 'double-count' and inevitably underestimate uncertainties (back to the newspapers analogy). But observations are used throughout the process of climate-model development in such an ad hoc way that it is impossible to disentangle the influence of any particular data set. Leaving aside the double-counting problem, how can we be sure that our forecasts depend primarily on observations (which, although uncertain, tend to be revised and updated relatively slowly) and not on the earlier choice of models or perturbations (which are subject to the whims of expert opinion)?

To be reliable, probabilistic climate forecasts must begin with a perturbation analysis of one or more climate models to identify consistent relationships between observable quantities and forecast variables of interest. We then weight the individual perturbed model-versions to ensure that the entire ensemble accurately represents both current knowledge and uncertainty in these observable quantities. We can then infer future probabilities from the weighted forecasts, comfortable in the knowledge that — provided these relationships are consistent across physically reasonable models of varying structure and resolution — our results should depend on observations, and not a dubious earlier choice of models. It is a subtle but profound change in our attitude to climate models. Rather than providing surrogates for reality, they become tools for teasing out useful relationships between things we can observe and things we want to forecast.

This is straightforward enough when the relationships in question can be represented



World view: desktop computers are uniting to reduce uncertainty in climate simulations.

Climate forecasting

Climate modellers need to start saying what changes can be ruled out as unlikely, rather than simply ruled in as possible.

by simple, low-order functions, but most forecast quantities of interest will be related to several independent observable quantities. With realistic climate models, we not only require many simulations with different starting conditions for each model-version, but we also cannot 'aim' simulations at poorly sampled regions on the manifold. We just have to keep perturbing our model(s) until we fill out this multi-dimensional space of responses, potentially requiring hundreds of thousands of century-timescale simulations of a full-scale climate model.

Mapping the response manifold of a full-scale, non-linear climate model is a truly formidable challenge, well beyond the capabilities of conventional supercomputing resources. The only way to access sufficient resources is to use idle processing capacity on home and desktop personal computers. This is the climateprediction.net approach, proposed on these pages almost four years ago, following the successful launch of the SETI@home project (<http://setiathome.ssl.berkeley.edu>), which is now by far the largest single computation ever performed. Thanks to the support of the UK Research Councils for coupled modelling and the help of the Meteorological Office, the ingenuity of a dedicated group of scientists and the enthusiasm of a small army of beta-testers, we have configured one of the world's best climate models to run on almost any up-to-date Windows PC. If you own a PC and would like to take up the challenge of probabilistic climate forecasting, please join us at <http://www.climateprediction.net>. ■

Myles R. Allen is in the Department of Physics, University of Oxford, Oxford OX11 0QXJ, UK. This was written on behalf of the climateprediction.net project team, with assistance from D. A. Stainforth and D. J. Frame (University of Oxford), A. Kettleborough (Rutherford Appleton Laboratory) and M. Collins (UK Meteorological Office).

FURTHER READING

Houghton, J. T. *et al.* (eds) *Climate Change 2001: The IPCC Scientific Assessment* (Cambridge Univ. Press, Cambridge, 2001).
Schneider, S. H. *Climatic Change* 52, 441–451 (2002)
Allen, M. R. *et al.* in *Proceedings of the 2002 ECMWF Predictability Seminar* 275–295 (ECMWF, Reading; online at www.climateprediction.net/science/pubs).