

The probability of rapid climate change

If you look at a map of the air temperature of the surface of the Earth, you will see that North West Europe, including the UK, is warmer than Alaska, which is at the same latitude but on the Pacific rather than Atlantic Ocean. At school you were probably told that this was because of the Gulf Stream. However, there is a very similar current in the Pacific—the Kuroshio—which takes warm water north past Japan and then out into the Atlantic. Peter Challenor asks: What is the unique feature of the Atlantic that keeps us warm and could it change in the next few years?

The main difference between the currents in the two oceans is what is called the thermohaline circulation. Ocean currents such as the Gulf Stream and the Kuroshio are driven by the wind, and so are known as the wind-driven circulation. There is another part of the ocean circulation that is driven by density differences caused directly by solar heating, because the density of seawater is related to temperature and salinity; this is known as the thermohaline circulation.

Simple but crucial

The thermohaline circulation is fairly simple. Water cools in the Northern North Atlantic and, more importantly, becomes saltier, because sea ice is fresh, rejecting salt as it freezes. It therefore becomes denser and sinks. This sinking then draws warm water North through the Atlantic. The cold, salty water is now at the bottom of the ocean. It flows south through the Atlantic, warming as it travels, and finally comes to the surface in the Antarctic and the North Pacific. It then moves back into the Atlantic south of Africa and is drawn north again. This circulation is known as the ocean conveyor belt and is shown schematically in figure 1. To complete this circulation can take hundreds of years, but the net result is that heat is transported north in the Atlantic.




Figure 1 The Thermohaline Circulation

Is it possible that global warming could shut down this circulation and hence deprive Europe of this heat? If you add large amounts of fresh water to the North Atlantic it will stop the water becoming dense and sinking; this will shut down the thermohaline circulation and depress the temperature in Europe. Where might this water come from? Global warming will lead to increased evaporation and hence rainfall, and might also melt the Greenland ice sheet.

Simple dynamical models have two possible stable solutions: one with a thermohaline circulation and one without. These models are too simple to make predictions from, but may be indicative of the non-linear behaviour of the real climate system. There is evidence from ice cores that, in the distant past, very rapid changes in temperature have occurred around the North Atlantic. This can be seen in figure 2, where we can




Figure 2. This figure shows the amount of the isotope oxygen 18 in the GRISP ice core from Greenland against time before present (1950). More O18 means that the temperature of snow falling on the ice surface is higher, thus we can get a record of temperature back through the past. Note the rapid cooling associated with the Younger Dryas and the smaller 8.2k year event. These rapid climate changes are believed to be associated with the collapse of large freshwater lakes caused by deglaciation. The consequent input of freshwater into the North Atlantic caused the cessation of the THC and hence the rapid cooling.

see two severe cooling events since the end of the last ice age. The older of these is known as the Younger Dryas (named after an alpine flower) and the younger, more prosaically, as the 8.2k year event. Another final piece of evidence comes from the large general circulation models such as those run in the UK at the Hadley Centre (part of the Meteorological Office) and used by the Intergovernmental Panel on Climate Change (IPCC). While some of the models do not show a decreasing thermohaline circulation with global warming, most do. So-called hosing experiments, where large amounts of fresh water is poured into the Greenland Sea, show that it is possible to stop the circulation almost completely and to depress temperatures over North West Europe by 4–5°.

Dramatic climate change

Filmgoers may recognise this scenario as the premise behind the extremely rapid climate change in the Hollywood epic, “The Day After Tomorrow”. In reality, we are not talking about an ice age arriving in a few days, but rather a fall in temperature of a few degrees, superimposed on a global warming signal of a similar size. This may not sound much, but a UK with summers 2° warmer and winters 2° colder than at present would be a very different place: our climate would be more like the Eastern USA than our present climate, with problems of rapid adaptation for both human society and the natural ecosystem.

Dramatic climate changes such as this are known as low probability/high impact events. So what is the risk of a collapse of the thermohaline circulation and a consequent rapid climate change? This is where the statisticians come in.

Looking at low probability climate changes has much in common with other climate prediction problems. The main difference is that we are looking at the tails of the probability distributions of our predictions, rather than at means, and we therefore need to be more careful with our statistics. We cannot base our estimates purely on observed data. Although there are examples in the distant past of what we believe are thermoha-

line collapses, they are few in number and are in climate regimes very different to the present day. Thus, we cannot base any estimate of risk on past events. We therefore need to use dynamical models of the Earth's climate to make predictions of future climate states and use these as our data. Realistic climate models have tens of internal parameters and millions of state variables.

Statistical determinism

The climate models we are using are deterministic. If we run them with the same parameters (here we include external inputs with the internal parameters), they will give the same answers, give or take some numerical "noise". To make probabilistic predictions, we need to specify probability distributions for all these parameters and specify a set of initial conditions for the state variables with known uncertainty. This gives us huge problems of dimensionality. The problem of the initial conditions can be tackled by "spinning up" the model. This means running the climate model for so long that the initial conditions no longer matter and the final state depends only on the values of the models parameters. In our work, we are using Bayesian methods for statistical inference. To be fully Bayesian we would need to specify prior distributions for all the parameters. Getting climate scientists to specify a full prior is difficult and we may need to either use non-informative priors or use likelihood based methods.

The statistical analysis of such large deterministic models is known as SACCO (statistical analysis of computer code outputs). The use of statistical methods to design and analyse computer experiments was pioneered in engineering but is now spreading into the environmental sciences, where computer models are of increasing significance.

There are a number of points worth noting with this approach. In this way, by sampling from the probability distributions of the inputs, we can find the distribution of the outputs and, in principle, look at the probability of rare events such as the collapse of the thermohaline circulation. Note that, to use this method, we are only allowing uncertainty on the model parameters, not on the model structure. This is a serious limitation of all current methods for looking at the output of computer models.

Our second difficulty arises because the climate models are very computer intensive. Even the largest computers in the world take many hours to produce a few hundred years of output. There are two current solutions to this problem.

Distribution

One way is to find huge numbers of computers and run jobs on them over a number of months. This is the approach taken by climateprediction.net. Using climateprediction.net you can run a climate model on the PC on your desk. Running as a background job or a screensaver, in a similar way to seti@home, it can run several hundred years in a couple of months. So far, climateprediction.net has run several tens of thousands of climate runs. The model they are currently using does not include a full ocean, so they are not able to look at thermohaline collapse. Although utilising spare computer cycles in this way is an efficient use of resources it is unlikely that we will be able to run the hundreds of thousands of runs we need for Monte Carlo simulations.

Emulation

The alternative is to build what is known as an emulator—a statistical approximation to the full climate model. It is much faster than the full model: in our experiments we get a speed-up of 5 orders of magnitude, and our climate model is a relatively fast one—simulating 1000 years in a few hours. Because the emulator includes a measure of its own uncertainty, we can tell where the approximation is good and where it is less effective.

To build the emulator we run the full climate model in a designed experiment. The usual design is a latin hypercube. This is a design that covers a multi-dimensional space with the minimal number of model runs. An alternative is to use a sequential design, where we choose where to run the model next according to where the emulator uncertainty is large. The emulator is the key to this sort of analysis, as it allows us to use Monte Carlo based methods, which would be unfeasible if we had to run the full model (even using the unused cycles of home computers).

Emulators are only effective where we are interpolating within our original design; extrapolation can produce non-physical results. This issue becomes critical as the dimension increases. For high-dimensional

problems the hyper-volume contained within any design is a vanishingly small proportion of the entire hyper-volume. We are actively looking at ways of reducing the dimensionality of the system. Is there some low-dimensional manifold within the full high-dimensional parameter space within which all plausible climates lie?

Specifying the distributions of climate parameters

To calculate the posterior distributions we have two sources of data. There is the model output—but how close to data are the models? We therefore include data on the real climate. First, there is instrumental data taken from weather stations, ships, balloons, aircraft and satellites. Most of these data come from the last 20 years or so, although there are instrumental records, of varying quality, dating from the nineteenth century and earlier. Although there have been variations in the climate over this period, they have not been dramatic. Small differences in the climate model parameters give big differences in future climates but can still replicate the current climate fairly well.

There is another source of data that may be more useful in specifying the distributions of the climate parameters. Many natural phenomena are related to climate variables. For example, the width of tree rings may be related to rainfall, or the isotope ratios in shells preserved in ocean sediments may relate to sea surface temperature. By measuring these proxies, we can infer (with greater or lesser uncertainty) climatic conditions in times before the advent of meteorological instruments. The traditional way to use such data is to invert an equation describing the way the proxy relates to a climate variable; we are working on methods where this inversion is not required. Although these proxy data are more uncertain than the modern instrumental data, they are from different climate regimes, and so may be able to better define the distribution of the model parameters.

So, what is the risk of a thermohaline circulation collapse? Unfortunately, we are not able to give an answer to that question yet; we are still struggling with the methods described above. We hope to have practical, believable results within the next year or so. Our aim is to produce probabilities conditional on various future emissions scenarios. Such probabilities will enable policy makers to make better decisions on what are “acceptable” levels for emissions post-Kyoto.

Peter Challenor is a statistician and oceanographer at the Southampton Oceanography Centre. He is a member of the steering committee of the Natural Environment Research Council’s Rapid Climate Programme, and is keen to see more statisticians interested in environmental science, particularly climate.

This is a preprint of an Article accepted for publication in Significance © Royal Statistical Society