Numerical Models for Weather Forecasting and Climate Projections

To be able to predict how the weather, and climate, will evolve, requires both an understanding of the physical processes which govern atmospheric behaviour, and a range of observations of what the atmosphere is doing now.

To predict the weather even a short time ahead, it is essential to know the current state of the atmosphere. Regular observations on land, at sea (ships and buoys), in the air (radiosondes – weather balloons, and aircraft) and in space from satellites provide us with information about pressure, temperature, wind speed and direction and humidity. An experienced forecaster can predict the weather a few hours ahead simply by looking at the current state of the atmosphere.

There are several possible approaches to producing a longer term forecast. You could make a 'persistence' forecast; "it will probably be warm tomorrow because it is warm today", or a forecast based on long-term statistics (climate); "it will probably be warm tomorrow because it is usually warm on this day each year". However, a more skilful forecast, requires an understanding of how weather systems develop and the fundamental physical laws which govern the behaviour of the atmosphere, e.g. "it will be warm tomorrow because the high pressure area will persist with prevailing winds from the south".

In order for scientists to predict the future behaviour of a system they require a model of that system. Models in general are simplified versions of reality. Numerical models are collections of interacting equations that we know govern the behaviour of real systems. It may be useful to have a really simple numerical model to help understand or visualise a process – for example what happens when a low pressure system is caught in the wind field of a high pressure system, or they may be collections of hundreds or even thousands of equations, attempting to capture all the processes operating in something as complicated as the Earth's atmosphere. Model predictions can be compared with observations of the real system to evaluate the 'skill' of the model. Models also allow us to look into the future, and to ask 'what if?' in a way that it is not possible to experiment with the real Earth.

In the case of the Met Office in the UK, the same numerical model is used to predict the weather of the next week or so, and to make simulations of past, present or future climate – where climate is defined as being a 30 average of day-to-day weather.

A numerical model used to predict the weather or climate more than a few days ahead must simulate the whole climate system (atmosphere, ocean, cryosphere etc.) for the whole world, as the weather that you are experiencing now may have had its origins thousands of kilometres away a few days before. Such models are commonly refered to as GCMs – standing either for General Circulation Models or Global Climate Models. The surface area of the Earth is about five hundred million square kilometres and the model needs to follow circulations at least to forty kilometres up into the atmosphere.

Five basic equations, known as 'primitive equations' lie at the heart of any numerical model of the atmosphere. These govern the relationship between 5 fundamental properties – wind speed and direction, air temperature, atmospheric pressure, air density. The five equations are:

1. Equation of motion in the x direction, taking into account the forces acting due to the pressure gradient, gravity, friction and the rotation of the Earth.
2. Equation of motion in the y direction, taking into account the forces acting due to the pressure gradient, gravity, friction and the rotation of the Earth.
3. Thermodynamic equation governing the conservation of energy.
4. Conservation of mass, ensuring no air is lost or gained.
5. Hydrostatic equation which is an approximation to the equation of motion in the z (vertical) direction: $\delta p/\ \delta z=-\rho g$ i.e. the rate of change of pressure with height is related to the density of air and the gravitational constant.

State-of-the-art models include many, many more equations e.g. equations that control the conservation of salt and water, the formation, transportation and destruction of ozone, the emission, transportation, incorporation into clouds and eventual removal from the system of atmospheric pollutants etc.

All these equations cannot be solved continuously over the surface of the Earth, but must be discretised. For this purpose, the atmosphere and oceans are divided up into boxes, with each equation being solved once in each box. The box size can vary. The height of each box is usually smallest at the levels of the atmosphere we care most about – usually near the surface of the Earth. The horizontal extent can also vary. Some models define boxes according to a set number of degrees of latitude and longitude (e.g. 2.5º by 2.5º) in which case the boxes will be smaller near the poles, where the lines of longitude converge. A typical current weather forecast model will have a resolution of 50 km over the 500 000 000 km^2 surface of the Earth: in each global layer there will therefore be 200 000 boxes. To get a reasonable approximation to the atmosphere there need to be about 50 layers, giving 10 million boxes.

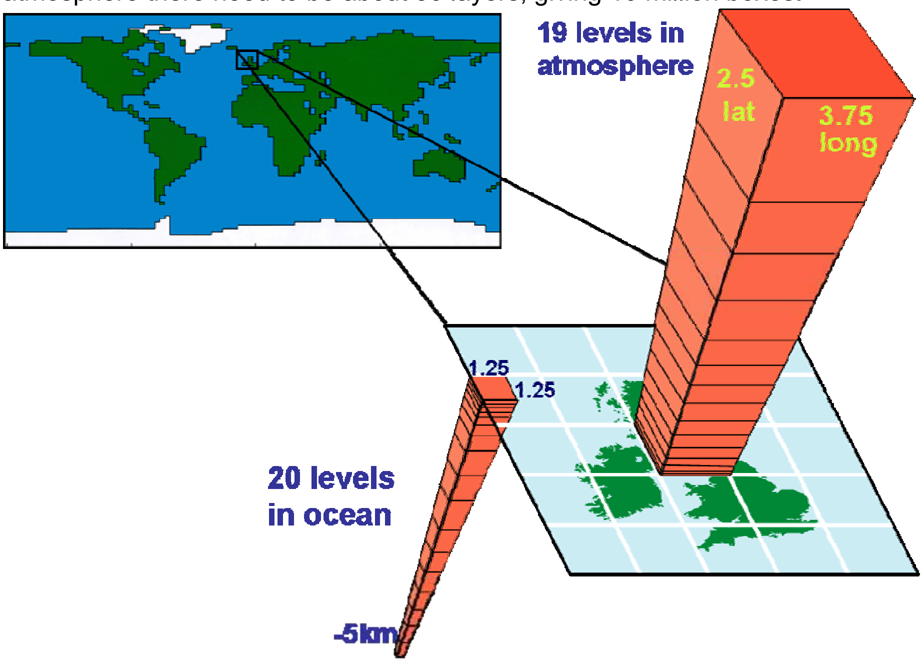

Figure 1: Typical resolution of a GCM used for a climate simulation.

Although GCMs have to be global in extent, it may be appropriate to enhance the resolution (decrease the box size) in a region that you are particularly interested in. So, for example, a climate modeller interested mainly in the future climate of Europe could increase the resolution over the North Atlantic to better capture the processes generating the weather systems that bring European weather. Some very clever models are able to adapt their resolution to better represent key weather features as they develop and move.

Similarly, the equations cannot be solved continuously in time. The models are advanced in time increments – typically half an hour for a climate forecast. So, to build up a forecast, the equations are all solved to simulate the weather in 30 minutes time, then again for 30 minutes after that etc. until a 5 day weather forecast, or 100 year climate forecast is built up. Obviously, you would generate a forecast fastest with a larger time increment. If you want a 100 year climate forecast, why not advance the forecast 24 hours at a time? The answer is not trivial, but in simple terms, you will rapidly lose the accuracy of the prediction if any information (e.g. energy, momentum, water vapour) can be carried (by advection by the wind, or by waves in the atmosphere) further than one model box in one time increment.

The fact that many hundreds of equations have to be solved in several million boxes at each time increment explains why numerical weather models depend on some of the most powerful computers in the world to run them.

There is always a compromise that has to be made between the resolution of the model, and the computing resources available to run the model. Fewer boxes (lower resolution) permit more rapid simulations, but these may not correctly represent important weather features at smaller scales. Therefore, a weather forecaster only interested in a 5 day forecast might have the computing resources to run a model at much higher spatial and temporal resolution than a climate modeller interested in hundreds of years. A model of a given resolution can only hope to describe weather phenomena on length (and related timescales) greater than the size of its grid boxes – so few models can correctly simulate the development of one small cumulus cloud, and climate resolution models cannot even accurately simulate the development of a hurricane. In terms of climate simulations, a rule of thumb suggests that it is misleading to look at anything smaller than the average of 10 model boxes. This means that we can at best hope to understand how the climate of Northern Europe will change, and can say little about changes specific to the UK.

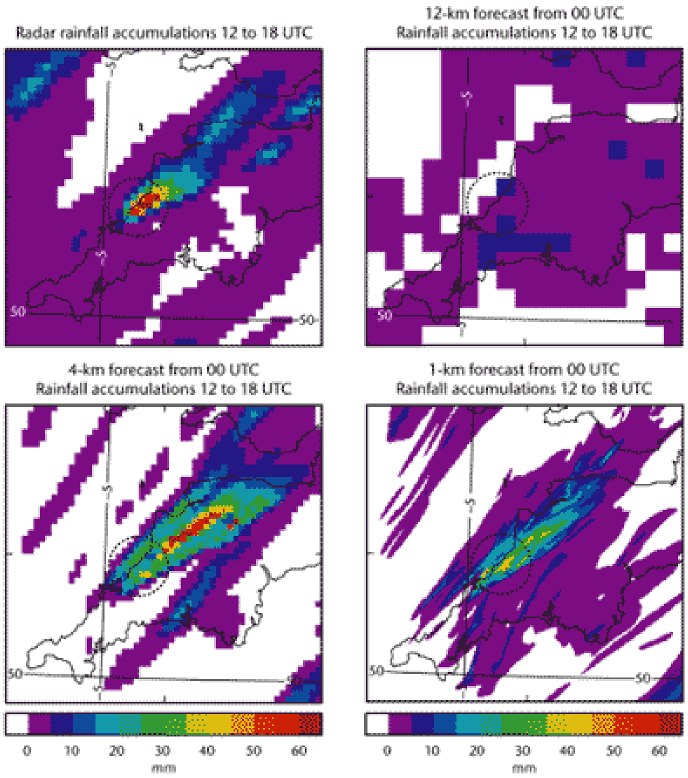

Figure 2: The effect of resolution on forecast rainfall. The top left figure shows observed rainfall (itself at fairly coarse resolution) over a 6 hour period at the time of the Boscastle floods. The remaining images show the results of a forecast initialised 12 hours before the event, using models with 12km, 4km and 1km resolution. Figures obtained from the Met Office.

To make a skilful forecast, three things must be right:

1) The initial conditions

Each time a weather or climate simulation is started, a set of data or 'initial conditions' is required. This is the best guess at the state of the atmosphere, oceans etc. at the starting point in time of the simulation. For a weather forecast starting now, initial conditions can be established much more completely, thanks to satellite and other observations, than they can be for a simulation of the climate of, for example, the last ice age. State-of-the-art numerical weather forecasting models continuously assimilate new observations from satellites to maximise the accuracy of the forecast.

In the early 60s, Ed Lorenz noticed that if he reduced the precision of a number that he supplied his numerical model as initial conditions (specifying three decimal places rather than 6), the model would simulate a very different development of weather patterns. The difference in the simulated weather grew exponentially through the forecast, so that a very small difference in the starting conditions led to completely different weather conditions a few days into the simulation.

This became talked about as the ‘butterfly effect’ – a small perturbation such as a butterfly flapping its wings in the Amazon rainforest could (although it would be very unlikely) lead to a tornado in Texas.
Since the state of the atmosphere, ground surface and oceans will never be known precisely, the weather forecasting problem may seem hopeless, but there are times and places when differences grow more slowly than at others; often when the weather is in a particular pattern. At these times the atmosphere is said to be more ‘predictable’ and there is higher level of confidence in a forecast. Predictability can be assessed by running the numerical model with a set, or ensemble, of initial states, each of which is equally consistent with observations to within suitable errors. This leads to a probabilistic, rather than absolute, forecast. The sophistication of current numerical models is such that five or six day weather forecasts are regarded as having a useful level of skill and a target of around fourteen days is regarded as the maximum which can be achieved.

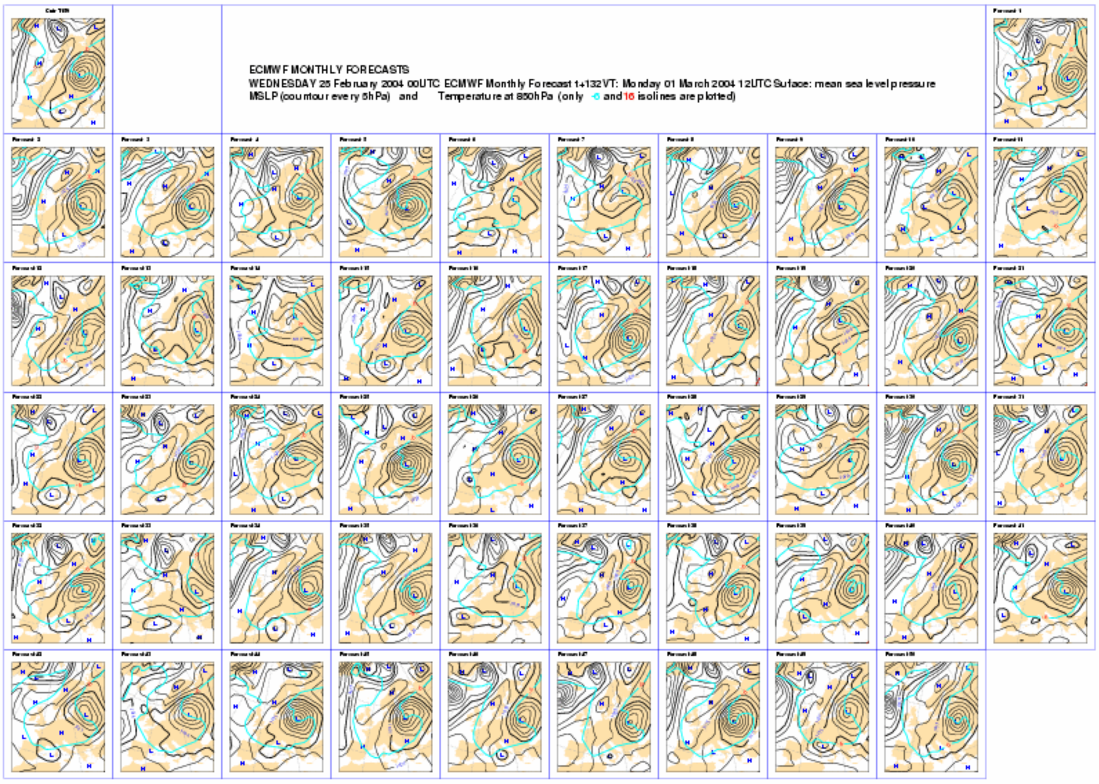

Figure 3: An ensemble forecast of March mean temperature and pressure. Each ‘stamp’ sized image shows one model prediction. Taken together, the probable and possible weather of the month can be established. Copyright ECMWF.

2) The model

Any forecast will be limited by the quality of the model that is producing it. The quality of models can be limited in two main ways: by the completeness of our understanding of the atmosphere, and by the resolution of the model. The coarse resolution of the models means that many processes operating on small and micro-scales cannot be represented explicitly, but must be represented by parameterisation schemes. Such processes include; the detailed absorption, scattering and emission of solar and terrestrial radiation; the effect of small-scale convection and mixing within the atmosphere; the formation of clouds and precipitation; the interaction of the atmosphere with a rough or smooth solid surface, with vegetation or with the oceans; the formation and melting of ice, particularly on the sea; and many other processes, including chemical reactions.

For example, the physical processes operating within clouds are relatively poorly understood, but additional problems arise because the scale of these processes is usually much smaller than the resolution of the model. The effects of clouds on the weather are therefore usually represented by parameterisation schemes.

3) The 'model forcing'

The scientist using a model to produce a weather or climate simulation must specify what will change over the period of the forecast – for example, for a weather forecast, you need to take into account how the amount of incoming solar energy changes as day turns to night. For a climate forecast, information about changes in volcanic activity, changes in atmospheric composition, changes in solar activity, changes in land use, even changes in the position of the continents and the Earth's orbit must be taken into account.

As a result of the factors mentioned above, a scientist with limited computer resources planning a numerical weather simulation must find an appropriate compromise between resolution, complexity, duration of the simulation and ensemble size. Someone planning a 5 day weather forecast would prioritise resolution (so as to minimise the growth of errors in the forecast, and to represent as many physical processes as possible) and ensemble size (so as to be able to get a feel for the accuracy of the forecast). Someone planning a 100 year climate forecast, on the other hand, needs to take into account those aspects of the climate system which vary on longer timescales, such as the land surface and ocean, thereby increasing the model complexity, at the expense of model resolution and probably ensemble size.

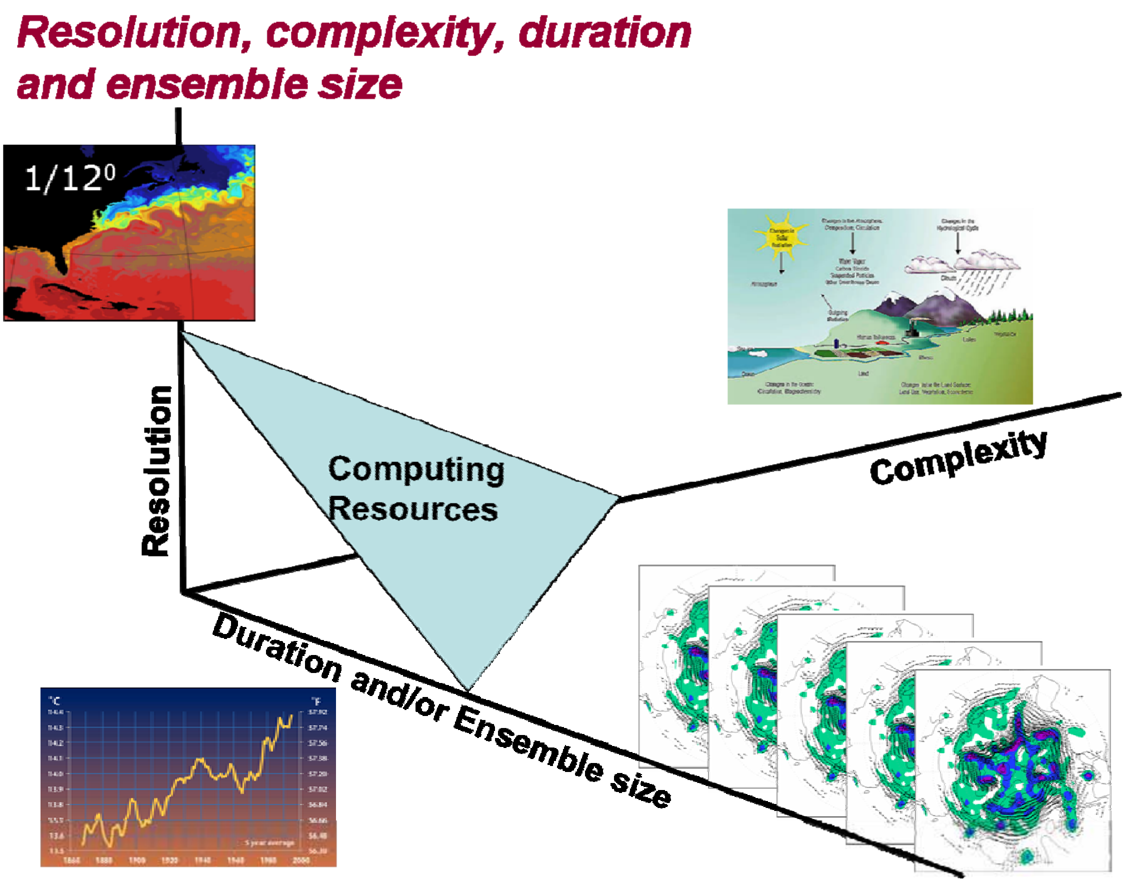

Figure 4: Given limited computing resources, meteorologists must make compromises between the resolution of the model, its complexity, the duration of the simulation and the number of simulations.

The first attempt at a numerical weather forecast was made by a British mathematician, Lewis Fry Richardson. Starting from a set of observations of the atmosphere, he attempted to make a weather forecast based entirely on numerical calculations, using only a slide-rule. His aim was to determine what would happen over the next few hours using only the laws of physics to calculate the movement of air and the redistribution of heat. The attempt was a complete failure as a forecast and took much longer to complete than the time interval covered by the forecast, being published in a book in 1922, a decade after the period to which it related! Consequently, Richardson's book made little impact at the time. Nowadays it is possible to recreate Richardson's forecast in a few minutes on a home computer, showing that, although his sums were correct, his forecast failed mainly because of errors in the initial conditions and because he didn't have the computational resources to optimise his numerical model. The first successful numerical prediction of the weather was completed by an electronic computer in 1950, and weather forecasts have been improving in skill, resolution and duration ever since.

Figure 5: Richardson's Human Forecast Factory: Lewis Richardson dreamed of using thousands of people to be able to produce real time weather forecasts.

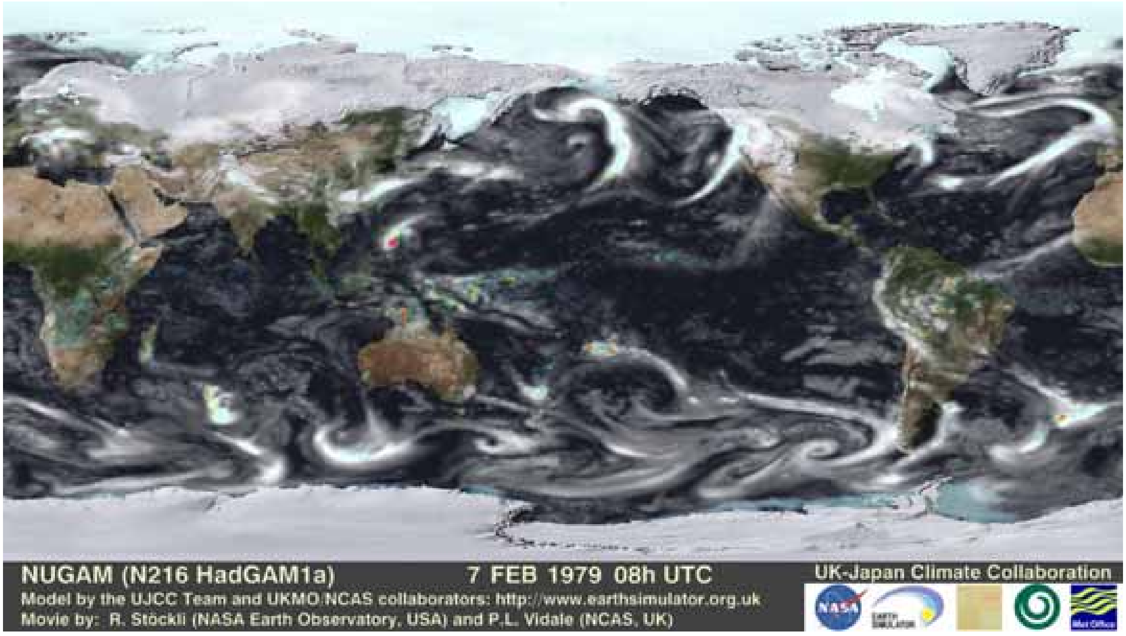

Figure 6: State of the art climate model resolution. This figure was taken from a simulation made on the Earth Simulator in Japan – currently the most powerful computer available to meteorologists.