Weather forecast models and apps

Rupert Holmes unpicks the tech behind the data that underpins all weather apps and highlights his favourite software

NWP (numerical weather prediction) can't resolve small, sharp features such as this squall in the channel between the Greek island of Chios and Turkey

Rupert Holmes

To understand the forecasts provided by different weather apps and websites we first need to have a basic appreciation of the data they present. Obviously they can only be as good as their data inputs, but where does that come from?

The answer lies in the raw data output of numerical weather prediction (NWP) models. We therefore also need to understand the inherent strengths and weaknesses of these forecasts.

As there aren't that many different models that provide free or relatively low-cost data, often different apps or websites are simply presenting the same data in different ways.

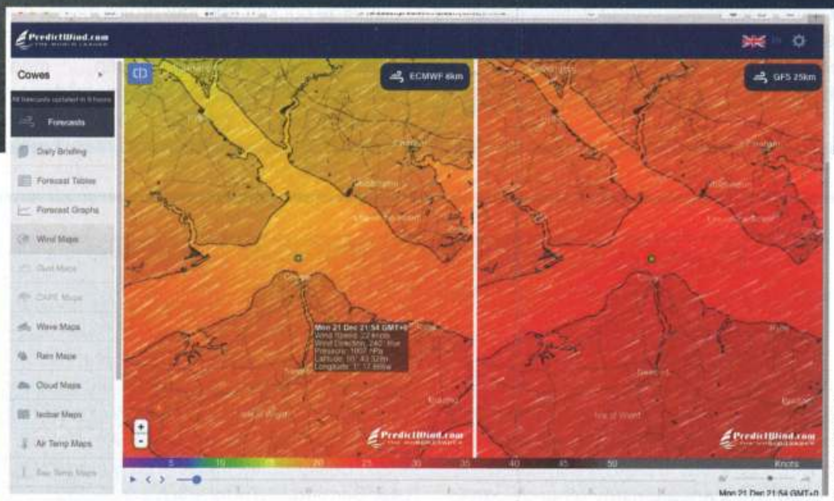

PredictWind

PredictWind's split screen display allows a clear comparison of different weather models

Understanding numerical weather prediction

The key premise behind weather modelling is that the mathematical equations governing motion, fluid dynamics and thermodynamics can be used to predict the future state of the atmosphere.

Key variables within each model include wind strength and direction in three dimensions, air density, temperature, pressure and humidity.

This makes it a hugely data intensive process, which necessitates dedicated use of some of the world's most powerful supercomputers, which are often replaced on two- or three- yearly timescales. Even then, a single model run can take up to six hours to process.

The accuracy of model output depends on many factors including the size of the horizontal grid and the number of vertical atmospheric 'slices' considered.

A feature, whether meteorological or geographical, needs to cover at least three grid points in order to register in the model. Therefore, a model with a 22km grid will ignore weather features that are less than 65km across.

However, there's a trade-off as fine-grained models can't be used for longer range forecasts – the smaller the grid the sooner the output becomes unreliable. Therefore medium-term forecasts from global models tend to use a large grid size, while a small grid is used only for short-term predictions.

Still, it's worth noting that in general grid sizes are becoming smaller, without loss of longer term accuracy.

Weather models

The Global Forecast System (GFS) model currently uses a 28km grid for the first seven days, expanding to 70km for days 8 to 16. This data is freely available, which means it's the most commonly used, especially in free apps.

The ECMWF (European Centre for Medium term Weather Forecasting) global model has a 9km grid and is a common option for paid-for apps.

The German ICON model grid is 13km, reducing to 6km for European coverage.

The MetOffice's Unified Model has a 16km grid, but is a lot more expensive and is therefore seems to be rarely, if ever, used by consumer-facing third party operations. This is disappointing as, along with ECMWF, it is also one of the most accurate.

In addition there are many more models that cover smaller areas in more detail, but for shorter timelines.

The MetOffice short term UKV model, for instance, has a horizontal resolution varying from 1.5km (over the UK) to a maximum of 4km and is therefore capable of resolving features down to 5km across. This model is run more frequently than the global one – once every three hours – but only to 36 hours into the future. Pricing means it's not viable for most app developers to access this model.

'The parameters for the start of each model run are enormously important in achieving a reliable forecast'

Rupert Holmes

The French AROME model (2km grid size, six hour refresh rate) covers most of the UK, other than the far north of Scotland, and consistently produces exceptionally good fine-grained forecasts.

However, to put all that in perspective: neither the MetOffice nor AROME yet offers a sufficiently small grid size to be able to predict local effects in places such as the Needles Channel, Portland Bill or the Menai Straits.

Atmosphere in 3D

As sailors we can easily fall into the trap of thinking of the atmosphere in a two-dimensional 'plan' view. Reality is very different, of course, and it's critical to have a three-dimensional appreciation of the atmosphere.

Weather models consider up to 70 vertical slices, representing altitudes of up to 80,000m for medium term forecasts, and 40,000m for shorter range predictions. In addition, as land masses play a big part in determining weather outcomes, terrain maps are incorporated into the models.

The initial parameters for the start of each model run are enormously important in achieving a reliable forecast, and requires a complete set of up to date observations from around the globe, and at different altitudes.

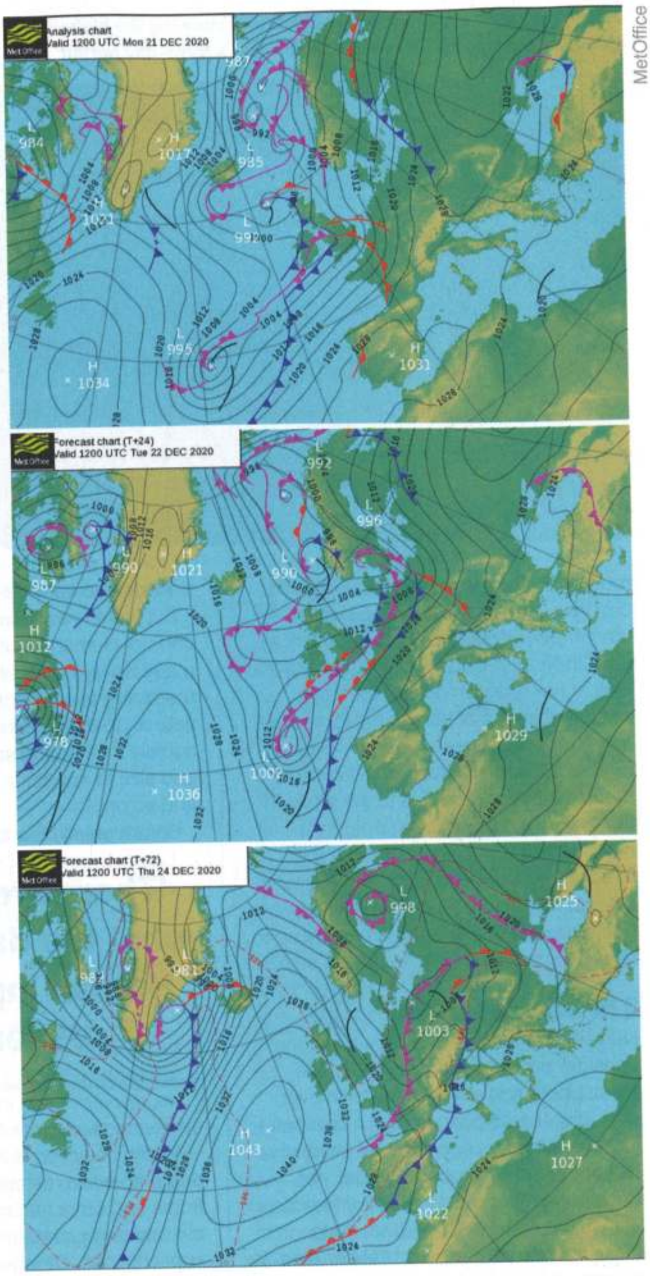

RIGHT Don't discount the importance of analysing synoptic charts to make sense of the bigger picture

MetOffice

Roughly half the computing power used for numerical weather prediction – and a considerable amount of research time – is therefore expended on this aspect. Different accuracy levels in the initial conditions used can also explain the different outcomes between the outputs of various models.

The problem is any small errors in the input state can magnify into much larger errors in the forecast. Fortunately, this can also be used to our advantage thanks to ensemble forecasting. This repeats each model run multiple times, with the initial conditions varied slightly. This gives forecasters an indication of the probability of different outcomes occurring – covering a whole gamut of parameters including wind strength and direction, temperature, precipitation, cloud cover and so on.

The ECMWF's ensemble model, for instance, runs a control based on the main model output, plus 50 further calculations, each with slightly different parameters to the control. If the output of these is tightly grouped, a high degree of confidence can be placed on the forecast. Similarly, the greater the spread of outcomes, the more uncertain the forecast and the more you need to plan for different potential scenarios.

What we're normally presented with is the 'deterministic' forecast. This is the output of a single relatively high resolution model run using unperturbed data.

However, it can sometimes be useful to also see the full complement of ensembles, which can help to visualise how closely matched the majority are to the deterministic output.

Unfortunately, this level of analysis is rarely available to end users, although I've used the full 21 GFS ensembles, as well as the ECMWF deterministic, for

BELOW **Windy's accumulated wind feature, showing the highest gusts expected in the Thames estuary over a three-day period**

windy.com

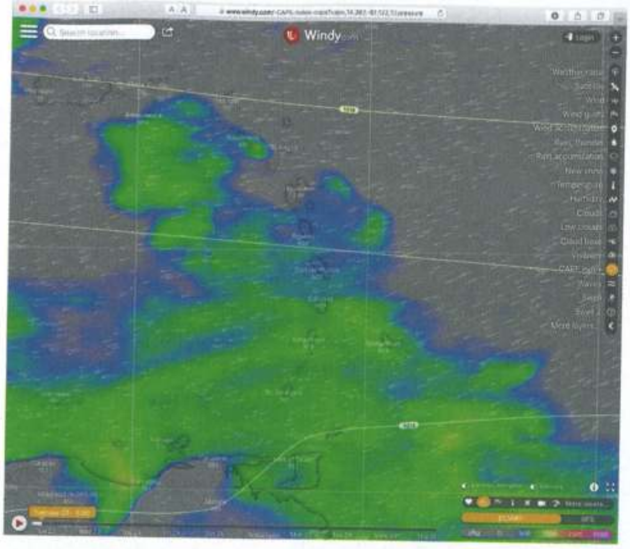

windy.com

ABOVE **CAPE index chart for the Caribbean showing values of up to 800kJ**

routing in offshore races and an Atlantic crossing. We can be sure the skippers in the Vendée Globe are doing similar right now. Indeed, professional forecasters may look at half a dozen different models.

In the absence of being able to view a complete set of ensembles, comparing forecasts that originate from different models makes it possible to get a feel for the likely accuracy of each one. If they all more or less agree, then it's clear the forecast can be relied upon, but if they differ considerably it's wise to keep an open mind, and make sure you stay up to date with developments.

The human factor

Numerical Weather Prediction has transformed our understanding of the atmosphere and ability to forecast over ever-longer time periods. Since its advent just over 40 years ago forecast accuracy has improved by around one day per decade. In other words, 48-hour forecasts at the time of the 1979 Fastnet Race disaster were no more accurate on average than today's six-day forecasts.

Raw model output also gives us tools that make weather forecasting at sea and passage planning far different to only a decade ago for most of us.

But there are still limits to what this data can do. Human input is still essential in many ways: to identify fronts, to predict the timing of wind changes on a race course, to insert discrete features such as squalls that are too small relative to the grid size to appear in the model output, and so on.

Unfortunately, few app providers have the resources to include human input.

'There are limits to what the data can do. Human input is still essential'

Windy, however, is a notable exception and employs a meteorologist to input fronts and other features on the forecasts available with some of its paid-for subscriptions. Otherwise you simply see the rain and windshifts averaged out over time, rather than a more accurate representation of a distinct frontal system.

Raw output from global models can even miss complete weather systems, especially recently developed secondary lows or tropical storms, both of which can be very small in diameter. This was the case during the outbound leg of the 2019 Azores and Back (AZAB) race, where named storm Miguel crossed the fleet as a new secondary depression, before going on to wreak devastation on the Atlantic coast of France. The newly developing storm was mentioned on the GMDSS High Seas forecast. This is compiled by human forecasters and in this case included enough information to plot the track of the storm, allowing competitors to route around it.

This is also a reason to always check the Inshore Waters and Shipping forecasts. In general there are now much better ways to figure out the likely winds for a day's sailing or a longer passage, but they don't have the benefit of expert insight that may highlight a key safety-related issue that's not apparent from the raw model data.

Which app is best?

I often hear people saying they prefer a certain app or website, 'because it's always the most accurate,' but this misses the point that accuracy (or inaccuracy) is not intrinsic to the app. Weather apps and websites are simply interfaces that deliver data in formats that we can easily

BELOW **Photo sequence showing the development of a storm cell**

Rupert Holmes

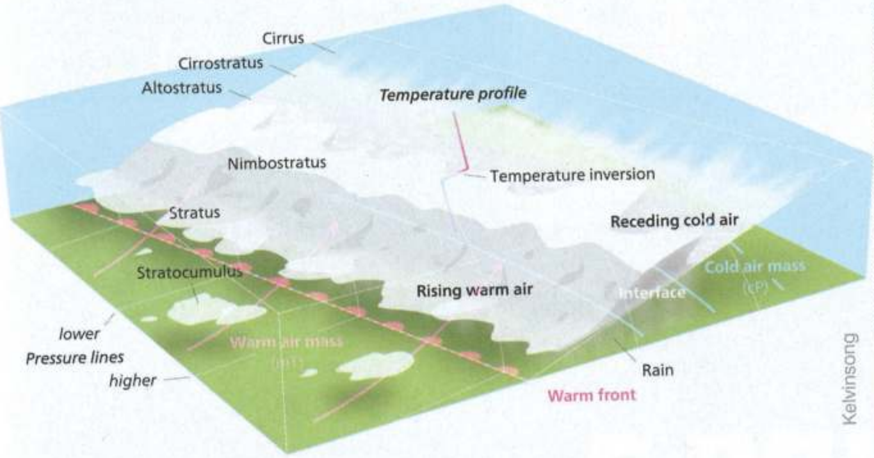

Kelvinsong

ABOVE Try to get used to thinking about this atmosphere in three dimensions – this is a slice through a warm front

understand. But the potential accuracy of that data is entirely dependent on the NWP models used.

Therefore, before using any app the first step should be to find out the sources of the data used. The best services offer a choice of models – global ones for medium term forecasts, and more fine-grained local ones for additional detail in the short term.

So how is this data best used in practice? A few days before sailing I start looking at synoptic charts online – usually from the UK MetOffice and ECMWF. I also look at the MetOffice's spot location forecasts, even though as these are on shore they don't give a full indication of the wind strength offshore.

■ Windy.com

Windy is my top pick for more detail – viewed if possible on a laptop or larger screen – although the company's apps are extremely good. These use a variety of models, including ECMWF and AROME, as well as graphics that clearly show different winds and rain. In addition, the 'wind accumulation' function is useful – it shows the maximum gusts predicted by any model over a range of time periods.

Windy also includes CAPE (Convective Available Potential Energy) forecasts. This predicts the amount of energy in the atmosphere that's available to create thunderstorms and violent squalls. CAPE figures below around 200 are not a worry, but it pays to keep a close eye on developments if it exceeds 800.

■ PredictWind

PredictWind has its roots in Alinghi's 2003 America's Cup challenge in New Zealand. In addition to GFS and ECMWF forecasts, plus its own 50km global models, for a long time the company has offered 1km and 8km predictions in conjunction with the Australian CSIRO CCAM model for 450 popular sailing areas around the world, including many in UK waters. It has also just incorporated AROME modelling on a 1.3km grid for French coastlines, as well as similar high-resolutions models for North America.

PredictWind also includes gust forecasts and CAPE index, as well as offering routing and departure planning options.

■ Squid Sailing

Another favourite, especially among offshore racing teams and delivery skippers, Squid Sailing offers a variety of well-developed tools for both mobile devices and laptops. It also offers a choice of models, including ensemble outputs.

■ RASP

Those who engage in other activities, whether as commercial airline pilots or paragliders, have a different perspective and think of the atmosphere in three dimensions. Using some of their tools can help us too – paragliders, for instance, need to know when updraughts are likely to form, and where to find them.

It's not surprising then that there are apps to help them with this very specific aspect of weather forecasting.

RASP (Regional Atmospheric Soaring Prediction) is an iOS app that offers high-resolution forecasts aimed at glider pilots. Given updraughts over the land are a key part of the mechanism that creates sea breezes, on a near calm summer's morning this information can help us more accurately predict when a sea breeze will fill in and how far offshore.

Although there are plenty of free apps available, a paid-for service can offer many advantages, especially in marginal weather or when needing to identify a weather window for a longer passage.

Interpreting forecasts...

It's generally easy to interpret forecasts in settled weather when there's strong agreement between different models. However, divergence between models need not be a reason to despair – it's often possible to make sense of apparently different forecasts.

The longer the forecast period, the greater the variation in the timing of when specific features such as the passing of a low pressure system. A 12-hour difference could mean one model predicting a dull rainy day, while another forecasts a wet night followed by a crisp sunny day. Both are in agreement on what's going to happen, but disagree about the timing.

That's useful information that helps us plan (while crossing fingers that the rain does actually fall overnight, leaving a sunny day in its wake). PBO

Top weather tips

1 Always start by looking at synoptic charts. This will quickly give you a broad-picture understanding that you can then build by studying finer grained data online or using apps.

2 Make sure you know which model's data you're looking at in any app or website.

3 When comparing data from several different apps beware of the tendency to cherry pick forecasts that best suit your plans. Instead, consider all the possibilities, including those that are more unlikely, so that you can make contingency plans if necessary.

4 Raw model output tends to underestimate wind speeds over the sea – adding 20-50% to wind speeds is often prudent.

5 Look at gust data – not just mean wind speeds.

6 Don't ignore Inshore Waters, Shipping and High Seas forecasts – they may include important messages that aren't apparent when looking at an app.

7 Remember to also look at wave forecasts – generally speaking leisure boats are more likely to encounter problems caused by sea state than by the sheer force of wind alone.