ECMWF Feature article

from Newsletter Number 121 – Autumn 2009

METEOROLOGY

Tracking fronts and extra-tropical cyclones

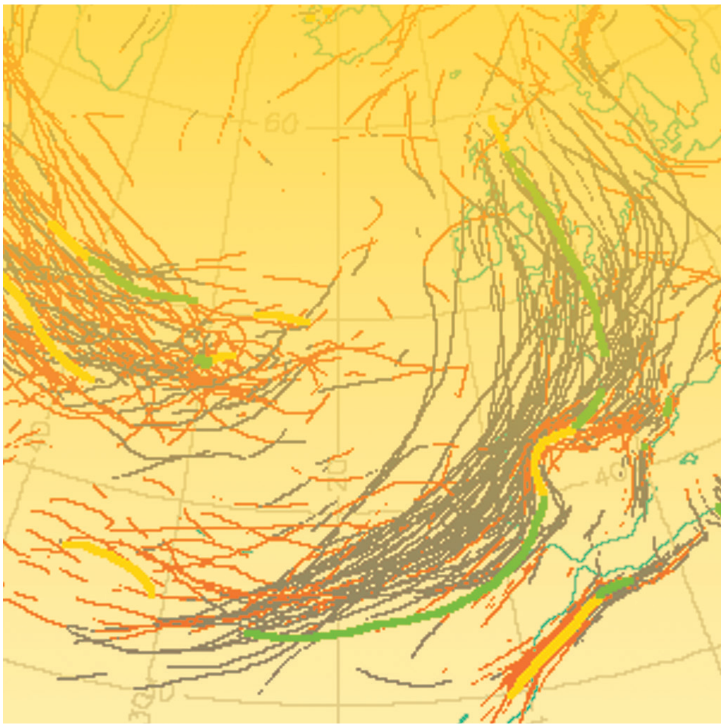

www.ecmwf.int/en/about/news-centre/media-resources

doi:10.21957/un8liw3vjs

This article appeared in the Meteorology section of ECMWF Newsletter No. 121 – Autumn 2009, pp. 9–19.

Tracking fronts and extra-tropical cyclones

Tim Hewson

Forecasters have long recognized the merits of tracing the evolution of meteorological 'features' in the atmosphere – originally in observational data, and increasingly, nowadays, in model output. The motivation for following such features is that, directly or indirectly, they correlate with significant and sometimes severe weather. For example, line convection, a feature seen on radar images, directly relates to heavy precipitation, reduced visibility and squally winds, whilst upper-level PV (potential vorticity) anomalies, features seen in model output, can be an indirect precursor to rapid cyclogenesis and all the adverse weather that entails. The two features most widely recognised in the extra-tropics are, arguably, fronts and cyclones. Analysis of these has historically been a time-consuming manual process; in operations, when there was but one model run to deal with, this was just about tractable, but nowadays, in the era of ensembles, it is plainly not.

Conveniently, mathematical algorithms have been built up, over a number of years, to identify and track such features in an automated fashion. Initially these algorithms were developed at the Met Office and applied to MOGREPS (Met Office Global and Regional Ensemble Prediction System) data. Related output is now part of the suite of operational real-time products provided to Met Office forecasters. More recently this code has been upgraded and applied to forecasts from ECMWF's Ensemble Prediction System (EPS) to develop products for ECMWF customers. This article describes those products and their origins. We also provide advice on how to use the products in an operational setting, in conjunction with guidance from the higher-resolution 'deterministic' model.

The key benefits for the user of this new feature-based approach to ensemble processing are summarised below.

- Identifying where warm fronts, cold fronts and cyclonic features of different types are in the ensemble, how these propagate with time, how uncertainty (i.e. positional spread) develops in their handling and how this uncertainty varies, on a given day, across the domain of interest.
- Seeing whether the deterministic and control runs provide 'mid-range' solutions, in their handling of fronts and cyclonic features, relative to the rest of the ensemble.
- Visualising 'synoptic-chart-style' animations that show fronts, cyclonic features and mean sea level pressure, for all the model solutions.
- Distinguishing between cyclonic features of different synoptic types, such as 'frontal waves', 'diminutive frontal waves' and 'barotropic lows', and also distinguishing cold front features from warm front features.
- Visualising the evolutionary history, in the deterministic and ensemble solutions, of a particular cyclonic feature in a 'feature-plume' format. Components represented include 12-hour movement vectors (i.e. tracks), mean sea level pressure at the feature point, and low-level wind maxima within fixed radii of the feature point.
- Visualising, as animations, storm track strike-probabilities using three different thresholds for feature intensity. This allows 'the potential for cyclonic activity' and indeed 'potential storminess' to be diagnosed out to 15 days.

Fronts

Both forecasting practice and theoretical ideas indicate that warm and cold fronts should mark the transition zone between airmasses of different thermal characteristics. Such fronts thus denote regions of large thermal gradient. For consistency with the typical vorticity signature – as denoted at the surface by a wind shift – fronts are actually identified *along the warm air boundaries* of the regions of large gradient.

Two key choices have to be made when identifying fronts – which thermal variable should one use and on which atmospheric level or levels? Following extensive testing, which involved comparison of objective fronts with fronts found on manually-produced synoptic charts, the choice was made to use wet bulb potential temperature (θ_w) on a terrain following co-ordinate that is 1 km above the model topography.

 doi:10.21957/un8liw3vjs

θ_w has valuable conservative properties; in particular it tends to give continuity to frontal progression across elevated topography. A level of 1 km is reasonably representative of the lower troposphere and has a close connection with surface weather, as required, but at the same time is not 'contaminated' too much by the underlying surface. One important aim was to avoid seeing semi-permanent fronts around many coastlines – at levels lower than about 1 km this becomes an issue. A pressure level approach was not adopted because surface processes and sub-surface extrapolation would have had a contaminating effect around high topography.

To plot the fronts a number of diagnostics, based on horizontal derivatives of the chosen thermal field, are computed; these are then plotted using standard graphical algorithms. The fronts themselves are merely a contour plot on which some contour segments have been erased. The red and blue colouring, for warm and cold fronts respectively, is dependant on another variable; here the sign of the geostrophic thermal advection is used. This is broadly consistent with forecasting practice in Europe.

Comparison of objective and manually-analysed fronts

Automated and manual analysis charts for one time have been interlaced on Figure 1 to enable synoptic feature positions to be compared. Manual fronts are shown conventionally, objective fronts are in colour: good general agreement is immediately apparent. There are some discrepancies however. One concerns smoothness; this is much greater for the manual fronts. An analyst will deliberately smooth out fronts to provide a product that is more aesthetically pleasing, though may at the same time sacrifice some useful local detail. For the automated product there is only slight smoothing of the input fields. A second point is that manual occlusions tend to be denoted by objective warm fronts. Thirdly, note that a few fronts, both manual and objective, do not have counterparts. This will be partly because an analyst can take into account additional factors, such as cloud bands on imagery, or visibility gradients, when identifying fronts. As discussed above, the objective method uses only wet-bulb potential temperature, in the expectation, from experience and frontal theory, that these other delineations will be broadly coincident most of the time.

New animated web product – 'spaghetti fronts'

Figure 2 shows example frames from an animated web product being developed for users, which we call a 'spaghetti fronts' animation. It depicts objective fronts in all members, at different lead times, from one EPS run in summer 2009. The evolution from a coherent pattern, with good inter-member agreement at short leads (Figure 2a) to a random-looking pattern by long leads (Figure 2c) is typical. At long lead times only climatological aspects – such as the absence of fronts in the tropics – tend to stand out. What we are effectively doing here is representing trends in the EPS spread in a synoptically meaningful way (i.e. as they relate to frontal positioning). Questions the user can address with this product include the following.

- How confident are we in the positioning of different fronts at particular lead times? In Figure 2a confidence in the position of the warm front southwest of Greenland is evidently higher than confidence in the fronts northwest of Norway.
- How representative of the ensemble are the control and deterministic runs? Control run fronts are always over-plotted in green and gold. Figure 2b indicates that a cold front will likely be crossing Northwest Europe around T+180 hours; the control run seems to be mid-range and so should be 'reasonably representative' of the ensemble for this particular feature, and by implication for the general weather in this region around this lead time.
- At what lead does making any sort of deterministic prediction become futile? One could argue that, away from arctic regions, there is some confidence in identifying airmass boundaries at T+180 (Figure 2b), justifying a modicum of determinism in forecasts issued for that time. Conversely, by T+360 (Figure 2c) the confidence in where the airmass boundaries will be is virtually nil.

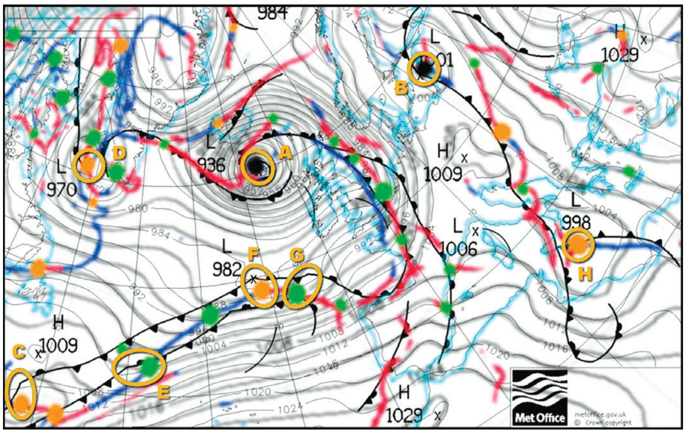

Figure 1 Met Office subjective (black) and ECMWF automated (grey and colour, control run) synoptic analysis charts for 12 UTC on 22 January 2009, blended for comparison. On the objective charts warm fronts are red and cold fronts blue, whilst diminutive waves, frontal waves and barotropic lows are shown, respectively, by green, orange and black spots. Weaker frontal features are shown by smaller spots. Rings highlight cyclonic features that are equivalent on the two chart types.

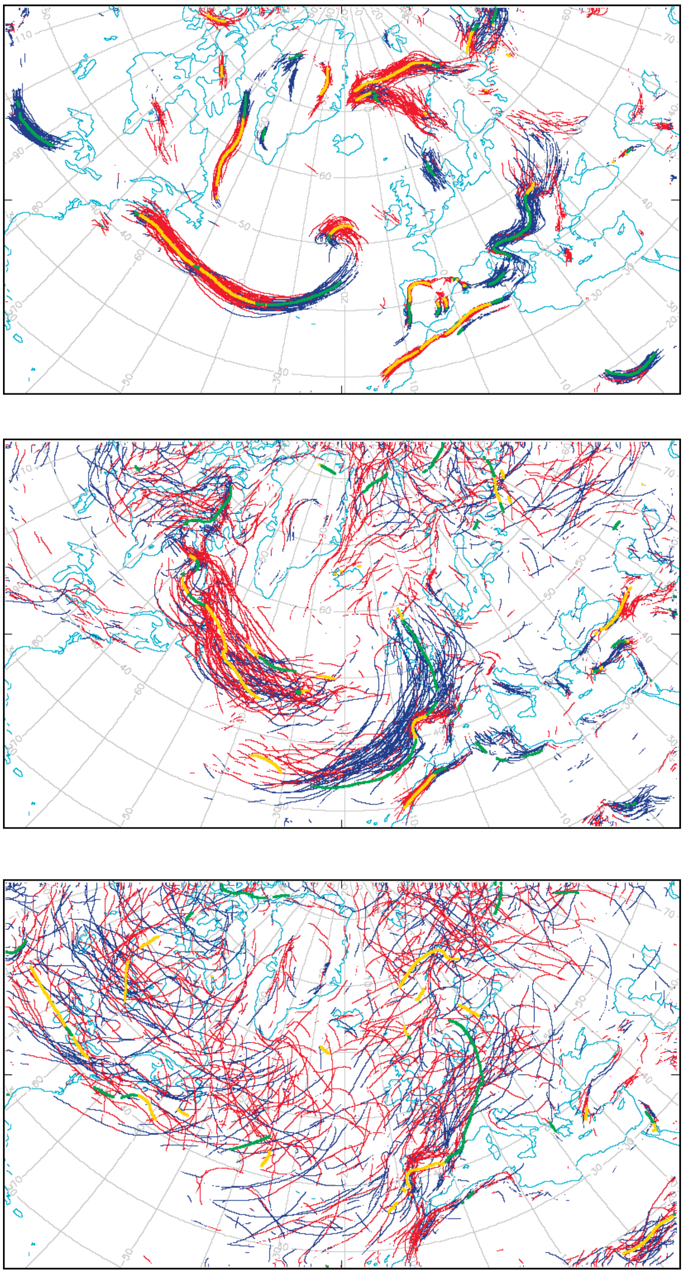

Figure 2 Objective front spaghetti plots, from one set of ensemble runs, data time 00 UTC on 23 July 2009. Lead times are (a) 60 hours, (b) 180 hours and (c) 360 hours. Control run fronts are shown in gold (warm) and green (cold). Fronts from the deterministic run are not shown here, but will be incorporated onto standard products in the near future (up to T+240).

doi:10.21957/un8liw3vjs

Cyclones

Feature types

Cyclones in the extra-tropics can take on many forms, but commonly they start out as a small cyclonic disturbance on a pre-existing front. Most forecasters would refer to this as the *'frontal wave'* stage. On a synoptic chart such features lie at the meeting points of cold and warm fronts, with the extra condition that the implied frontal rotation must be in the correct (cyclonic) sense.

Recent work has identified a new type of frontal cyclonic feature, which has been named a *'diminutive wave'*. This is typically a very minor disturbance on a front, which, in the cyclone life-cycle, tends to precede the frontal wave stage. It can be recognised, synoptically, by a slight opening out of the isobars along a front. Dynamically, this opening out signifies the development of a local positive-negative couplet in the vorticity of the cross-front geostrophic wind; it is the maximum in this quantity – the positive part – that pinpoints the diminutive wave (see also Box A). If a diminutive frontal wave develops further, then a frontal segment on one side will change to a different (warm or cold) type, and at this point it will have transitioned into a frontal wave. Then if the frontal wave further intensifies it will eventually lose its thermal (frontal) signature as it evolves into a non-frontal low pressure centre. This we refer to as a *'barotropic low'*.

The dynamical significance of diminutive frontal waves — A

At first sight diminutive frontal waves seem to be such small features that one might imagine that they are dynamically inert. Conversely the fact that many can be tracked in a coherent fashion, and that some develop substantially, suggests otherwise. One ensemble example of a modest development that could be tracked in time, and that led to severe weather, is shown in the figure.

In terms of dynamical forcing, the locally splayed out isobaric pattern that one tends to see around diminutive waves (e.g. stage 1 in Figure 3 and the inset in the top-left panel of the figure below) can be interpreted in two ways.

(a) This pattern is akin to a deformation pattern that, when acting on isotherms that are typically front-parallel, will be frontogenetic. Such patterns are associated with forced ascent. For a diminutive wave, where the along-front extent of the deformation pattern is limited, there will preferentially be forced ascent just in the region of the deformation (e.g. area 'A' on the figure). Other things being equal this would lead, for example, to higher precipitation rates near to the associated frontal segment.

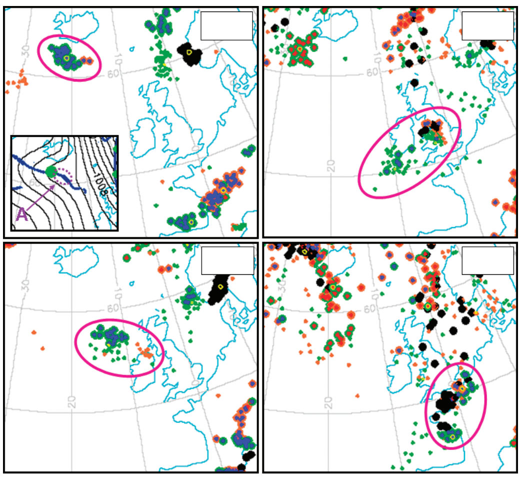

(b) For a warm front diminutive wave there will be a local maximum in warm advection on the high pressure side of the wave. For a cold front diminutive wave there will be a local minimum in cold advection on the low pressure side. In a standard quasi-geostrophic omega equation treatment of dynamical forcing, the Laplacian of the thermal advection field represents one part of the forcing for ascent. In these cases the maximum in warm advection and the minimum in cold advection both contribute positively to this Laplacian term, and thereby also contribute positively to locally forced ascent. So again, other things being equal, one would expect higher rain rates near to the respective front segments.

It is tempting to also try to apply the above arguments to the case of cold front waves (see Box B), though the complicating effect of variations in along-front thermal gradient, particularly for more developed features, makes this less appropriate.

The lesson is that forecasters should not readily dismiss diminutive waves as insignificant, inert features. Some, certainly, will turn out to be, but others can be very important even if they do not develop markedly.

Synoptic cyclonic features. A cold front diminutive wave in the ensemble (green spots with a blue centre) moves southeast towards the UK, and in some members develops further into a cold front wave (orange spot with blue centre) and/or a barotropic low (black spot). This feature gave rise to a very rare early snowfall over Southeast England on the evening of 28 October 2008 (between about T+66 and T+72). Top left inset shows the automated synoptic chart from one member; 'A' denotes a region of inferred forced ascent.The legend in Figure 5 gives the meaning of all symbols.

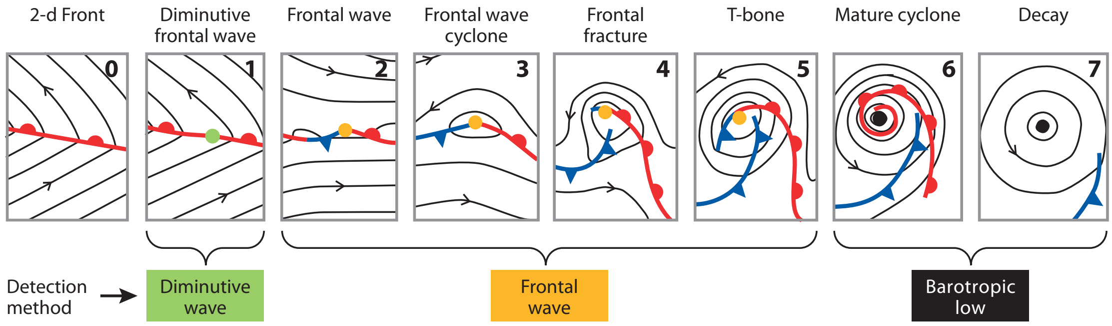

Figure 3 Conceptual model of the life-cycle of a cyclonic feature that develops on a warm front. Stages 3 to 6 come from the Shapiro-Keyser conceptual model; others have been added. Spot colours signify the detection methods used for the different stages.

Figure 3 depicts a conceptual model of the life cycle of a cyclonic feature that initially forms as a diminutive wave on a warm front, then passes through the stages mentioned above, before eventually decaying. A similar model could portray a feature that initially develops on a cold front.

The above discussion suggests that a 'typical' cyclone life-cycle in the extra-tropics can consist of a number of stages. At the same time, however, it is important to recognise that most cyclonic features will not evolve in the way shown by Figure 3 – for example, some will decay at a very early stage. Others meanwhile could start out as a barotropic low – given, say, large upper-level forcing in the absence of surface fronts. The identification and tracking strategies detailed below cater for all such eventualities.

Feature identification

In order to objectively recognise and track, in model output, cyclonic features, and at the same time make a strong connection with synoptic forecasting practice, separate sets of mathematical algorithms have been developed to identify each of the three feature types (annotated at the foot of Figure 3). On the automated synoptic charts these cyclonic features are denoted by spots (this approach was used for Figure 1 and the same convention is used on charts referred to later in the case study). Post-processing forces there to be a minimum separation of ~300 km between features and concurrently enforces a decision hierarchy: frontal waves take precedence over barotropic lows and diminutive waves, and barotropic lows also take precedence over diminutive waves.

Comparison of objective and manually-analysed features

On Figure 1 agreement between objective and subjective cyclonic features (see rings) is generally good. On the automated chart barotropic lows (A, B) tend to be in exactly the right place, and frontal waves are usually close to their counterpart (C, D, E, F, G, H). However there are some differences between the charts. The most striking is that there are rather a lot of unmatched objective features, most of which are diminutive waves. One reason for this is that as yet the synoptician has no symbol available to denote such a feature. Occasionally, if they think a cyclonic disturbance is developing, a frontal wave can be 'concocted' on a front in between the isobars, where arguably a diminutive wave should have been shown – waves E and G are perhaps cases in point.

Feature C on Figure 1 denotes the beginnings of windstorm 'Klaus', which hit France and Spain 1–2 days later. Analysis in the vicinity of such systems can be difficult, so it is reassuring that both objective and subjective means picked this up.

Feature tracking

After identifying where the cyclonic features are at each time, we next want to connect up associated features on successive charts to construct tracks. This is a non-trivial process and requires accurate estimates of 'association probability' to derive meaningful tracks that the synoptician will believe. The first stage of the process involves progressing and retrogressing features on consecutive charts, to try to meet up at 'half time'. The translation vectors applied are based on a 'steering wind' (equal to a fraction of an upper-level wind vector) and on previous movement if available. The association probability is then calculated for all possible feature pairings. This depends on three parameters:

- Feature separation at 'half-time'.
- Type transition.
- Thickness change (1000–500 hPa) at the feature point.

doi:10.21957/un8liw3vjs

In short, each of the following increases the association probability: close matches in space, type transitions that historically (in a manually-analysed training period) have been relatively common, and small thickness changes. Ultimately, through an iterative process, those pairings that have the highest association probabilities are 'matched'; thence they form part of a track, though at the same time care is taken to not track two features into one, and vice versa. In addition if association probabilities for a given feature are all too low then potential matches are discarded, which ensures that tracks have both start and end points.

Tropical cyclones

Though initially developed for the extra-tropics, our algorithms also succeed in locating and tracking tropical cyclones, which are classified as barotropic lows. Tracking works in the tropics, and also through 'extra-tropical transition' should that occur. During such a transition a cyclone typically evolves into the frontal wave feature type.

User products

Once we know the feature points and the feature tracks, there are many ways that this information can be usefully conveyed to the forecaster. The simplest way is to just display automated synoptic charts, in either postage stamp or animation format, and products like this are indeed a key part of the output (an example is given in Figure 4 that is used later in the case study). A recent innovation, to complement the animated front 'spaghetti plots', is an animation that shows just the feature points from all the EPS members. We call these 'dalmatian plots' (see the figure in Box A, and also Figure 5 used in the case study).

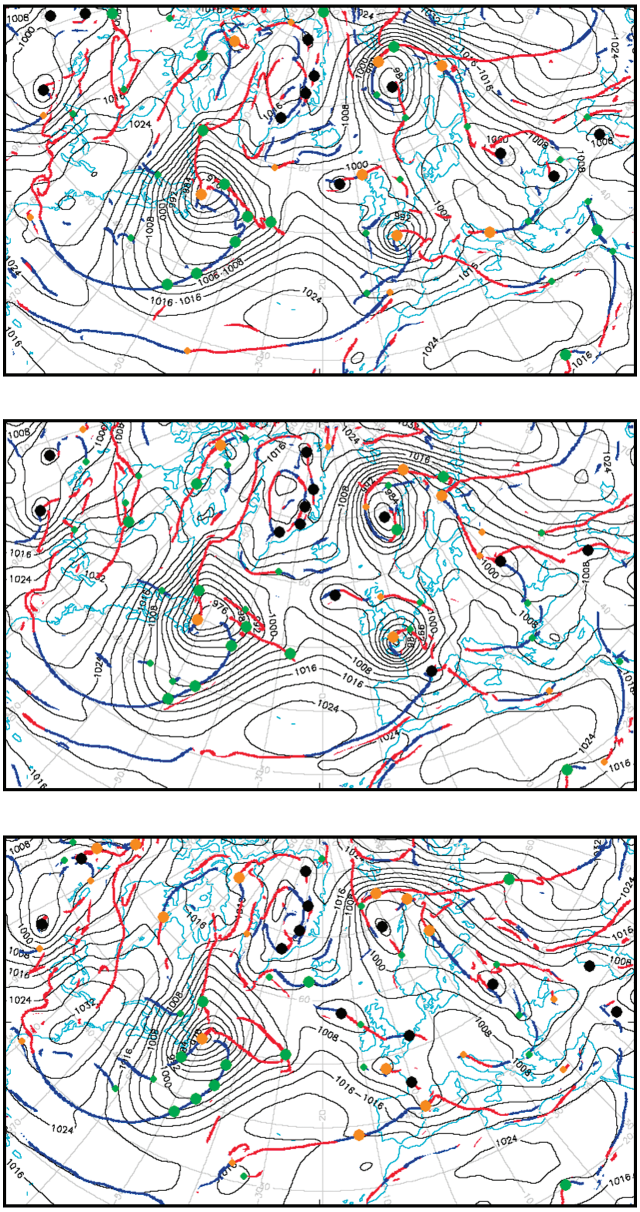

Figure 4 'Automated synoptic charts' at T+120 (00 UTC on the 10th) from data time 00 UTC on 5 February 2009 for (a) deterministic run, (b) control run and (c) ensemble member 21. Colour convention is as specified in Figure 3 and the mean sea level pressure is in black. Weaker frontal features are denoted by smaller spots.

Two different strategies are adopted for presenting track information to the user. Which one is most appropriate depends on the lead time. At time zero, it is usually very easy to cross-reference features in the members; in turn this enables ‘feature-specific plume diagrams’ to be put together, showing how a particular synoptic feature, and various characteristics that it possesses, evolve within the ensemble (see Figure 7 used in the case study). For upcoming significant (cyclone-related) weather events this approach aims to usefully portray the expected life-history of the responsible feature.

At longer leads, as the ensemble spreads, cross-referencing becomes increasingly difficult, and so instead we create storm-track strike-probability plots. In this approach all the tracks, of all the features, are divided up into three sets. The first set – ‘all features’ – contains every track. The second set – ‘stronger features’ – consists only of those features for which the wind within their circulation exceeds a certain threshold. The third and final set – ‘storms’ – is like the second set, but with a higher threshold. The final product comprises a separate animation, for each set. Using colour shading each of these shows the probability that a cyclonic feature, around which the wind exceeds a certain threshold, will pass within a distance of 300 km within a time window of ±12 hours of frame time (see Figure 6 used in the case study).

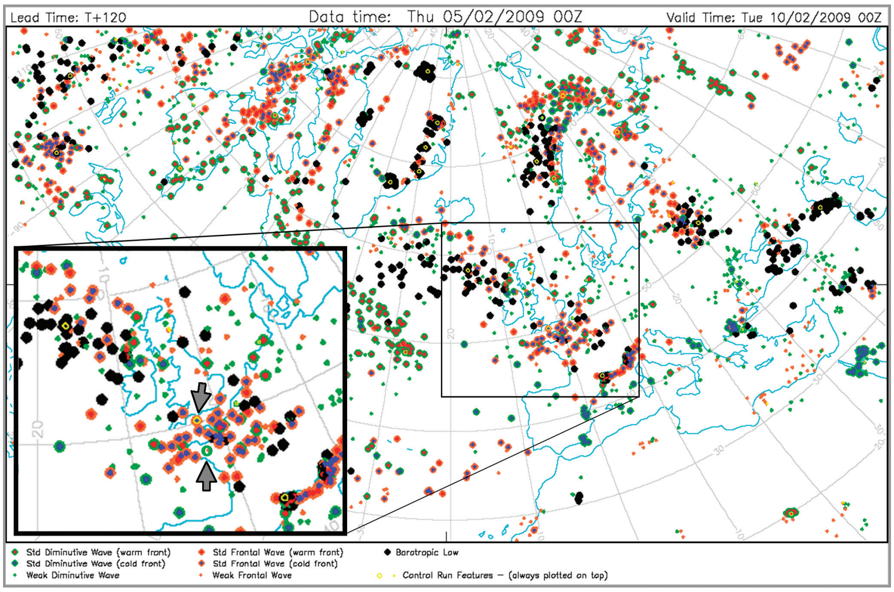

Figure 5 ‘Dalmatian chart’ showing positions of all synoptic features, in all EPS members, at T+120 (00 UTC on the 10th) from data time 00 UTC on 5 February 2009. Symbol style and colour denote type of synoptic feature as in the key below. A portion of the chart has been enlarged, with the position of the main feature in the deterministic run added - see green ring. Other features from the deterministic run are not shown here, but will be incorporated onto standard products in the near future (up to T+240).

 doi:10.21957/un8liw3vjs

How to use the new products – a case study

Here we present a case of 'potential severe weather'. Consider a forecaster who has to prepare forecasts for a Monday night, 9–10 February 2009, for France and the UK.

Thursday 5 February

First imagine that it is Thursday the 5th, and that the 00 UTC products have just become available. It is assumed below that the forecaster would naturally also incorporate information from other ECMWF products and other models as necessary.

What do the deterministic and control runs show?

First there is merit in looking at the automated synoptic chart sequences from the deterministic and control runs. Figures 4a and 4b show snapshots from these for the night in question. There is evidently potential for a deep cyclone to develop, bringing severe winds on its southern flank, as well as heavy rain generally, and also snow on the northern flank if the airmass there is cold enough. Both animations show the cyclone developing rapidly from a cold front wave, though the control run provides the most intense solution.

How typical are the control and deterministic runs?

To address this question we can refer to the 'dalmatian charts', showing where cyclonic features are at different leads. Figure 5 shows the T+120 chart. Evidently a high proportion of members have cyclonic features over northern France and the English Channel. Most are standard frontal waves; some are barotropic lows. Mostly the frontal waves are classified as 'cold front' features, suggestive of strong winds on the western and southwestern flank, as is commonly seen with destructive European cyclones. Box B explains this connection.

The deterministic run feature (shown as a bright green ring and arrowed on Figure 5) is clearly at the southern edge of a feature cluster, whilst the control run (in yellow and also arrowed) is at the northern edge. Thus in this instance, presuming that we gave equal weight to deterministic and control, and some weight also to each of the other members, the feature would be forecast to lie slightly to the east of a point midway between the two arrowed features. In this way dalmatian charts can be particularly helpful for constructing a 'consensus forecast chart' for a particular time.

Cold front wave or warm front wave? B

For both frontal waves and diminutive waves the identification system is able to classify whether they are cold front features or warm front features. So how is this done?

For the diminutive waves it is trivial, as only one objective front type is involved. For the frontal waves however classification is more involved – it depends on the front-normal geostrophic flow 200 km out from the feature. If this flow, on the cold front side, is greater than it is on the warm front side then the feature is classed as 'lying on a cold front' and vice versa. This broadly accords with synoptic practice. The asymmetry implied in this distinction can be helpful to the forecaster – stronger flow on the upstream side, as with cold front waves, can relate to interaction with a trough and/or PV anomaly at upper levels, and incursion of a dry slot on water vapour imagery. Damaging surface winds can occasionally result.

The dynamical explanation for this behaviour is that upper troughs generally exhibit positive vorticity advection ahead and negative vorticity advection behind. In turn, alluding to the omega equation, we see that this relates to a forcing couplet with forced ascent ahead and forced descent behind. The close juxtaposition of these different forcing regions can help enhance the pressure gradient at the surface, and so as an upper trough catches up with a low, it is commonplace for the aforementioned cold front wave characteristics to develop.

Meanwhile warm front waves (which a dynamical meteorologist might refer to as a 'diabatic Rossby wave') tend to be more confined to the lower troposphere, are more driven by warm advection forcing, and though they may produce copious amounts of precipitation, tend not to develop rapidly. The strong warm advection forcing relates to the stronger flow on the warm front side.

Thus the strong connections between surface synoptic feature type and upper-level interaction should make the provision of classification information to the user, via the 'dalmatian charts', all the more useful to the forecaster.

What is the likely track of the cyclone, and what is the probability that destructive winds will develop?
Whilst many members seem to be showing features that closely resemble the control and deterministic runs, we do not know from the dalmatian chart exactly how many there are, nor their intensity, nor their direction of movement (though animating can help). So the next product to examine would be a storm track strike probability plot for the highest threshold level ('storms') – see Figure 6. This suggests 80% probability of a windstorm-inducing cyclone tracking eastwards in the vicinity of northern Brittany, which could be sufficient justification for warning issue for the parameters discussed above. In general if the control and/or deterministic runs are outliers, then this probability would be less valid, though from Figures 4 and 5 that was clearly not the case here.

What are the other scenarios?
For completeness one should also examine other possibilities. This can be done, at a glance, by referring to synoptic postage stamp charts for T+120 (not shown). This process showed a number of members in which cyclonic development was much more muted. The user can click on these postage stamps to see a frame in more detail – Figure 4c shows one of the weaker members. Note the lack of development, and the stark contrast, at least in terms of implied winds, to Figures 4a and 4b. Issued forecasts could refer to this possibility.

Sunday 8 February
Next we will we move forward in time and imagine that it is Sunday evening with products from 12 UTC on that day now available.

Is there a change in the evolution?
We are now sufficiently close to the time of a potential storm that the responsible feature is apparent in both observation and model data. This means that we can use the 'feature specific plume' facility. This is accessed on the web site by clicking the feature spot on the control run T+0 frame (top-left inset on Figure 7, arrowed). This brings up another web page, shown in the body of Figure 7.

The top-left panel of Figure 7 (which can be animated) shows the tracks of the feature in the ensemble at 12-hour intervals, as well as numbers of members in which the feature was tracked for the first 96 hours. There now looks to be a strong signal for the feature to track further north than previously expected, across the far south of England. Note also that the control run track, plotted on top in pink/green, is now mid-range, suggesting that its evolution (or that of the deterministic run, which was almost identical) could usefully provide deterministic guidance should that be required.

What do the feature characteristics indicate?
Feature characteristics are shown by the smaller panels on Figure 7. Again there is quite good consistency in central pressure (top-right, spread is ~ ±5 hPa), and maximum 1-km winds in the circulation (within both 300 km and 600 km radii; lower-left and lower-centre panels respectively). Note also that within a 600 km radius the maximum winds exceed, slightly, those within a 300 km radius. This implies (a) a large system that is having an impact well to the south, and (b) near to the track winds may not be as strong. The fact that the low-level vorticity (centre-right panel) reaches a maximum before the pressure reaches a minimum may also signify an expanding slacker core to the system. This information can all be helpful for warning provision.

The maximum 1-km wind diagnostics can provide an *approximate* guide to gust strength at the surface, though note that other factors, such as stability, can play a substantial modulating role. Such diagnostics were mainly designed to provide a general metric of cyclone intensity; in this regard it is advantageous that stability over land and the diurnal variations thereof are not having any modulating effect. Users requiring point forecasts of gust strength, in which stability and wind shear are accounted for, can utilise the standard ECMWF 'maximum 10 m gust' diagnostic.

Finally, on Figure 7, upper-level jet strength is included in the lower-right panel as an alarm bell for when the forecast surface developments have a higher potential to go awry. Stronger jets are dynamically more active, and there are well-documented cases of 'forecast busts' for cyclones associated with these. In extreme cases upper jet strength reaches ~200 knots, so with a 120-knot forecast the error potential here is less than it could be.

doi:10.21957/un8liw3vjs

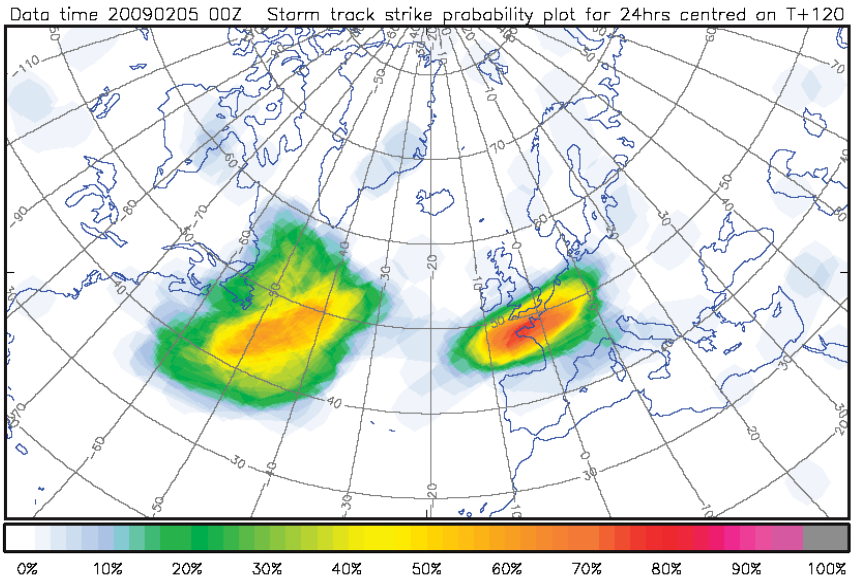

Figure 6 'Storm track strike probability plot' for ±12-hour window centred on T+120 (12 UTC on the 9th to 12 UTC on the 10th) from data time 00 UTC on 5 February 2009. Colours show the probability that a cyclonic feature meeting certain wind threshold criteria will pass within 300 km. The threshold used is that the 1-km wind, within a 300 km radius, must exceed 60 knots during the time window; this is called the 'storm' threshold.

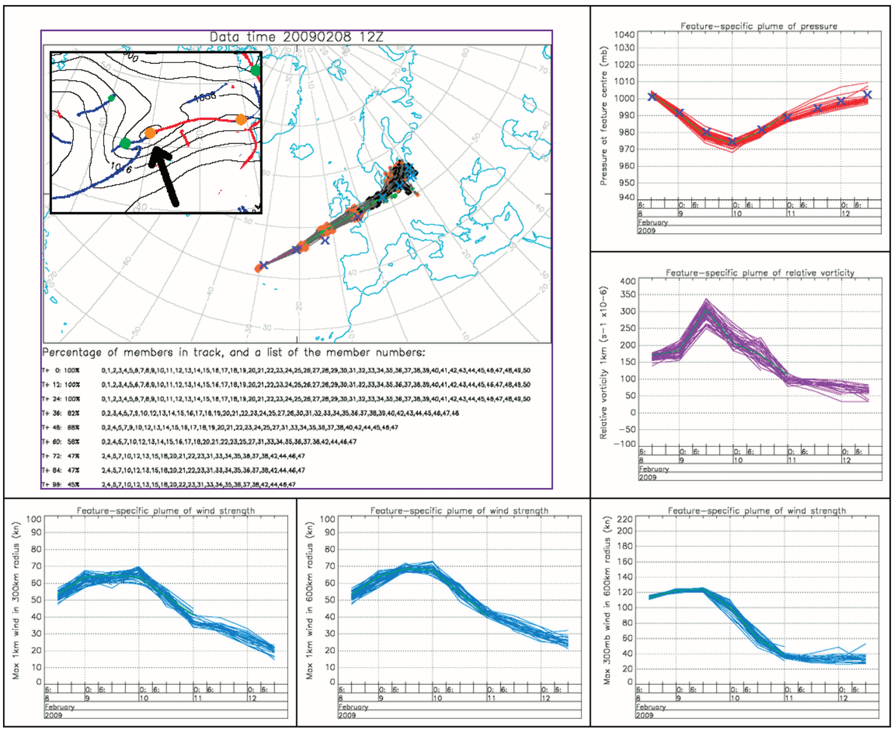

Figure 7 Top-left inset: Segment of the control run automated synoptic chart analysis (T+0 frame) at 12 UTC on 8 February 2009; colour convention is as specified on Figure 3, and mean sea level pressure is in black. The interactive web site allows the user to click on cyclonic features to see their behaviour in the ensemble. Remaining panels show feature plume diagrams for the arrowed ('clicked') feature, with verifying data added with blue crosses. Top-left panel: the feature track in the ensemble (this can be animated if the user clicks on the panel) with feature spots at 12-hour intervals. Top-right: feature point mean sea level pressure. Centre-right: relative vorticity at 1 km above the model topography. Lower-left: maximum wind at 1 km within a 300 km radius of the feature point. Lower-centre: same as lower-left but for a 600 km radius. Lower-right: maximum wind at 300 hPa in a 600 km radius

What is the probability of a major windstorm?

Finally using Figure 7 the forecaster should note the percentage of members in which the feature has been tracked (top-left panel), but at the same time be mindful that in some cases of cyclogenesis it is not always clear *which* of two successive waves might be developing. All members might be 'developmental', though there may be a split decision on which wave will develop. In the inset to Figure 7, for example, the diminutive wave southwest of the arrowed frontal wave could in some members develop instead. The closer one is to feature genesis the more of a problem this can be. So it would be incorrect to infer immediately from Figure 7 that there is only a 68% chance of a major feature existing by T+48.

To support their analysis the forecaster should refer also to strike probability plots from the same data time, which are not feature specific, and which in this case actually show a 100% probability of a major windstorm (Figure 8), with high confidence in its track (much higher, incidentally, than is confidence in the track of storm(s) well to the west).

What actually happened?

Verifying data is plotted on parts of Figure 7 as crosses, showing good agreement, except for feature position at T+24 that lay outside the ensemble range, though not by much. The verifying Met Office analysis for 00 UTC on the 10th is shown in Figure 9. The storm, named 'Quinten', is a large feature with a relatively slack core, and shows a nice resemblance to the 'T-Bone' stage of the conceptual model in Figure 3. Inland gusts of over 30 ms-1 (60 knots) were widespread between 46° and 48°N over France (tallying quite well with the 1-km wind maxima on Figure 7), whilst there was disruption due to heavy snow and flooding north of the track, over England.

In summary, as this severe event approached, the new products told us that there was an increasing risk of a major storm system in our area, they highlighted the track the system was likely to take, and they also indicated increasing confidence in that track.

Verification

By using, as truth, the cyclonic feature tracks identified in a sequence of analysis frames from either the EPS control run or the high-resolution deterministic run it is possible to verify many of the products illustrated here in an automated fashion. For features tracked from T+0, which have feature plumes, many characteristics can be verified, such as displacement errors, system velocity errors, deepening/filling errors and intensity errors (as denoted by wind strength in the circulation). As well as providing a picture of year-on-year performance changes, such statistics can also guide forecasters regarding how much confidence to attach (on average) to a forecast of a given type of feature, in a given area, at a given lead time. In turn, by comparing the spread in handling of a feature in the current forecast with such statistics one can gain insight into the predictability, in relative terms, of the current situation.

More generally, the strike-probability charts lend themselves to verification using a Brier Score approach. At the Met Office, for example, verification of these strike probabilities for the MOGREPS system, over a two-year period, suggested that there was a small degree of skill in predicting the more extreme storms beyond day 10. For those involved in warning provision for windstorms this provides a clear message – that for certain customers who are sensitive to small changes in probability there would be justification for issuing probabilistic severe event warnings at very long leads.

To re-iterate, such verification activities intrinsically focus on aspects of the forecast that are directly connected to adverse surface weather, and so have clear advantages, from a user perspective, over some of the more traditional measures (such as root mean square error in the 500 hPa height). The 'null cases', anticyclones in this case, are implicitly left out. Moreover, by focussing on features which are correlated with adverse weather, we are effectively using a proxy for severe weather verification that, conveniently, does not have associated with it the problems of observation representivity, observation quality control and variable reporting practices that arise when one directly verifies the weather parameters themselves.

 doi:10.21957/un8liw3vjs

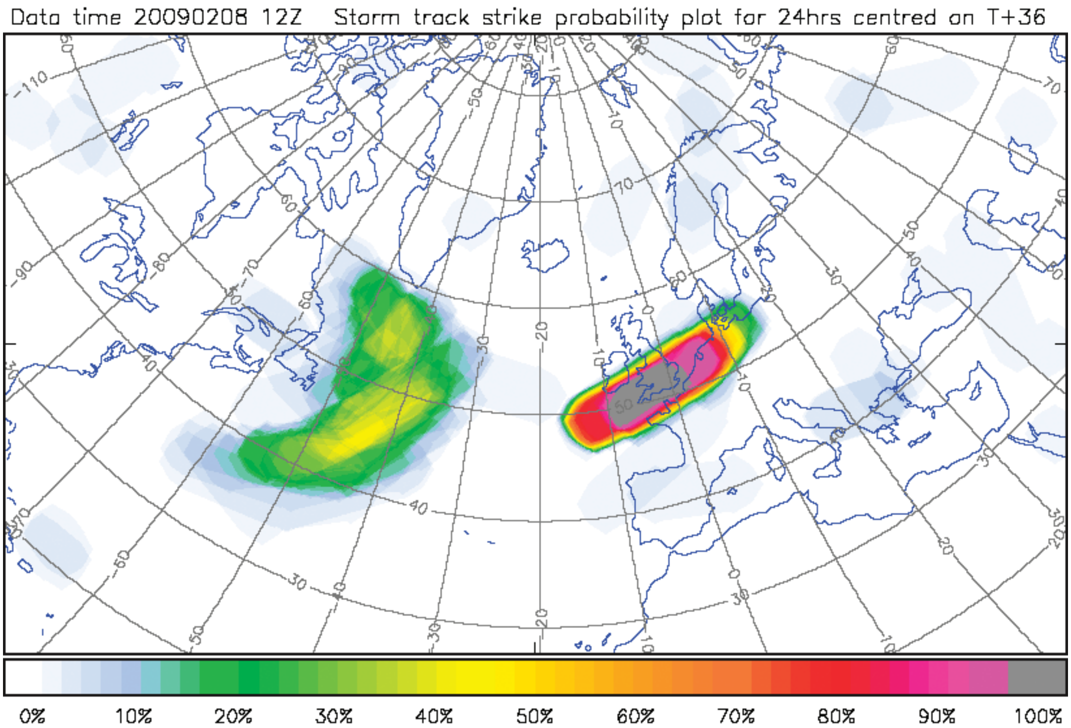

Figure 8 'Storm track strike probability plot' for ±12-hour window centred on T+36 (12 UTC on the 9th to 12 UTC on the 10th) from data time 12 UTC on 8 February 2009. The meaning of the colours is as specified on Figure 6.

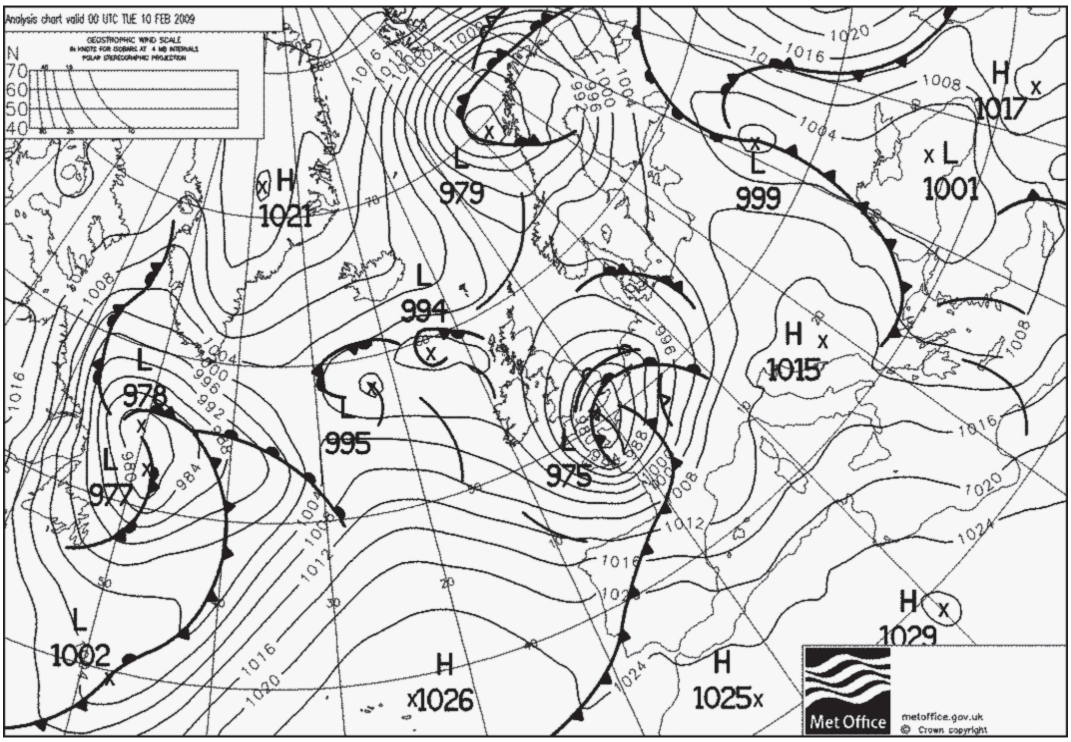

Figure 9 Met Office analysis chart for 00 UTC on 10 February 2009 showing storm 'Quinten' (975 hPa).

New opportunities and further developments

Through a new feature-based approach to post-processing ECMWF is now developing a suite of new products that provide fresh insights into ensemble handling.

A key attraction, for the user, of this new strategy is that the products use the synoptic 'language of forecasters', by focussing on fronts, cyclonic features and cyclonic feature tracks. The inherent automation should vastly reduce the amount of time the user needs to spend analysing the ensemble and deterministic output. Because 'features' have historically been used, in part, to highlight the potential for severe weather to occur, these new products inherently focus, by proxy, on this potential. This is despite the fact that the finite resolution of the model precludes a direct representation of many of the adverse weather phenomena themselves (e.g. line convection on a cold front). Verification of severe weather 'by proxy' is another significant opportunity that arises from this work.

From a product perspective further developments are planned. Firstly we will incorporate deterministic model output where appropriate. Then on the plume diagram web page other 'attributes' such as precipitation maxima and 10-metre gust maxima are likely to be added, with direct access to 'representative member' animations also provided. The dalmatian plot range should be expanded to signify other aspects, such as feature point mean sea level pressure, using a colour scale.

In concluding we acknowledge the significant contribution of Helen Titley at the Met Office, who first developed many of the track-related products described here.

Operational implementation of the products discussed in this article is expected in the first half of 2010.

Further reading

Hewson, T.D., 1997: Objective identification of frontal wave cyclones. *Meteorol. Appl.,* **4,** 311–315.

Hewson, T.D., 1998: Objective fronts. *Meteorol. Appl.,* **5,** 37–65.

Hewson, T.D., 2009: Diminutive frontal waves – A link between fronts and cyclones. *J. Atmos. Sci.,* **66,** 116–132.

Hewson, T.D. & **H.A. Titley,** 2009: Objective identification, typing and tracking of the complete life-cycles of cyclonic features at high spatial resolution. Submitted to *Meteorol. Appl.,* May 2009.

 doi:10.21957/un8liw3vjs