PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Weather Forecasting by Interactive Analysis of Radar and Satellite Imagery [and Discussion]

B. J. Conway, K. A. Browning, A. M. E. Runacres and P. L. C. Jeynes

Phil. Trans. R. Soc. Lond. A 1988 **324**, 299-315
doi: 10.1098/rsta.1988.0021

Phil. Trans. R. Soc. Lond. A **324**, 299–315 (1988)
Printed in Great Britain

Weather forecasting by interactive analysis of radar and satellite imagery

By B. J. Conway and K. A. Browning, F.R.S.
Meteorological Office, London Road, Bracknell, Berkshire RG12 2SZ, U.K.

[Plates 1 and 2]

At the Meteorological Office, Bracknell, quantitative rainfall maps from a network of ground-based radars, augmented by cloud images from *Meteosat*, are used to produce analyses and very-short-period forecasts of precipitation. These remotely sensed images provide the only way of presenting the current weather situation quickly enough and with the required spatial resolution and areal coverage. The processing of the radar and satellite data is highly automated, but there are some tasks that require judgements based upon many strands of information and an understanding of meteorological processes. To this end, forecasters use a specially developed display system to interact with the imagery. The facilities for interacting with the pictures have been optimized so that the forecaster, who is kept very busy in active weather situations, can keep pace with the flow of real-time data. Even so, as more radars are added to the network, ways must be found of reducing the burden of the forecaster's interpretive and judgemental functions by automating some of them and making others easier to perform.

1. Introduction

Remotely sensed images from a ground-based radar network and from the European geostationary satellite, *Meteosat*, are being used in combination as a routine part of the U.K. Meteorological Office's precipitation forecasting operation (Browning & Collier 1982). Every half hour, quality-controlled current rainfall maps and a series of extrapolation forecasts for a few hours ahead are generated for a limited area embracing England and Wales. Such information is required for flood prediction and as an input to detailed forecasts that are of value to a variety of users including farmers, the construction and leisure industries, and aviation. Some of these products are already being disseminated; others are still under development.

No other type of observing system can provide data with the spatial and temporal resolution required, but the image data, although our primary source, are not sufficient by themselves. Pure image-processing, however sophisticated, cannot compensate for all the ambiguities and errors that occur in these data, which must instead be interpreted in the light of knowledge from a variety of sources, including conventional observations, numerical model predictions and conceptual models of mesoscale atmospheric processes (Browning 1986). Weather forecasters are used to combining different data sources to construct forecasts, but they need new methods if they are to be able to work with images as easily as they have traditionally worked with isolated observations and line-drawn charts.

Appropriate presentation of the images is essential if the human operator is to stand any chance of recognizing the important relations contained within the pictures. It is particularly important to be able to view images in relation to others and as members of sequences; organization concealed in a single picture becomes instantly apparent when images are viewed as a rapid temporal sequence. All the images, both radar and satellite, must be presented on a common projection, and means of alternating between images and of combining one picture with another must be provided. Operations on single images, to delete, add or highlight selected features, must be available so that the forecaster can not only view the images but modify them. It must be remembered that images are not just input data used by the forecaster; they are also the natural means of presenting the final product to the customer.

The value added to the raw-image data by careful interpretation would be quickly lost if this resulted in a large delay in distributing the products to users; the total forecast period is only a few hours and the value of current rainfall maps decays quickly. None of the tools in a traditional forecast office equips the forecaster to handle images in the ways required, much less to do it quickly enough to be of any practical use.

The Meteorological Office's answer to these demanding image-handling requirements has been the development of a special-purpose interactive display system called FRONTIERS, which allows an experienced forecaster to examine, modify and combine the images, finally transmitting them to users throughout the country. FRONTIERS provides a range of facilities to help the forecaster and relieve him of routine, low-level tasks, while still leaving him in control, with the flexibility to modify the images as he wishes and to decide which functions to perform. With such a system it is possible to analyse and correct the radar rainfall images within minutes of receiving them and very shortly afterwards to have generated and distributed a sequence of extrapolation forecasts. The immediacy of this process and the dominance of the current situation in these very short term forecasts are encapsulated in the term 'nowcasting' which has been coined to describe operations such as these.

2. The FRONTIERS system

FRONTIERS was designed and built by Logica to a Meteorological Office functional specification and delivered in initial form in 1983. Experiments at the Meteorological Office's Radar Research Laboratory had already shown the kinds of task that the forecaster would need to perform when analysing the images and constructing the forecast, and work at other laboratories (Suomi *et al.* 1983) had demonstrated the use of interactive display systems for meteorological analysis, so it was possible to specify with confidence the facilities needed (Browning 1979). FRONTIERS is truly interactive; the forecaster uses it to modify images as well as to view them, and it is designed to put him in direct contact with the images without being conscious of the computer as an intermediary.

FRONTIERS comprises an interactive work station (figure 1) supported by a DEC VAX 11/750 computer. Within the work station the images are presented on high-performance colour displays, one of which is fitted with a transparent touch-sensitive panel so that the forecaster can operate on the images simply by pointing to features of interest with his finger. This colour monitor is flanked by menu screens on which are presented the options available to the forecaster at that point. Again, these are fitted with touch-screens so that selections are made by pointing to the desired options; there is never any need to remember file-names or computer commands or to type these at a keyboard. Other input devices, a joystick and a

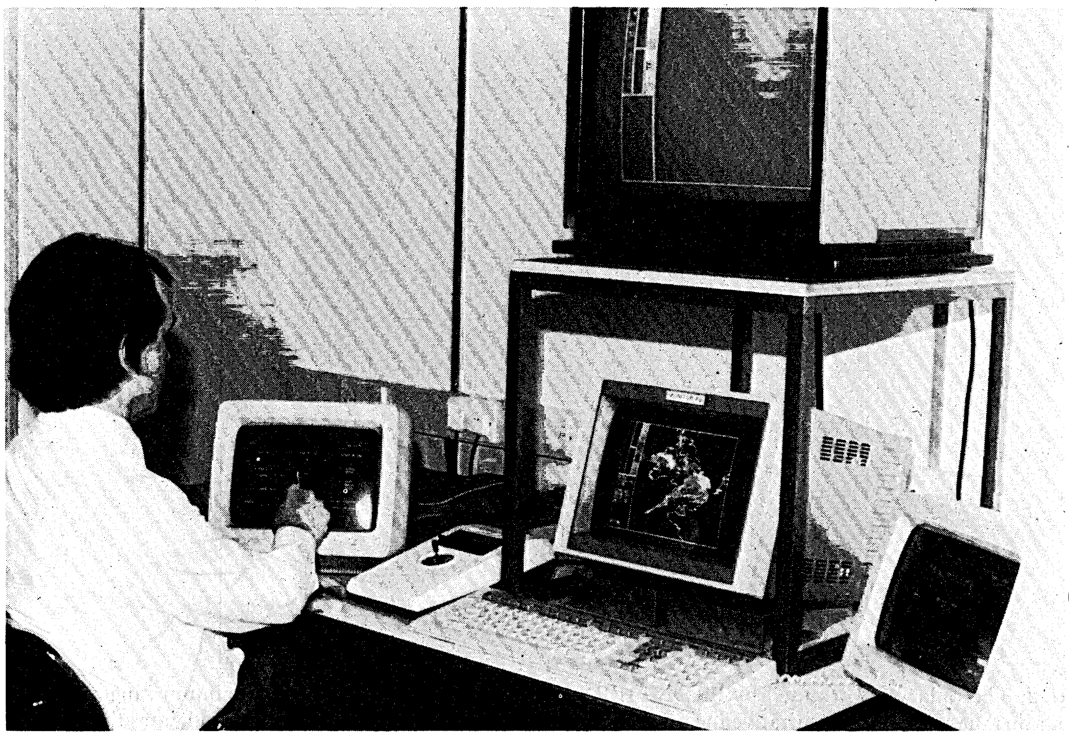

FIGURE 1. The FRONTIERS work station.

graphics tablet, are available for those functions to which they are well suited. Great attention is paid to making the forecaster at ease with the equipment: the system remembers which forecasters are left handed and swaps the menu panels, and every forecaster has his own preferred colour table in which the images are presented. The central idea is that the forecaster can give his entire attention to the job in hand, rather than to operating the equipment.

The FRONTIERS input data consist of quantitative rainfall maps from a network of ground-based radars and images from the European geostationary satellite *Meteosat*. Rainfall measurements in the U.K. network (figure 2) are made by the 'radar reflectivity' method; the rainfall rate is deduced from the strength of the return signal. Each radar site transmits its measurements as an array of 84 pixels × 84 pixels, each pixel showing the instantaneous rainfall rate averaged over a 5 km square. Readings from four or five telemetered rain gauges are sent with each radar image for calibration and as an independent check on the radar data. Measurements from all the sites are received every 15 min, but only the half-hourly pictures are used directly in FRONTIERS products; the intermediate ones are available to various users, but are used within FRONTIERS itself just to improve the continuity of replay sequences.

Digital images from *Meteosat* are received locally at half-hourly intervals. *Meteosat* scans the earth in three spectral channels: visible (0.4–1.1 μm), water vapour (5.7–7.1 μm) and infrared (10.5–12.5 μm). The horizontal resolution in the region of the southern U.K. is about 6 km × 3 km in the visible channel and 12 km × 6 km in the other two, but the images are sampled onto the same 5 km × 5 km resolution National Grid projection used for the radar images so that all the image products are readily compared. Although *Meteosat* images have poorer resolution

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

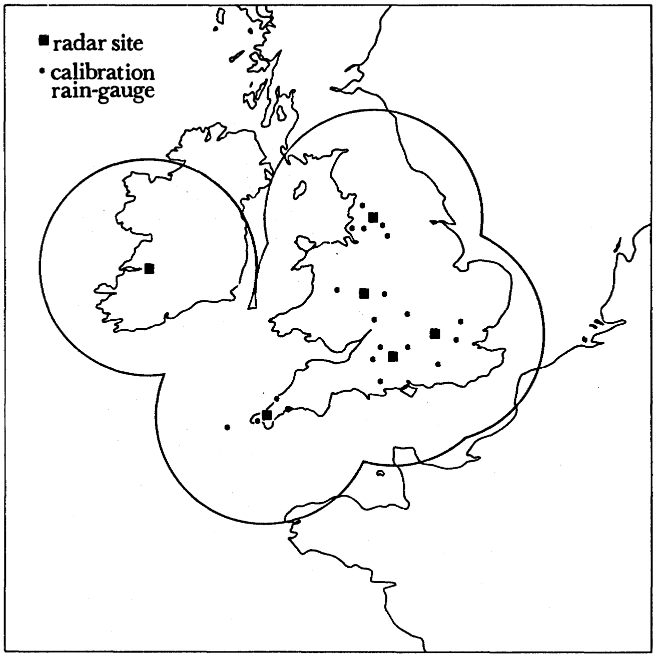

FIGURE 2. The radar network used by FRONTIERS in 1987. The circles show the maximum range for detection of rain; the range for quantitative measurements is smaller. The frame of the picture indicates the extent of the *Meteosat* image area used by FRONTIERS.

and are available in fewer spectral channels than the images from the NOAA polar-orbiting meteorological satellites, they have the overriding advantage for nowcasting of being repeated every 30 min rather than at intervals of several hours.

3. THE FRONTIERS OPERATIONAL CYCLE

FRONTIERS operates on a half-hourly cycle synchronized with the transmissions from *Meteosat*. In each cycle three major tasks are performed: a quality-controlled radar composite is produced, *Meteosat* images are used to extend the area of coverage and very-short-period rainfall forecasts are generated by extrapolation. Each of these tasks is divided into a sequence of smaller steps. Here we concentrate on a few selected operations which show the essential features of FRONTIERS and how it allows a forecaster to interpret and modify the image data in the light of his understanding of the weather situation, using a mixture of objective and subjective processes.

(a) *Radar analysis stage*

In this first stage the forecaster attempts to make the best possible interpretation of the latest radar network rainfall image, using the various data sources at his disposal to build a picture of the weather situation, of which the precipitation is one facet. Many of the errors in the radar images are not limitations intrinsic to radar, but are meteorological in origin, so that attempts to correct them are inseparable from an understanding of the weather systems being observed. Figure 3 shows some of the more important sources of error.

The first operation in the FRONTIERS cycle is to assemble a composite picture from the individual radar site data. Some corrections are applied automatically at the radar sites but

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

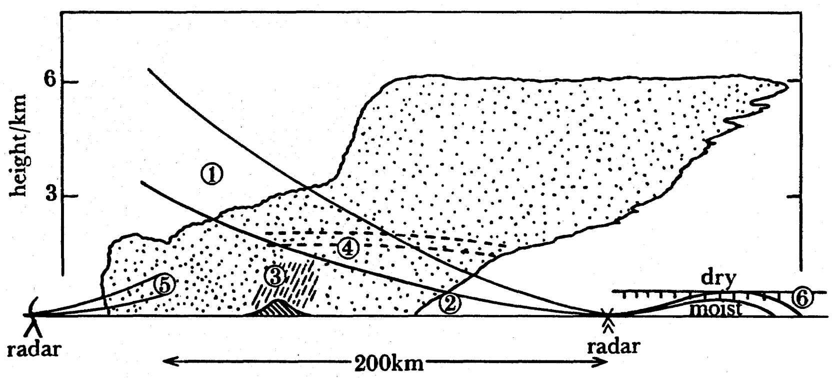

FIGURE 3. Cross section through an area of frontal precipitation showing six sources of error in the radar measurement of surface rainfall intensity: (1) radar beam overshooting shallow precipitation at long range, (2) low-level evaporation beneath the radar beam, (3) orographic enhancement above hills which goes undetected beneath the radar beam, (4) anomalously high radar signal from melting snow (the 'bright band'), (5) underestimation of the intensity of drizzle because of the absence of large droplets and (6) radar beam refracted by a strong vertical humidity gradient causing it to intercept land or sea (anomalous propagation).

most of these are removed to provide a fresh start for the more comprehensive FRONTIERS correction procedure. The operator's main task in overseeing the construction of the composite is to recognize gross errors caused by radar malfunctions and, if necessary, to remove completely the data from the offending radar. Such faults are usually instantly recognized by the operator when displayed as a rainfall map, though they are not easily identified by the radar site computer.

The next task is to try to identify and remove spurious radar echoes, that is any echo which does not correspond to rain at the surface. These include echoes produced by backscatter from the surface (ground and sea clutter) in conditions of anomalous beam propagation, and those resulting from rain which evaporates before reaching the ground. To do this job properly the operator must take into account information from a variety of sources, in addition to that in the radar images themselves.

It is difficult to recognize ground and sea clutter echoes in a single image, particularly if they are not at habitual clutter locations. Their appearance is often indistinguishable from that of genuine rainfall. The operator therefore uses FRONTIERS to replay repeatedly a series of images as a 'movie loop'. Most spurious echoes caused by ground returns then show up clearly because their stillness or erratic movement contrasts strongly with the organized movement of real rain systems. *Meteosat* cloud imagery is sometimes useful in the identification of spurious echoes. If no cloud can be detected in the region of the echo then it is unlikely to be rain. Spurious echoes that happen to occur in an overlap region between two radars may also be revealed if the image from the appropriate radar is alternated rapidly with the composite image with that radar removed; a particular ground clutter echo is usually seen by only one radar so it will flash on and off as the picture alternates. There is some prospect that spurious echoes from ground returns can be trapped by specialized hardware at the radar receiver, by using the temporal characteristics of the signals, but sea clutter is less easy to distinguish by such means and is likely to remain a task for FRONTIERS.

For most practical purposes we are interested only in rain that reaches the ground, so it is important to try to detect instances where rain detected aloft is evaporating in a low-level dry layer (figure 4). In doing this the operator will use many strands of evidence from different

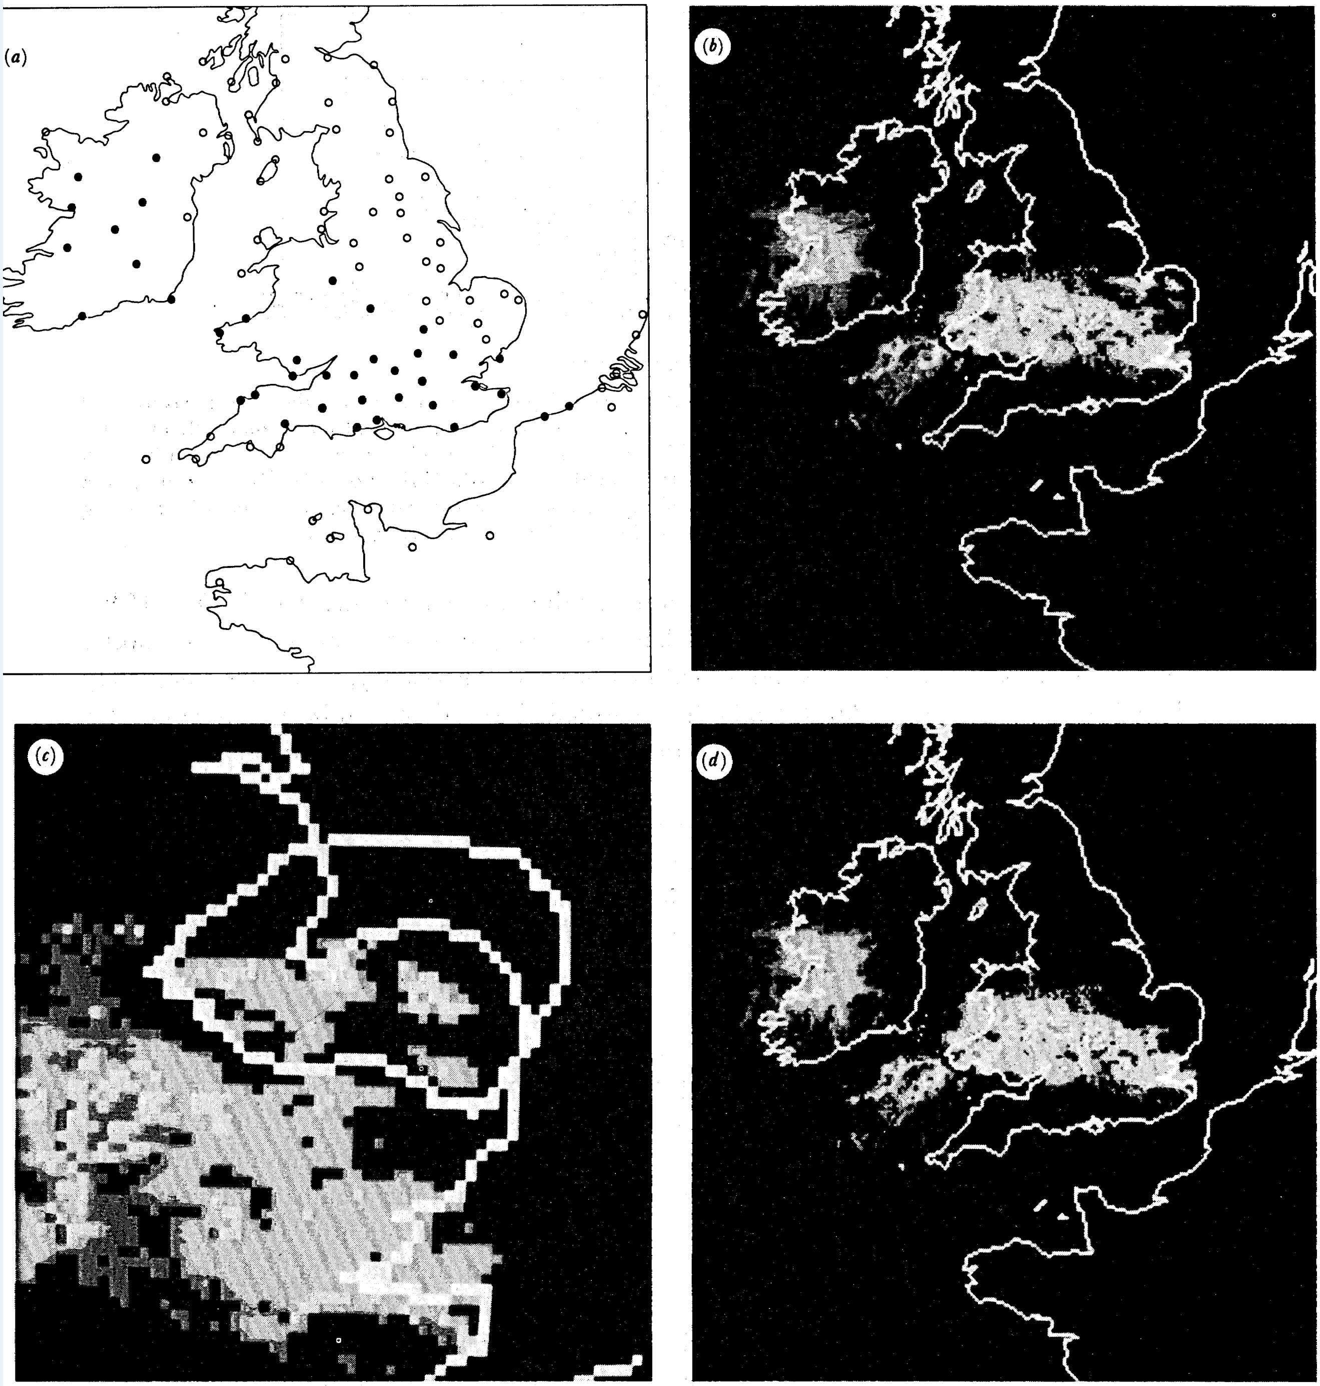

FIGURE 4. Correcting for low-level evaporation of rain in the FRONTIERS radar analysis stage. (*a*) Surface observations (from conventional meteorological stations) of rain (●) and no rain (○) at 03h00 GMT on 9 February 1987. (*b*) Radar composite for the same time as in (*a*). Rain is detected aloft, south of the Wash, by the London radar but is evaporating before reaching the ground. (*c*) The FRONTIERS operator zooms in on the evaporating rain and marks it for deletion by drawing a boundary round it on the screen. (*d*) The radar composite of (*b*) with the evaporating rain removed. The radar is missing some light rain and drizzle near the south coast because of the small drop sizes. The operator gets a chance later in the FRONTIERS cycle to insert rainfall here.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

sources. He will be helped in this by the fact that FRONTIERS is located in the Central Forecasting Office at Bracknell and is an integral part of a much larger forecasting operation. The operator will be aware of the synoptic situation (the large-scale pattern of weather) and will relate that to his knowledge of where the conditions for low-level evaporation are often to be found, for example just ahead of a warm front. Having identified the regions in which the effect is likely he will first see how far they are from the radar sites and, from this, whether the atmosphere is being sampled a long way from the ground. Unless he is convinced that rain is being measured low enough to be sure of reaching the surface he will go on to seek more specific evidence. Radiosonde ascents, which show the vertical temperature and humidity profiles of the atmosphere, may provide direct confirmation of the presence of a dry layer but will not normally allow precise definition of the area affected. Numerical forecast models will often provide a clearer indication of the pattern of humidity. The most useful information, however, is that from conventional surface observations and from telemetered raingauges. The relative importance of these different sources is not fixed, but depends on how they relate in space and time to the feature being considered; it is up to the operator to weigh them accordingly.

The time arrives when the operator must decide which pixels must be deleted as not representing surface rainfall. It is unlikely that the observational evidence alone will have been sufficient to permit an unequivocal delineation of the region of low-level evaporation and so the operator must use his judgement. By now he will have built up a three-dimensional picture of the atmosphere in the suspect region based on a mental model of the synoptic feature, fitted to the available observations in this specific case. He now uses the FRONTIERS facilities to draw boundaries on the screen within which echoes should be deleted, or alternatively to zoom in and delete individual pixels. The machine places him in direct contact with the images and makes it easy to modify them and to change his mind as often as time permits. Thus his full attention can be given to the changes he wishes to make rather than to the mechanics of the operation. Only when he is satisfied does he instruct the machine to implement the changes permanently and move on to the next step.

Four further steps remain before the FRONTIERS radar analysis stage is complete. If the presence of enhanced echo from melting snow (the so-called bright band) has been detected at any radar site the operator must examine corrections computed by FRONTIERS and decide whether to apply them. He must define areas of different types of rainfall, so that appropriate calibration factors can be applied and so that the range at which the beam will overshoot the rain and be of limited further use can be calculated. He must then look for evidence of individual radar calibration failures, where the sensitivity of a particular radar has changed significantly, and try to correct this. Lastly, he must decide whether intensification of rain over hills (so-called orographic enhancement) is likely to be important, and if so where. Orographic enhancement is a low-level phenomenon and often occurs beneath the radar beam. Thus it is necessary for the operator to instruct FRONTIERS to calculate and apply the appropriate correction factors.

In all of this the operator will be examining the radar measurements critically in the light of other information, just as he did for deletion of spurious echoes, and using FRONTIERS to implement modifications, calculate corrections and show quickly the results of his decisions. The operator has, at most, 17 min to complete all the steps listed above; if he has not finished by then FRONTIERS will take control and finish the radar analysis automatically, applying the corrections used in the previous cycle. If the operator can finish sooner he is rewarded by getting extra time for the satellite analysis step which follows, but if he takes less than 12 min he has to wait for the satellite images to arrive. This shows the extremely tight scheduling of the FRONTIERS cycle and the paramount importance in the design of the system of an efficient man–machine interface.

(*b*) *Satellite analysis stage*

In the second major part of its 30 min cycle FRONTIERS is used to estimate, from the satellite cloud imagery, where rain is likely to be falling outside the radar area, so that a less fragmentary glimpse of the overall pattern can be obtained and so that some warning can be given of approaching rain systems.

The *Meteosat* images do not show rain directly, so its presence must be inferred from cloud images (figure 5, plate 1). FRONTIERS allows this to be done by a mixture of objective and subjective techniques. Different information is available from the different spectral channels: the infrared image shows the temperature, and therefore the height, of the cloud tops, whereas the brightness of the visible image indicates the thickness of the cloud. Both are important in assessing the likelihood of rain falling from the cloud. At night, however, the infrared image must be used alone.

The preferred method of relating the cloud information to the rainfall pattern is to correlate both the visible and infrared images with the rainfall measurements within the radar area, and then to assume that this correspondence holds in other areas as well (Lovejoy & Austin 1979). The range of possible pixel values in each channel is divided into 16 equally sized classes so that the cloud at a given location will correspond to one of $16 \times 16 = 256$ possible visible–infrared classes. Each cell within the radar area is examined in turn and 'rain' or 'no rain' is entered against the cloud class at that location. At the end of this process we obtain, for each cloud class present in the data, the fraction of occurrences corresponding to rain, which can be taken as a measure of the probability of rain falling from that class of cloud.

The relation thus obtained is then applied, in reverse, to the whole area. All pixels in the cloud category with the highest rain probability are assigned as rain, then all those with the next highest probability, and so on, until the total number of pixels of inferred rain within the radar area is as close as possible to the number of radar rain pixels. This method prevents grossly inaccurate estimates of the amount of rain but, because the relation is imperfect, the pattern of the inferred rainfall will not exactly match that of the radar rain, and the degree of mismatch, both within the radar area and at its boundary, is used by the operator to gain a subjective estimate of the success of the operation and the acceptability of the correlation product.

FRONTIERS gives the operator the option of correlating rainfall against either the visible

---

Description of plate 1

Figure 5. Derivation of an extended area rainfall field from *Meteosat* cloud images. (*a*) Radar composite for 13h00 GMT on 4 August 1985, showing observed rainfall rates after correction for errors in the FRONTIERS radar analysis stage. Green shows lightest rain (less than 1 mm h^{-1}), blue shows heaviest (8–16 mm h^{-1}). (*b*) *Meteosat* visible (0.4–1.1 μm) image for the same time as in (*a*). The whitest areas correspond to the brightest and therefore the thickest cloud. (*c*) *Meteosat* infrared (10.5–12.5 μm) image for the same time as in (*a*). Colours show the temperatures and therefore the heights of the cloud tops. The coldest cloud tops in the picture are indicated by the orange patches over eastern England and Northern Ireland. Very dark areas are warm cloud or surface. (*d*) Field of rainfall at 13h00 on 4 August 1985, inferred from the *Meteosat* visible and infrared images by correlating them with the radar rainfall pattern within the radar area. This field shows only the extent, not the intensity, of the rain.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

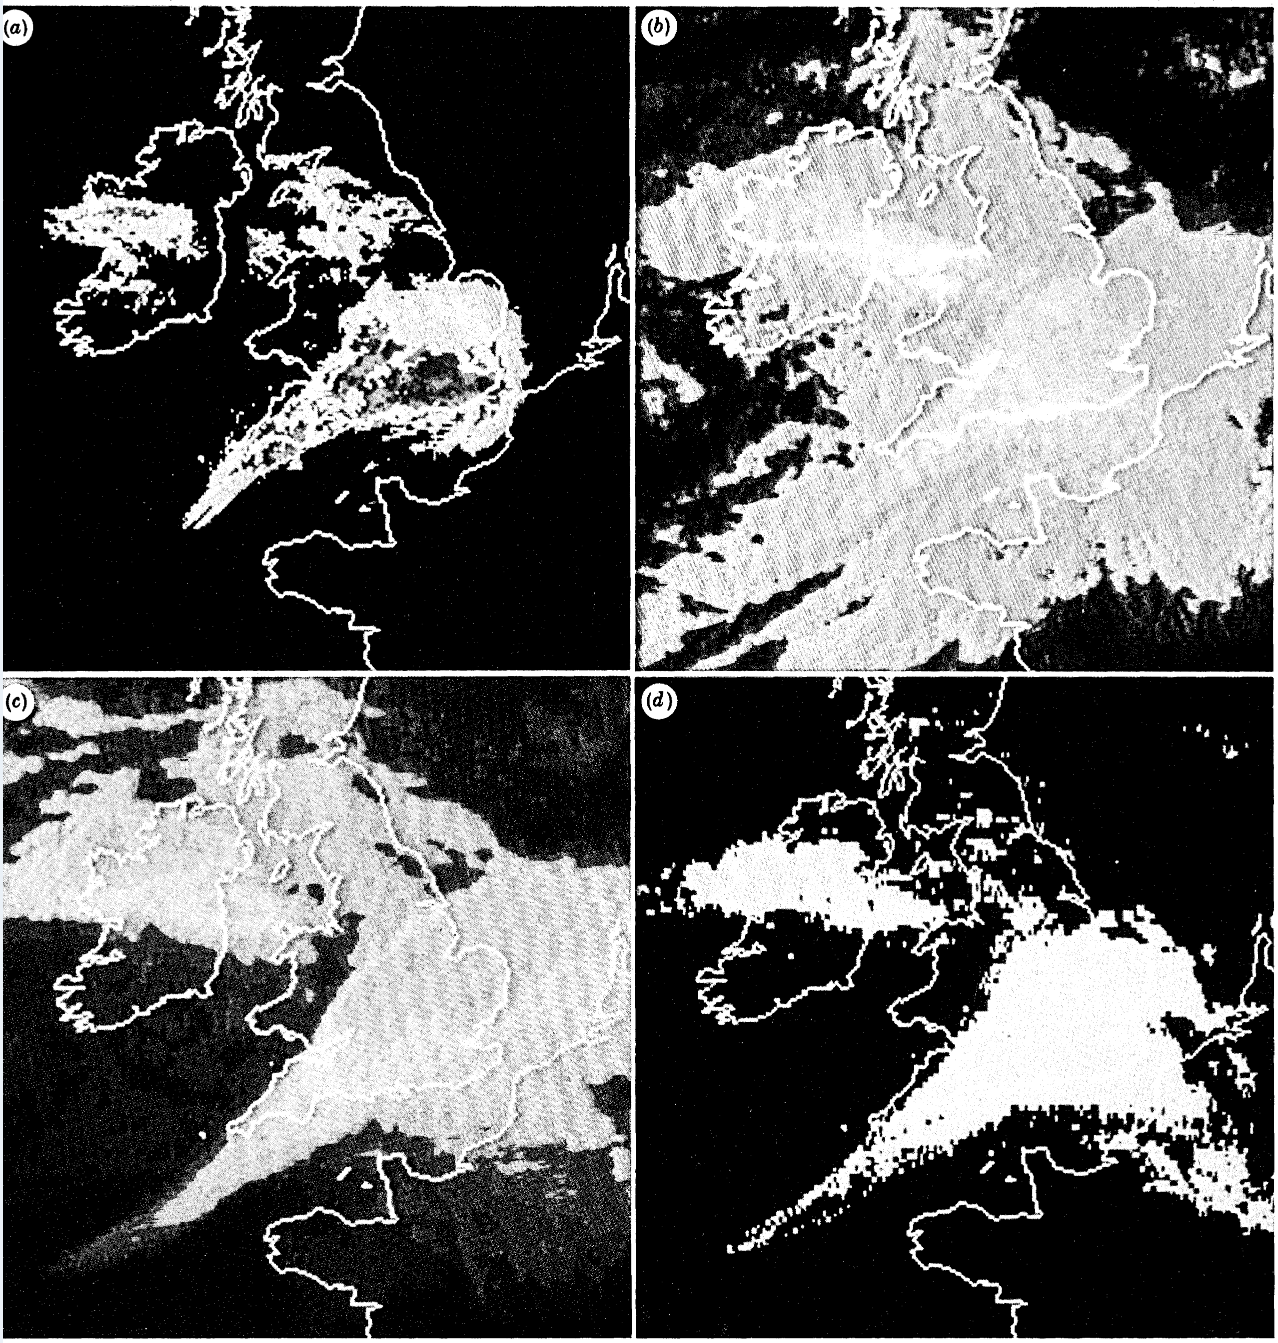

FIGURE 5. For description see opposite.

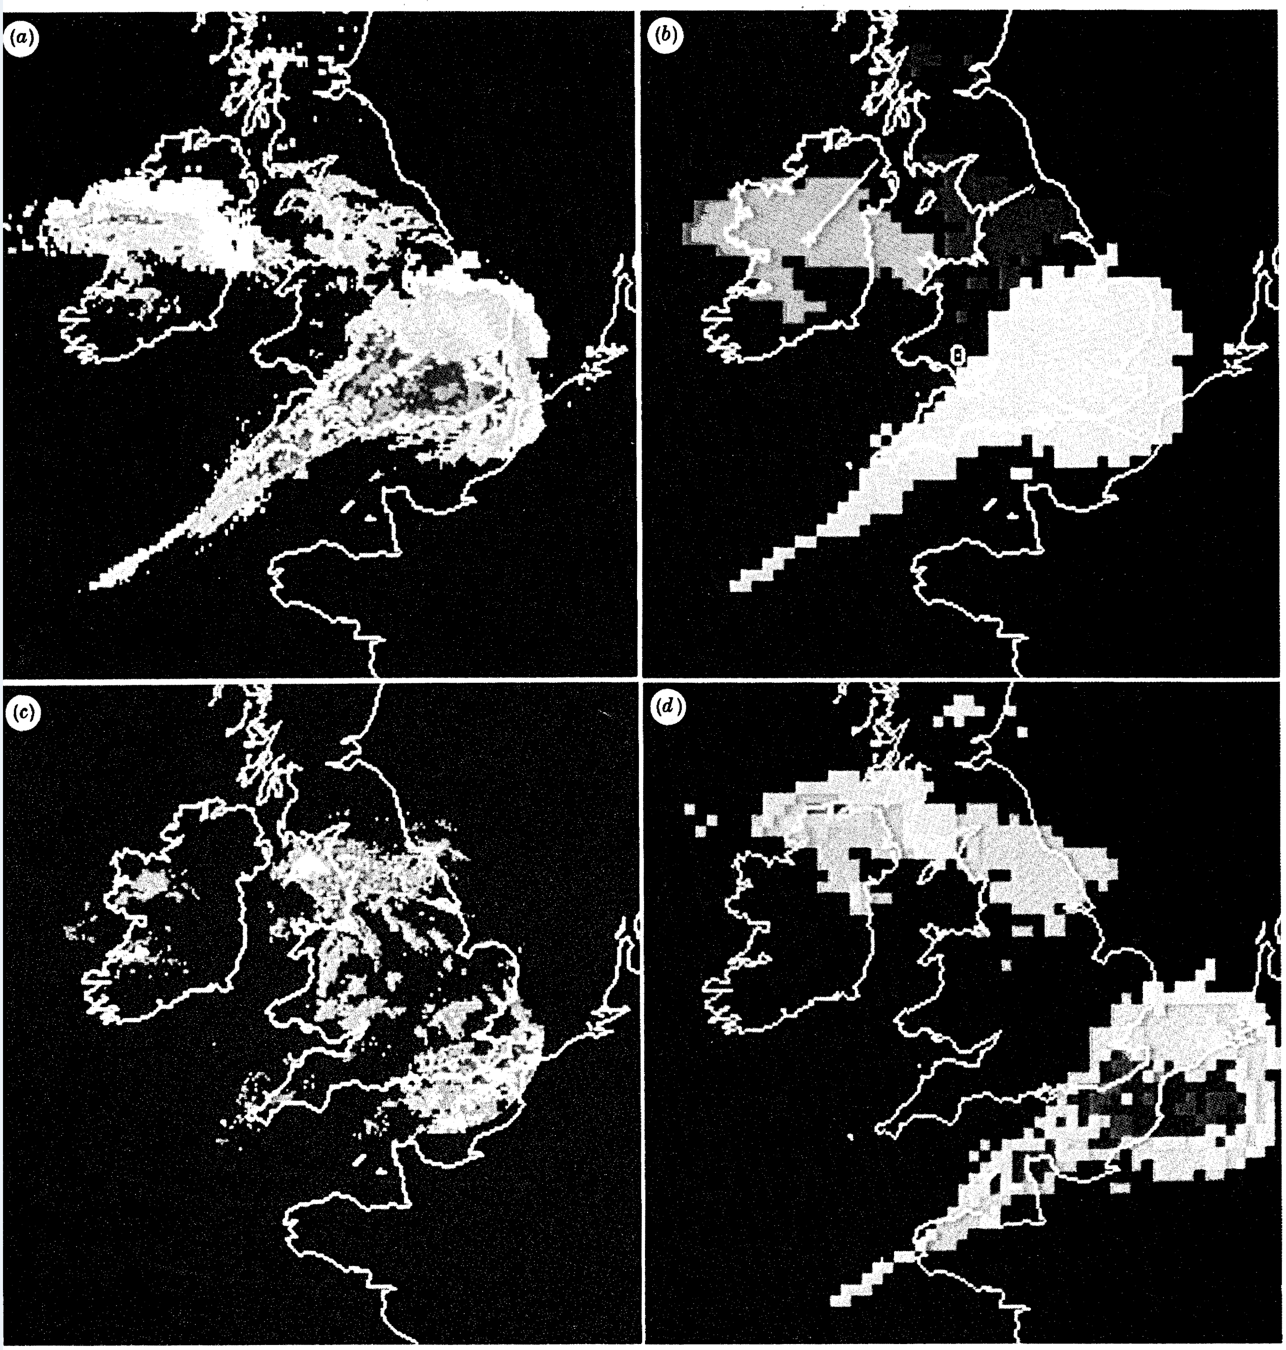

FIGURE 6. For description see opposite.

MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES
PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A

image, the infrared or, as described above, the two combined. Normally he would try the visible–infrared combination first, provided that both images were available, but if dissatisfied with the result he might try the other options. The choices available are presented on a touch-sensitive menu screen, so that exploring the different possibilities is quick and easy.

In most cases one or other of the correlation products will provide a reasonable approximation to the desired rainfall field, but the method is subject to a number of problems.

(i) If there is very little rain in the radar area, the correlation will be unreliable. This is countered by using a continuously evolving correlation table which consists partly of the present correlation and partly of the table used last time. If this fails, a fixed 'universal' table is available.

(ii) If the predominant types of rainfall inside and outside the radar area are different then the assumption that the same relation between cloud class and rain holds in both regions may be invalid.

(iii) The method is seriously weakened at night when only the infrared image is available, because the cloud-top temperature does not on its own bear a close or reliable relation to the presence or absence of rain, except for certain types of rain (convective storms).

(iv) A small misregistration of the satellite image with respect to the radar rainfall map can have a marked effect on the correlation result if the cloud and rainfall patterns contain much fine structure.

The operator must decide whether to accept one or other of the correlation products. He usually does this by comparing the satellite-derived field with the radar analysis product, using a FRONTIERS facility which allows the correlation rainfall field to be displayed with the radar field embedded. By using FRONTIERS to alternate rapidly between this and the complete correlation field, the operator can obtain a clear impression of how well the correlation product agrees with the radar measurements and whether the pattern outside the radar area is a realistic extension of what he sees inside it. In this he will also be guided by his knowledge of the synoptic situation and his expectation of the organization, within the large-scale pattern, of mesoscale features such as fronts and rainbands.

The operator may occasionally accept the correlation product unaltered, but more often he will decide to modify it either by deleting features he regards as spurious or by adding areas of rain where he believes that the correlation field contains unrealistic gaps or discontinuities. FRONTIERS provides the means to do this by direct interaction with the image, with pan and zoom facilities permitting accurate work on magnified areas. As usual with FRONTIERS, mistakes are easily cancelled *en route* to constructing a final product.

Description of plate 2

Figure 6. Generation of a rainfall forecast. (*a*) The radar–satellite combination rainfall field formed by adding the satellite-derived rainfall field of figure 5*d* to the quantitative radar rainfall field of figure 5*a*. The satellite-derived rainfall (shown in white) has been selectively trimmed by the FRONTIERS operator. (*b*) The rainfall pattern of figure 6*a* divided into clusters that will move independently in the forecast. Their velocities are indicated by the white and yellow vectors at their centroids. (*c*) The rainfall pattern 3 h later, at 16h00 as observed within the area of coverage of the radar network. (*d*) The 3 h extrapolation forecast for 16h00. The heavy rainfall in SE England is well forecast, as is the extent of the rain in northern England and the Scottish border region. The very heavy rain over the Irish Sea would have been missed almost entirely if the satellite data had not been available to bridge a gap in the radar coverage (figures 5*a* and 6*a*). The light rain over the Midlands and Wales developed *in situ* during the forecast period and so could not be forecast by extrapolation techniques. On this occasion the forecast actually gives better agreement with the surface observations in Northern Ireland than does figure 6*c* because the radar at Shannon (SW Ireland) is missing some rain in that area at 16h00 GMT.

Sometimes the operator will decide that none of the available correlation fields has represented the rainfall field well enough to be of any use, even as a starting point. Another option available to him is to choose a satellite image and to 'slice' it so that only a chosen range of pixel values is made visible on the screen, pixels of any other value being suppressed. By varying the size and position of this chosen range, the operator can selectively display features within the satellite picture, picking out clouds of a particular range of brightness or temperature. The operator controls the slicing by means of a joystick while watching the image evolve on the screen, so that from a very large number of initial possibilities he can quickly converge on the most realistic product. Experience has shown that a satisfactory product is more likely to be obtained by slicing the visible image rather than the infrared, so it is less easy to obtain good results by this method at night. As with the correlation field, the operator will often choose to make alterations to the sliced image, adding or deleting rain, until he is happy with the product.

As a last resort, if no satisfactory pattern can be obtained either from the correlation process or from slicing an image, the operator has the means to draw in the rainfall pattern from scratch, freehand, to produce what he considers to be a realistic pattern, agreeing with the radar data and with the current weather situation as far as he is able to determine it from the other sources at his disposal. This method produces a coarser rainfall field than the other techniques and is used on its own only when they fail. Nevertheless, it does provide a way of giving users an impression of the likely overall pattern of precipitation and a means of avoiding unnatural discontinuities associated with data gaps being advected into the radar area.

This completes the second stage of the FRONTIERS cycle, the extension of the area of coverage. The system provides the means to process this stage by a range of techniques, from the purely objective to the purely subjective, depending on the details of the particular case. The product forms the starting point for the next stage.

(*c*) *Forecast stage*

The last stage in the FRONTIERS cycle is the generation, by extrapolation of recent movement, of a forecast sequence showing the expected distribution of rainfall for a few hours ahead. The period of validity of such forecasts varies with the type of weather situation, but for the sake of delivering a uniform product, likely in most cases to show some skill throughout the forecast period, it has been decided for the present to generate a series of half-hourly forecasts stretching a fixed 3 h period ahead.

The extended area composite from the satellite analysis stage (figure 6*a*, plate 2) is first reduced to 20 km horizontal resolution by averaging neighbouring pixels. This was originally done to save computing time later in the forecast, but it has undesirable effects such as causing jerky motion of the rainfall patterns in the forecast sequence and reducing the contrast in the rainfall intensity field. We therefore plan to revert to using 5 km resolution throughout, and solve the problem of computing power in other ways (see §4).

Rarely will a single velocity vector adequately describe the motion of the whole rainfall field. The way this is taken into account in FRONTIERS is for the operator to divide the rainfall pattern into a number of clusters which he believes will maintain their identity during the forecast period (figure 6*b*, plate 2). He does this by simply drawing the boundaries between clusters on the screen using his finger. In choosing how to dissect the image the operator uses not only his knowledge of the pattern of weather producing the precipitation, but also the evolution of the pattern in the recent past, as revealed by replay sequences of the radar and satellite images.

Having selected the rainfall clusters that will populate his forecast, the operator must assign them appropriate velocities. Although a general impression of movement within an evolving weather system is easily obtained, measurement of the precise velocities of chosen features proves to be much more difficult. So far, our attempts to develop totally objective procedures have been less successful than our efforts to optimize the implementation of subjective methods.

One subjective technique is to calculate the velocities from the difference in position of the features at two known instants. Problems arise because if a short time interval is chosen, measurement of the small movement that has occurred is inaccurate, whereas if a longer interval is used to give a larger movement, changes in the shapes of the features will often make identification of corresponding points uncertain. Although FRONTIERS allows the operator to define velocities in the above manner, a much preferred technique which has been specially implemented in FRONTIERS for this purpose is the 'lagrangian replay' facility. This exploits the human operator's ability to separate a chosen feature from a confusing background and then follow it as both it and the background move and distort. This is possible provided he can see the sequence as a continuously moving picture, not as a series of disconnected 'stills'.

The operator first selects the type of image to use (radar-rainfall or satellite) and within it a feature whose velocity he wishes to measure. He chooses a short replay sequence, terminating on the latest image, throughout which the feature can be identified. He zooms in, centring on the feature, and touches corresponding points on the first and last frames so that the system can compute a 'first-guess' velocity. The sequence of images is then repeatedly replayed, with the system subtracting the first-guess velocity from the entire field so that the chosen feature appears roughly stationary. The operator then adjusts the compensating velocity by means of the joystick, while the replay sequence is running, until the feature's residual velocity is minimized, and in this way a reliable estimate of the required velocity is obtained. It is found that the lagrangian replay is the method that gives the most consistent results when the same case is processed by several different operators. Even when a feature significantly changes during a replay sequence, thereby rendering automatic procedures unreliable, there tends to be good agreement among different observers about when it appears stationary. The lagrangian replay thus proves to be a powerful subjective technique for measuring velocities in complex evolving systems.

When a number of velocities, characteristic of the identified separate components of the rainfall system, have been measured, they are assigned to the appropriate clusters simply by the operator touching the appropriate velocity on a displayed list and then touching the corresponding cluster. A particular velocity may be assigned to more than one cluster. After that the generation of the forecast sequence is largely automatic, although the operator has some choice about how orographic corrections should be applied in the forecast fields and whether he wishes to define multiplicative factors (in the light of other evidence) by means of which rainfall in the various clusters can be made to increase or decay with time.

Many aspects of the extrapolation forecast are undeniably crude. Clusters move in straight lines at constant speed without changing shape, and only the simplest representation of development and decay is included. Because the technique only depicts a continuation of observed recent trends, the development of completely new rainfall areas can never appear in the forecast. Nevertheless, in many situations, particularly in the frontal patterns which dominate rainfall events in the U.K., the method gives a good indication of the immediate movement of the rainfall pattern and constitutes useful precipitation forecast guidance in the very short term (figure 6*c*, *d*, plate 2).

4. Future developments

FRONTIERS is now undergoing operational trials in the Central Forecasting Office at Bracknell. The quality-controlled radar composite map is being routinely disseminated to users around the country but other FRONTIERS products await the availability of improved display and transmission facilities which the Meteorological Office plans to install over the next few years (Cluley & Hills 1988). In the meantime, FRONTIERS extended rainfall maps and forecasts are being used for guidance centrally.

Development of FRONTIERS continues. Even if the system were perfect in itself, the fact that it is part of a changing forecasting environment would mean that it must evolve to take advantage of new sources of data and to deliver products meeting practical needs. Another spur to development is the need to increase the throughput of the system. The operator is already kept very busy in active weather situations. With the imminent expansion of the U.K. radar network and the possibility of the use of data from continental radars, he will be extremely hard pressed to keep up with the flow of real-time data.

The further development of FRONTIERS will be approached in three principal ways: (*a*) through use of improved but standard computing techniques, (*b*) through interaction with a dynamical weather-prediction model and (*c*) through the application of expert-system techniques.

(*a*) *Use of improved but standard computing techniques*

Improvements will be made to the speed of the FRONTIERS computer system and to the ergonomics of some of the facilities. The machine is annoyingly slow when responding to some commands and, although it may be doing a lot of work at the time, this is not usually apparent to the operator, who is conscious only of having to wait before proceeding. These difficulties can be largely overcome by straightforward computing techniques (improved software efficiency, more powerful processors, faster input–output paths, etc.) once the specific causes have been identified. Similarly, observation of how the operators use the machine can show how particular facilities can be modified or redesigned to make them more convenient to use. Neither of these approaches requires any alteration to the FRONTIERS concept or any change to the division of work between the operator and the machine, but the margin for improvement by these means alone is unlikely to be sufficient for the foreseen increase in work load.

Ultimately, the point must come, though we have not reached it yet, when further enhancements to the computer system will yield negligible overall improvement because the remaining significant delays are due to the thinking time needed by the operator. We must therefore seek other ways of upgrading FRONTIERS so as to give the operator more help, and options (*b*) and (*c*) below offer two ways of doing this.

(b) *Interaction with a dynamical weather-forecasting model*

The second avenue of development for FRONTIERS takes account of the position of the system as a cooperating part of a much larger forecasting operation. Of particular significance is the current development by the Meteorological Office of a mesoscale numerical forecasting model that will cover the entire U.K. and use a horizontal grid spacing of 15 km (Golding 1987). The model, which solves the equations of motion, state and continuity for the atmosphere, will be run every 6 h and produce forecasts of many parameters, including precipitation, for periods of up to 18 h ahead. Although they cover similar geographical areas and both make short-period predictions, FRONTIERS and the mesoscale model have different strengths and so complement one another rather than compete. The first few hours of any detailed forecast are necessarily dominated by the current situation and recent trends, whereas for longer term forecasts these will be less important than nonlinear development. The mesoscale model predicts these longer term developments, which cannot appear in the FRONTIERS extrapolation forecasts, so that beyond a few hours ahead the numerical model will produce better guidance. However, the model cannot properly represent features with scales of less than three or four grid lengths, so the current situation and its immediate development are shown better by the FRONTIERS products. The point at which the cross-over in skill occurs will vary from case to case, but is likely to be in the range 2–6 h; it will be up to the forecaster who receives both products to judge which advice to follow when they differ.

To allow FRONTIERS and the mesoscale model to operate in isolation would be to throw away important synergetic opportunities for improving both. It is already planned to use FRONTIERS quality-controlled rainfall composites as one source of data for initializing the mesoscale model. Rainfall observations cannot be assimilated directly by the model but instead are used to adjust fields of relative humidity and vertical motion. Experiments will also be done to see whether output from the mesoscale model can be used to provide guidance for FRONTIERS. Among the possibilities are to obtain predictions of development and decay from the model and apply these to the FRONTIERS forecasts, instead of asking the forecaster to define constant development–decay factors as at present, and to use model wind fields to advect the analysed rainfall field rather than extrapolating rainfall clusters with constant velocity as at present. If successful, these methods would allow the automation of a substantial manual process, although the capacity for manual intervention would need to be retained as a safety net.

(c) *Application of expert-system techniques*

The remaining line of attack is even more challenging. This will be to try to copy, within the computer, some of the judgemental processes which have so far been left entirely to the operator. This is different from the example given in §4*b*, where we propose trying to replace an existing manual method with a different one more suited to automatic computation. Both approaches are likely to find application within FRONTIERS, but this third approach represents a fundamental change in the relation between the operator and the FRONTIERS computer.

FRONTIERS, as implemented at present, is not an 'intelligent' system. The job of the computer is to present control options and image data to the operator and to perform whatever processing functions are necessary to implement his decisions. It is the operator's responsibility to provide all the meteorological knowledge and understanding needed to interpret the data

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

and to make all the decisions, both about modifying individual images and about selecting which functions to perform. The man–machine interface thus acts as a boundary, not only between the man and the machine, but between two quite separate classes of function.

We now propose to try to move some of these functions across the boundary, into the machine. In this context it should be remembered that all the case-specific data used by the operator to guide his decisions reach him via computers, and are therefore, in principle at least, available to FRONTIERS. What are not yet available to the machine, and what we have not begun to formulate in the detail required, are the knowledge and experience which enable the operator to organize and select these disparate, incomplete and sometimes conflicting pieces of observational evidence, and fit them together to form a coherent picture.

The fundamental limit to this transfer of expertise will be set by our ability to understand and formulate the operator's reasoning paths. Until now, the use of the man–machine mix, in which are combined the versatility of the human forecaster and the data-handling power of the computer, has absolved us from the need to formulate the tasks in meticulous detail. Guidance on the use of the FRONTIERS system has been provided, but we have relied on the operator, an experienced forecaster, to be able to take sensible action, even in unforeseen situations.

The availability of a group of operators experienced in using FRONTIERS in its present form, for real analysis and forecasting, will be invaluable in testing and refining our understanding. In addition, FRONTIERS archives all its cases on magnetic tape so that they are available for further analysis; it will thus be possible to test attempts at automation against a large number of real cases and, particularly when they fail, to examine how one or more human operators tackle the same problem.

A valid criticism of interactive systems like FRONTIERS is that the quality of the products depends upon the skill of the individual operator and that operators are not uniformly experienced and skilled. Automation will improve the objectivity of the FRONTIERS processes and make the results less dependent upon the operator, even if he retains a supervisory role with the option of accepting or altering the automatically generated products. It will also allow more use to be made of forecasting rules which depend directly on calculations based on the physics of atmospheric processes. However, it will be important to measure the success of automated functions against the results achieved by human operators; automation is likely to be easier and more successful in some cases than in others and a pragmatic approach to its introduction is needed.

Success in automating FRONTIERS, using either method (*b*) or method (*c*), will shift some of the burden from the operator to the computer, and we must expect to encounter performance problems in the machine as a result. However, these are problems amenable to solution by method (*a*), enhancements to the computing system, rather than being absolute obstacles set by the speed at which a human being can work.

5. Conclusion

The availability of image data from radars and satellites has revolutionized our ability to measure the detailed distribution of precipitation and has led to the emergence of a new class of very short period extrapolation forecasts or 'nowcasts'. New methods have had to be developed to process these extremely perishable data and deliver them to users quickly enough to be of practical value. In the Meteorological Office this has been achieved through the development of a special purpose interactive display system, 'FRONTIERS', which allows us to merge the knowledge and experience of a professional forecaster with the data-handling and presentation capabilities of a computer. The forecaster modifies images and generates forecasts, using information from many sources about the current weather situation to help compensate for the unusual error characteristics of the radar and satellite imagery.

The forecaster is already kept very busy in active weather situations and, with the doubling in size of the radar network during the next few years, he will no longer be able to keep pace with the flow of real-time data. There is some scope for improving throughput by simply increasing the speed of the computer system, but this will be insufficient by itself. Means must be sought of transferring to the computer some judgemental and interpretive tasks, hitherto entirely the province of the human operator. This will be a difficult and challenging task, which will hinge on our ability to make a much more detailed and precise analysis of the forecaster's task than has so far been attempted. In it we shall be greatly helped by the experience gained, using the present FRONTIERS system, of how a forecaster tackles his job, and the ability to test ideas against the large collection of cases archived by FRONTIERS. Success in this enterprise will not be easy to achieve but is likely to have considerable impact, both on FRONTIERS and on many other aspects of the weather-forecasting process.

References

Browning, K. A. 1979 The FRONTIERS plan: a strategy for using radar and satellite imagery for very short-range precipitation forecasting. *Met. Mag., Lond.* **108**, 161–183.

Browning, K. A. & Collier, C. G. 1982 An integrated radar-satellite nowcasting system in the UK. In *Nowcasting* (ed. K. A. Browning), pp. 47–62. London: Academic Press.

Browning, K. A. 1986 Conceptual models of precipitation systems. *Weather forecast.* **1**, 23–41.

Cluley, A. P. & Hills, T. S. 1988 Meteorological Office outstation display system: from concept to reality. *Met. Mag., Lond.* **117**, 1–12.

Golding, B. W. 1987 Short range forecasting over the UK using a mesoscale forecasting system. In *Short and medium range weather prediction* (ed. T. Matsuno). Universal Academy Press Inc.

Lovejoy, S. & Austin, G. L. 1979 The delineation of rain areas from visible and IR satellite data for GATE and mid-latitudes. *Atmos. Ocean* **17**, 77–92.

Suomi, V. E., Fox, R., Limaye, S. S. & Smith, W. L. 1983 McIDAS III: a modern interactive data access and analysis system. *J. Clim. appl. Met.* **22**, 766–778.

Discussion

[*Questioner not identified.*] What is the economic justification for a system such as FRONTIERS?

B. J. Conway. The justification for FRONTIERS is in the value it adds to the weather radar network. The raw radar coverage is incomplete and the data subject to errors; the map of the network that I presented showed the maximum range for *detection* of rain, but the range for reliable quantitative measurements is much smaller and so there are gaps in such coverage. More radars will be added in the next few years but these will mainly serve to extend the outer boundary of the network so as to encompass the rest of the U.K. The cost of FRONTIERS is small compared with that of adding enough radars to fill all the gaps in quantitative coverage, yet it not only helps us to compensate for the deficiencies in the raw data, but it also adds an entirely new dimension to the products by allowing us to produce short-period extrapolation forecasts. The radar network is an expensive investment and we would be neglecting our duty if we failed to extract the maximum value from it; FRONTIERS is our way of getting the best from this facility.

A. M. E. RUNACRES (*Department of Geography*, *University College*, *Swansea*, *U.K.*). Is it possible for the FRONTIERS operator to alter the depiction of rainfall intensity, as opposed to adding or deleting pixels completely?

B. J. CONWAY. The forecaster can, within limits, subjectively adjust the overall calibration of any radar; he also has some freedom about how and where he instructs FRONTIERS to calculate and apply quantitative corrections (for example for orographic enhancements) but he cannot adjust the values of individual rainfall pixels just as he chooses.

A. M. E. RUNACRES. Would Dr Conway comment on the problems encountered in comparing and calibrating the weather radar, which gives instantaneous but areally averaged samples. with raingauges, which give point measurements integrated over a period of time?

B. J. CONWAY. This causes many difficulties in calibrating and evaluating radar rainfall measurements, particularly in convective situations where the rainfall has high spatial and temporal variability. There is no way of eliminating the problem with any practicable network of gauges. In situations of more uniform, widespread rain the comparison is easier, but even then topographic effects mean that individual gauges are valid as calibration tools only within limited local domains, and these domains vary with factors such as the wind direction. FRONTIERS has access to calibration gauge data but the operator must exercise caution about how they are used.

P. L. C. JEYNES (*Oxford Computer Services Ltd*, *U.K.*) What wavelengths do the ground-based radars operate at?

B. J. CONWAY. They operate at 5.6 cm. One or two older radars, now being replaced, operate at 10 cm.

P. L. C. JEYNES. So the radar system essentially relies on measuring the 'Rayleigh' back-scattering from the rain droplets?

B. J. CONWAY. Yes.

P. L. C. JEYNES. Would there be any useful Doppler information in the echo?

B. J. CONWAY. Given a sufficiently dense network of Doppler radars, we could obtain useful information about the horizontal divergence field of the air and therefore its vertical motion, but we do not usually have the severe tornadic storms in this country that make Doppler radar the natural choice in the U.S.A.

P. L. C. JEYNES. I was thinking more in terms of using the Doppler shift to measure the vertical free-fall terminal speed from which one might infer the rain droplet size.

B. J. CONWAY. This would be difficult with current radars which operate near horizontal. The high beam elevations which would be needed for fall-speed measurements would limit the real coverage of such measurements.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

P. L. C. Jeynes. Perhaps the vertical rain motion could be obtained from satellite-borne or other radars?

B. J. Conway. Perhaps, though satellite systems have their own problems, such as resolution, power requirements and, for polar orbiters, the interval between observations. With the ground-based radar system we have to choose carefully how we spend our limited resources. Our current priority is to extend the area of coverage to the whole of the U.K.

MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES
THE ROYAL SOCIETY A
PHILOSOPHICAL TRANSACTIONS OF

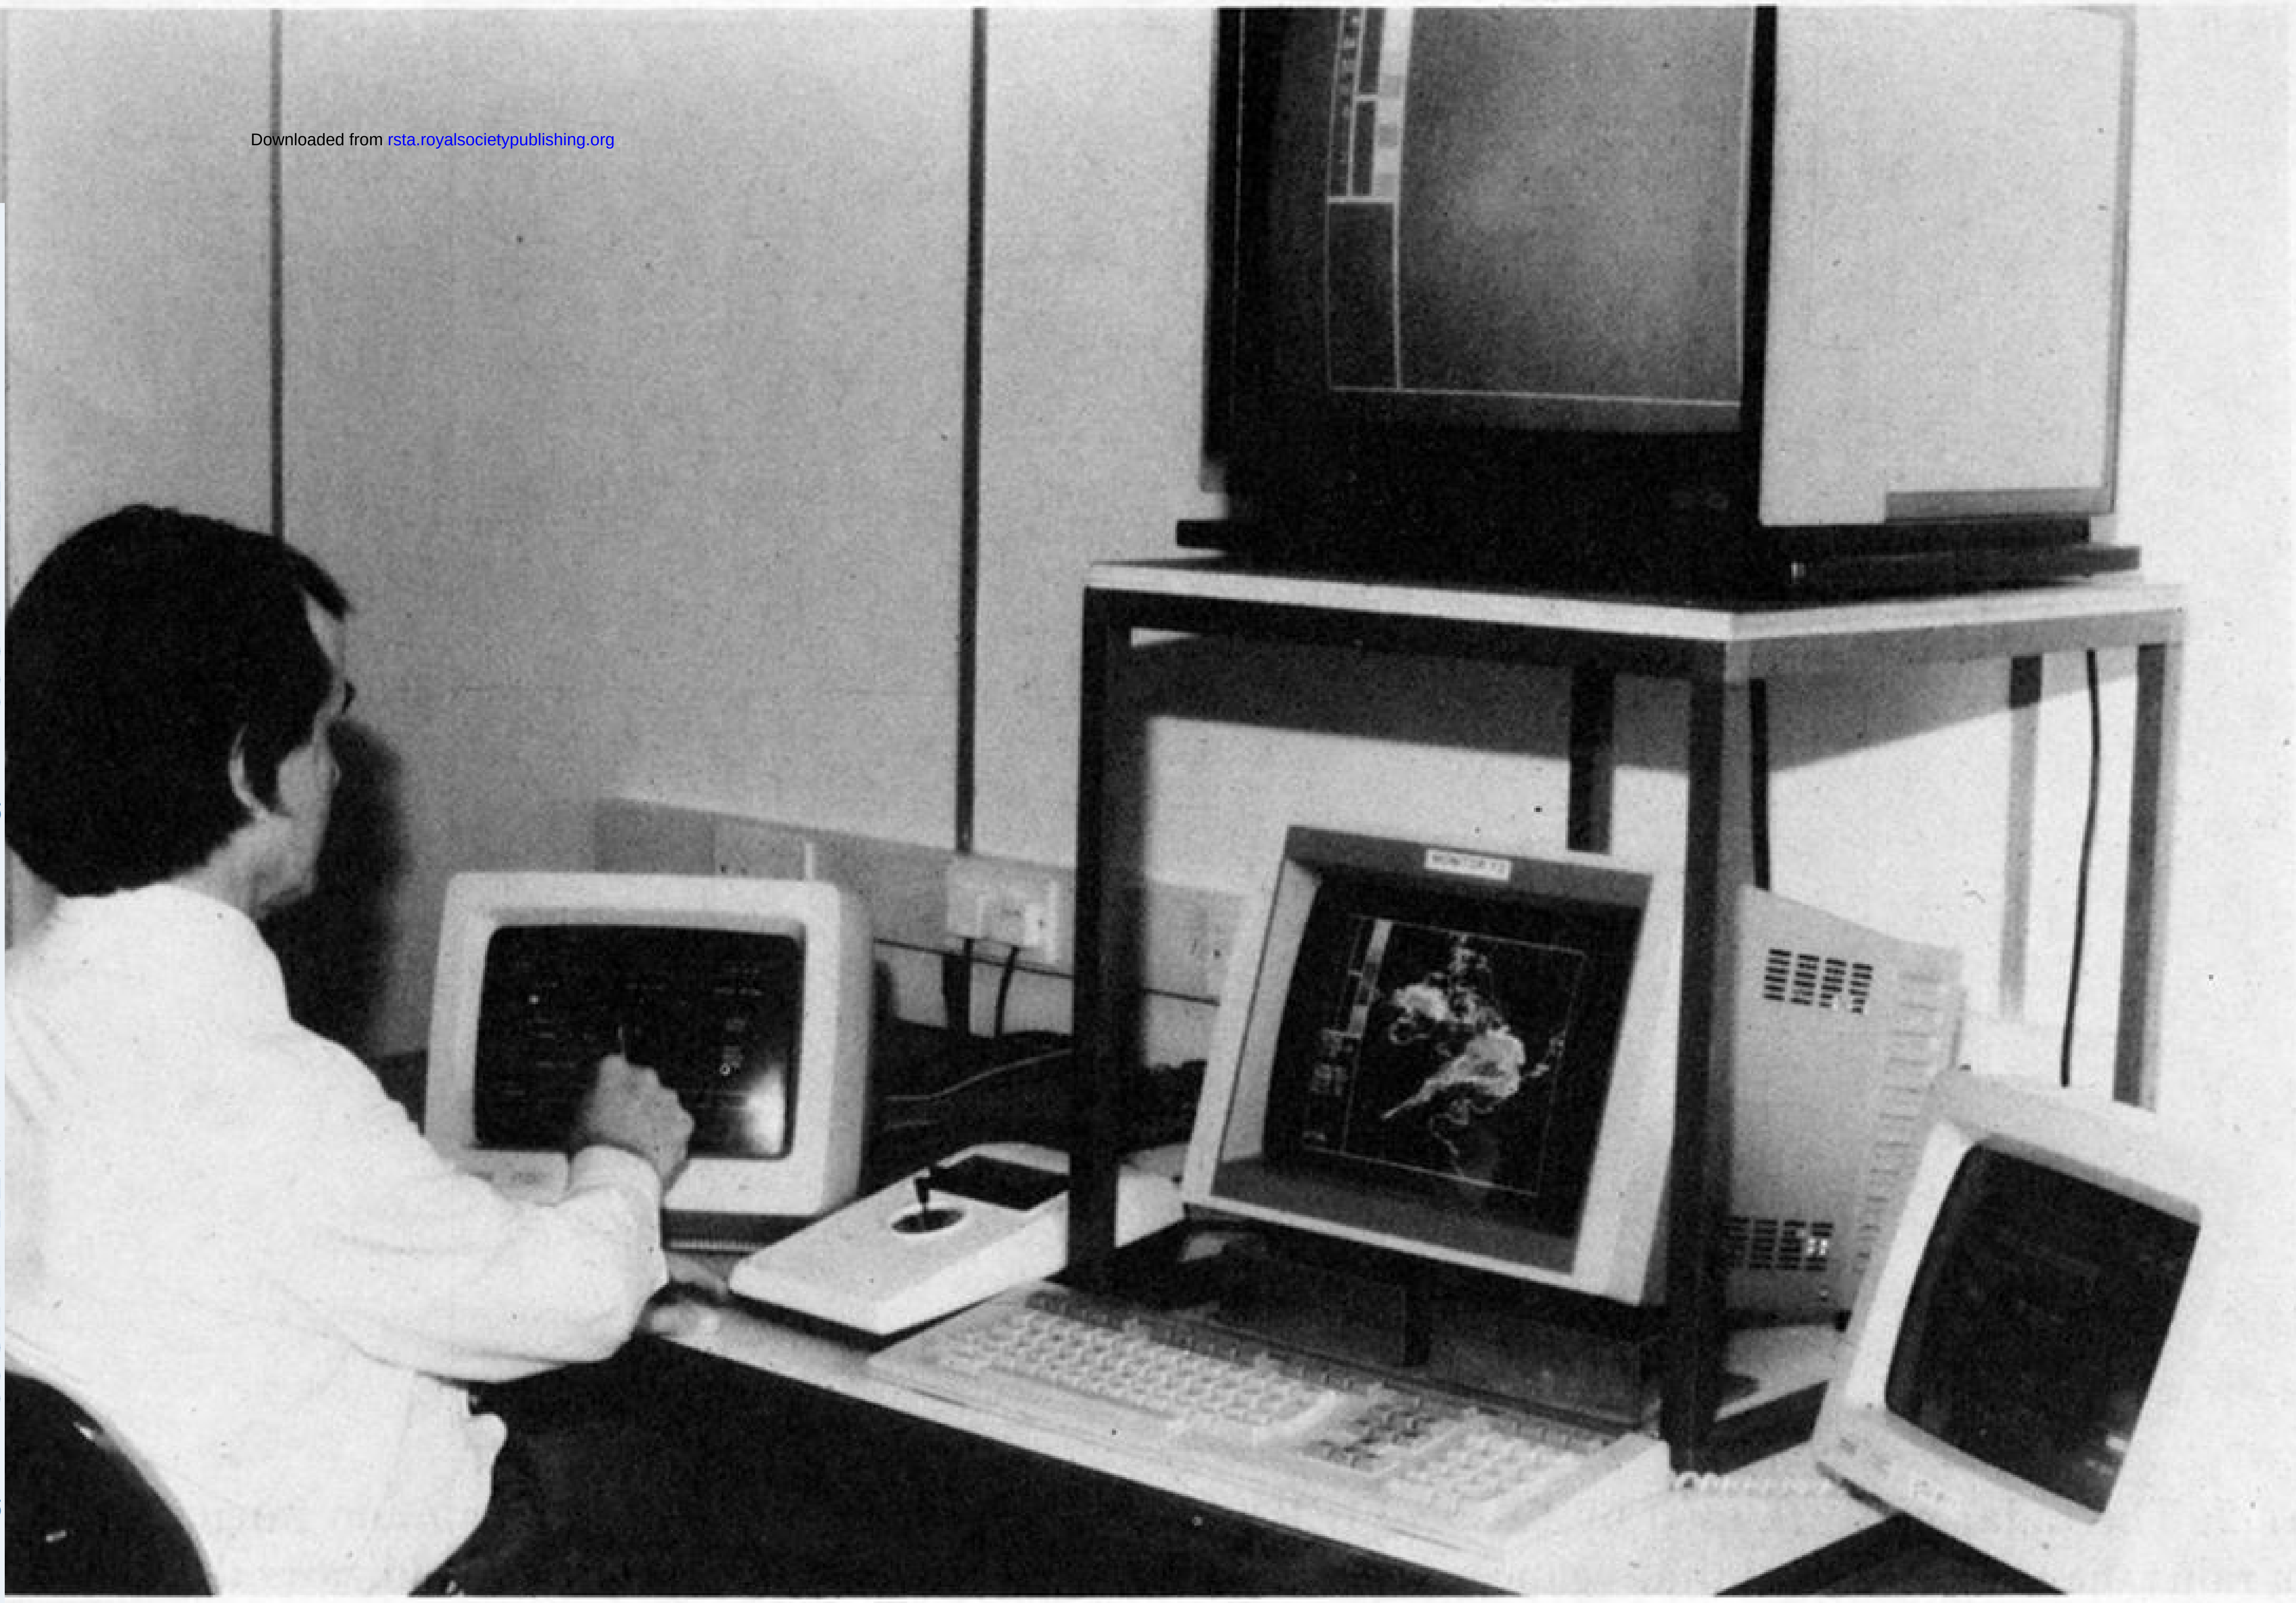

FIGURE 1. The FRONTIERS work station.

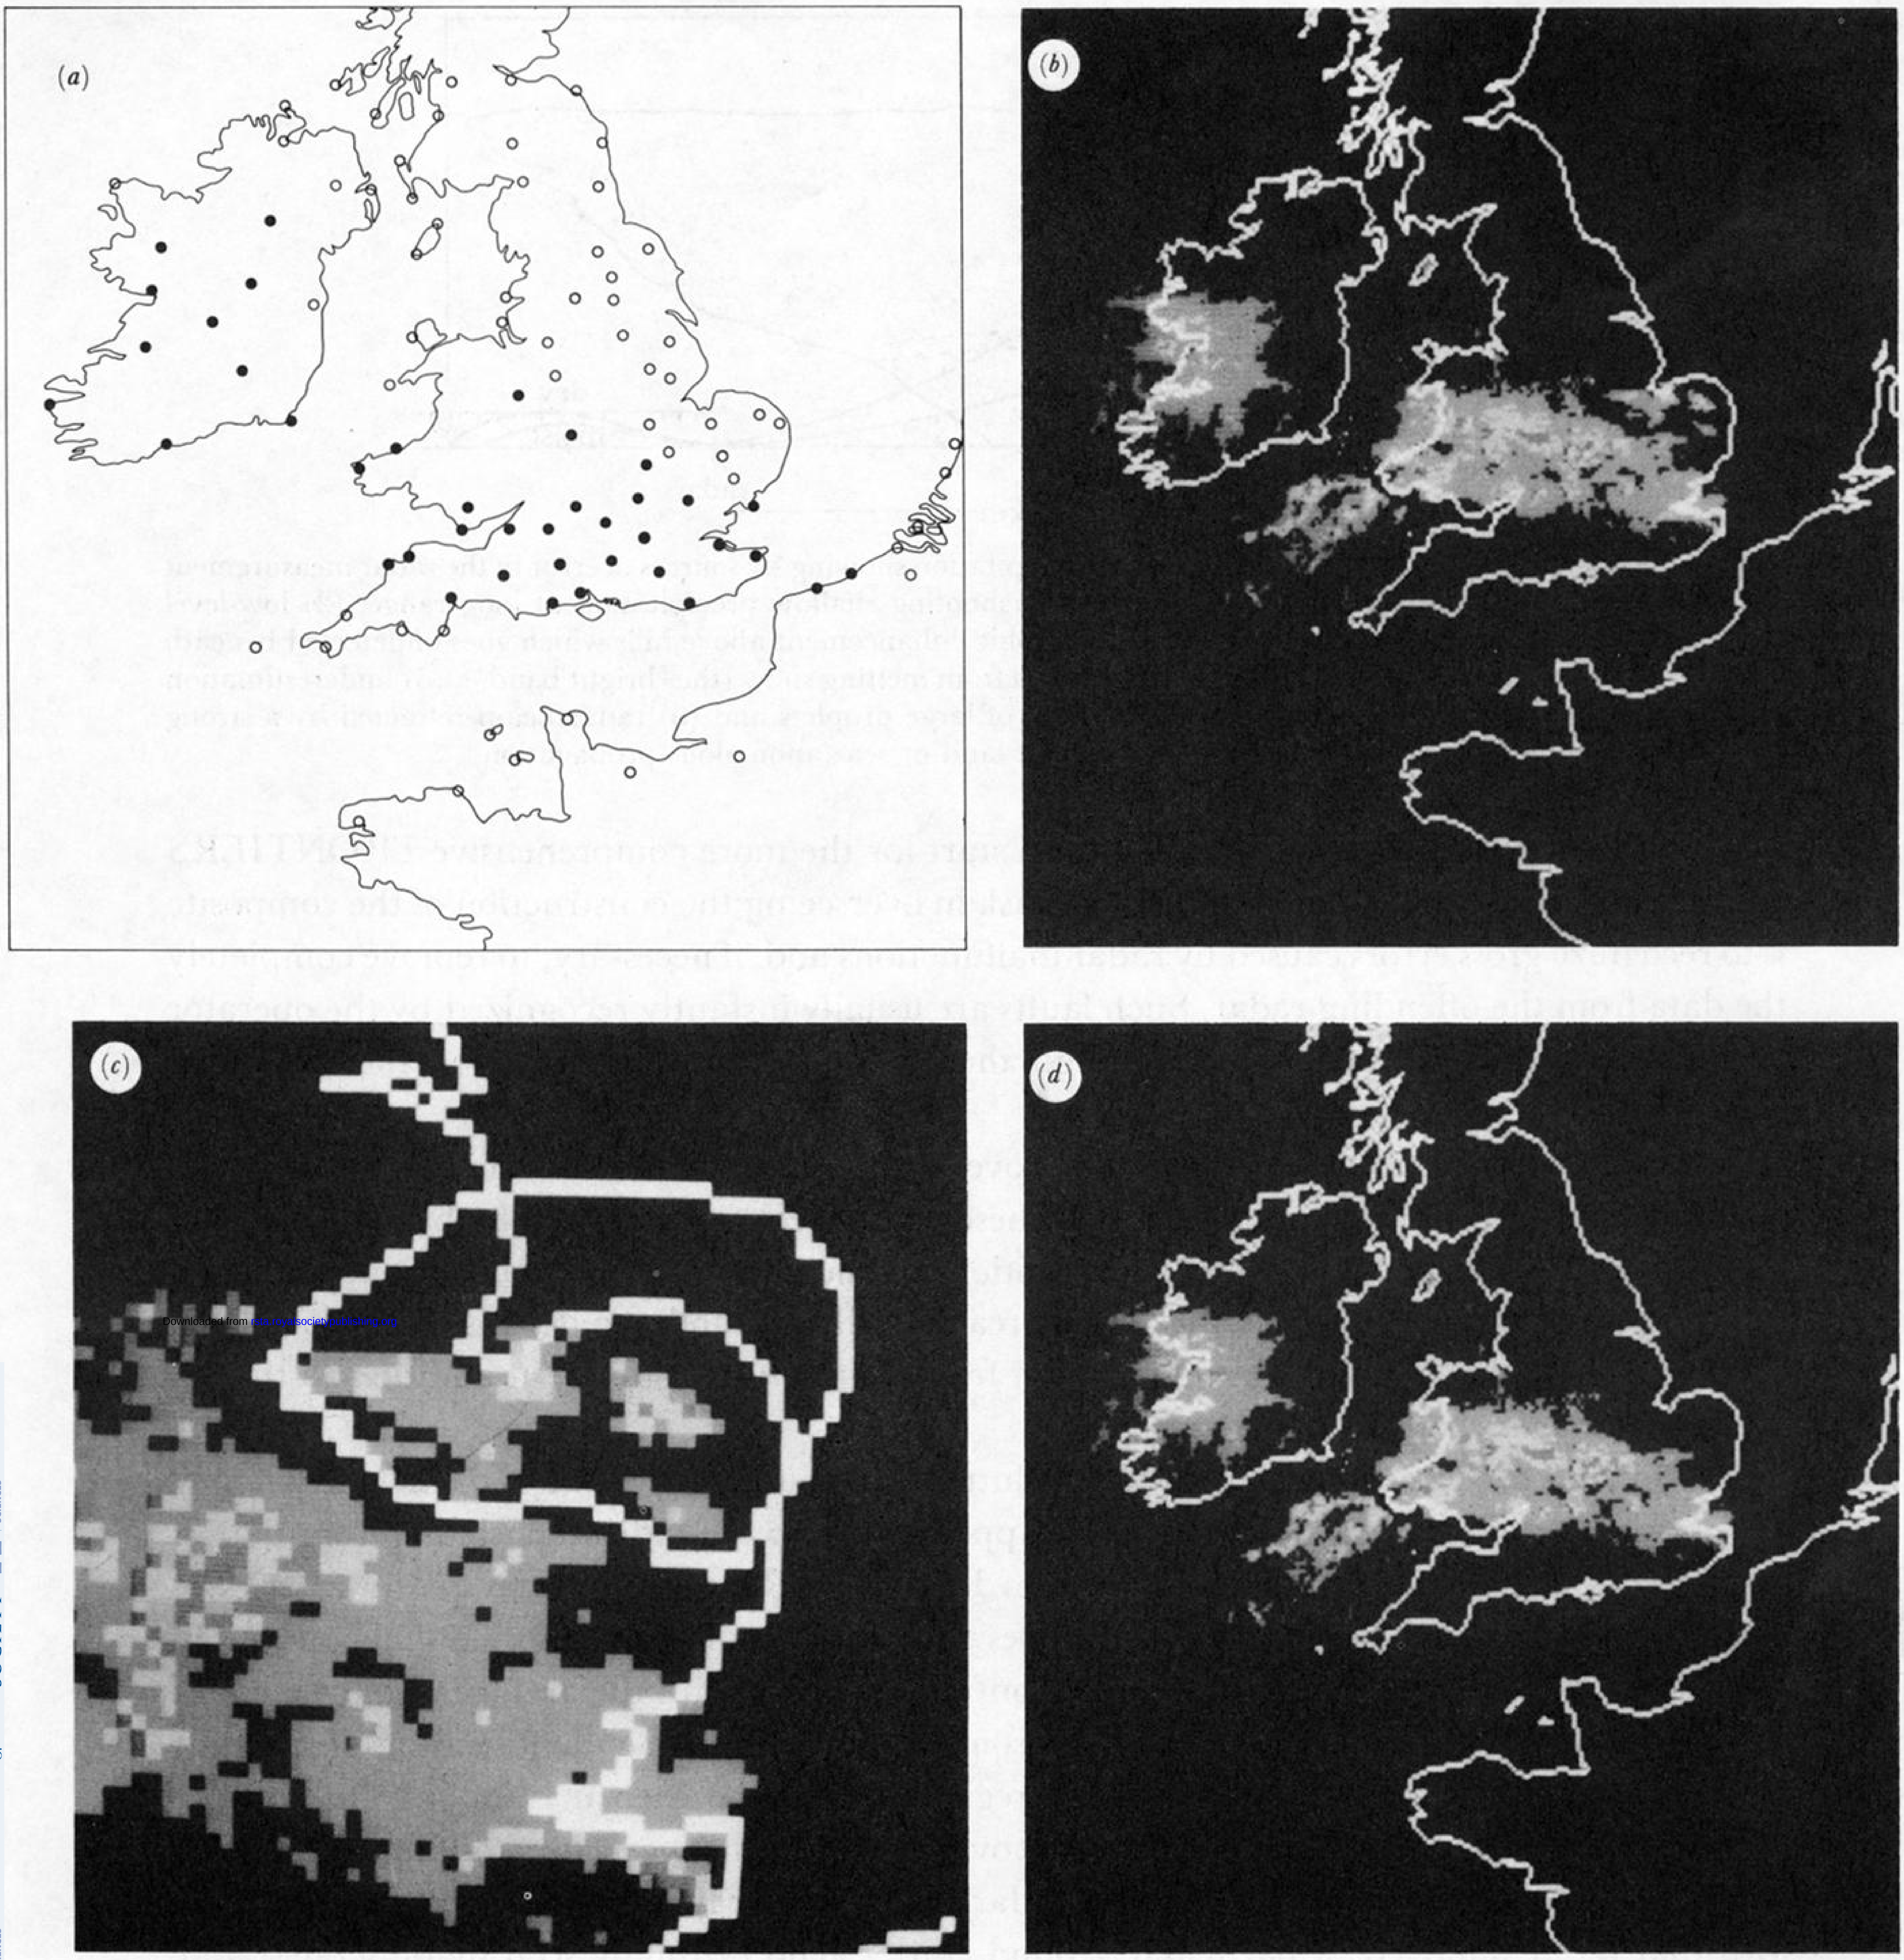

FIGURE 4. Correcting for low-level evaporation of rain in the FRONTIERS radar analysis stage. (*a*) Surface observations (from conventional meteorological stations) of rain (●) and no rain (○) at 03h00 GMT on 9 February 1987. (*b*) Radar composite for the same time as in (*a*). Rain is detected aloft, south of the Wash, by the London radar but is evaporating before reaching the ground. (*c*) The FRONTIERS operator zooms in on the evaporating rain and marks it for deletion by drawing a boundary round it on the screen. (*d*) The radar composite of (*b*) with the evaporating rain removed. The radar is missing some light rain and drizzle near the south coast because of the small drop sizes. The operator gets a chance later in the FRONTIERS cycle to insert rainfall here.

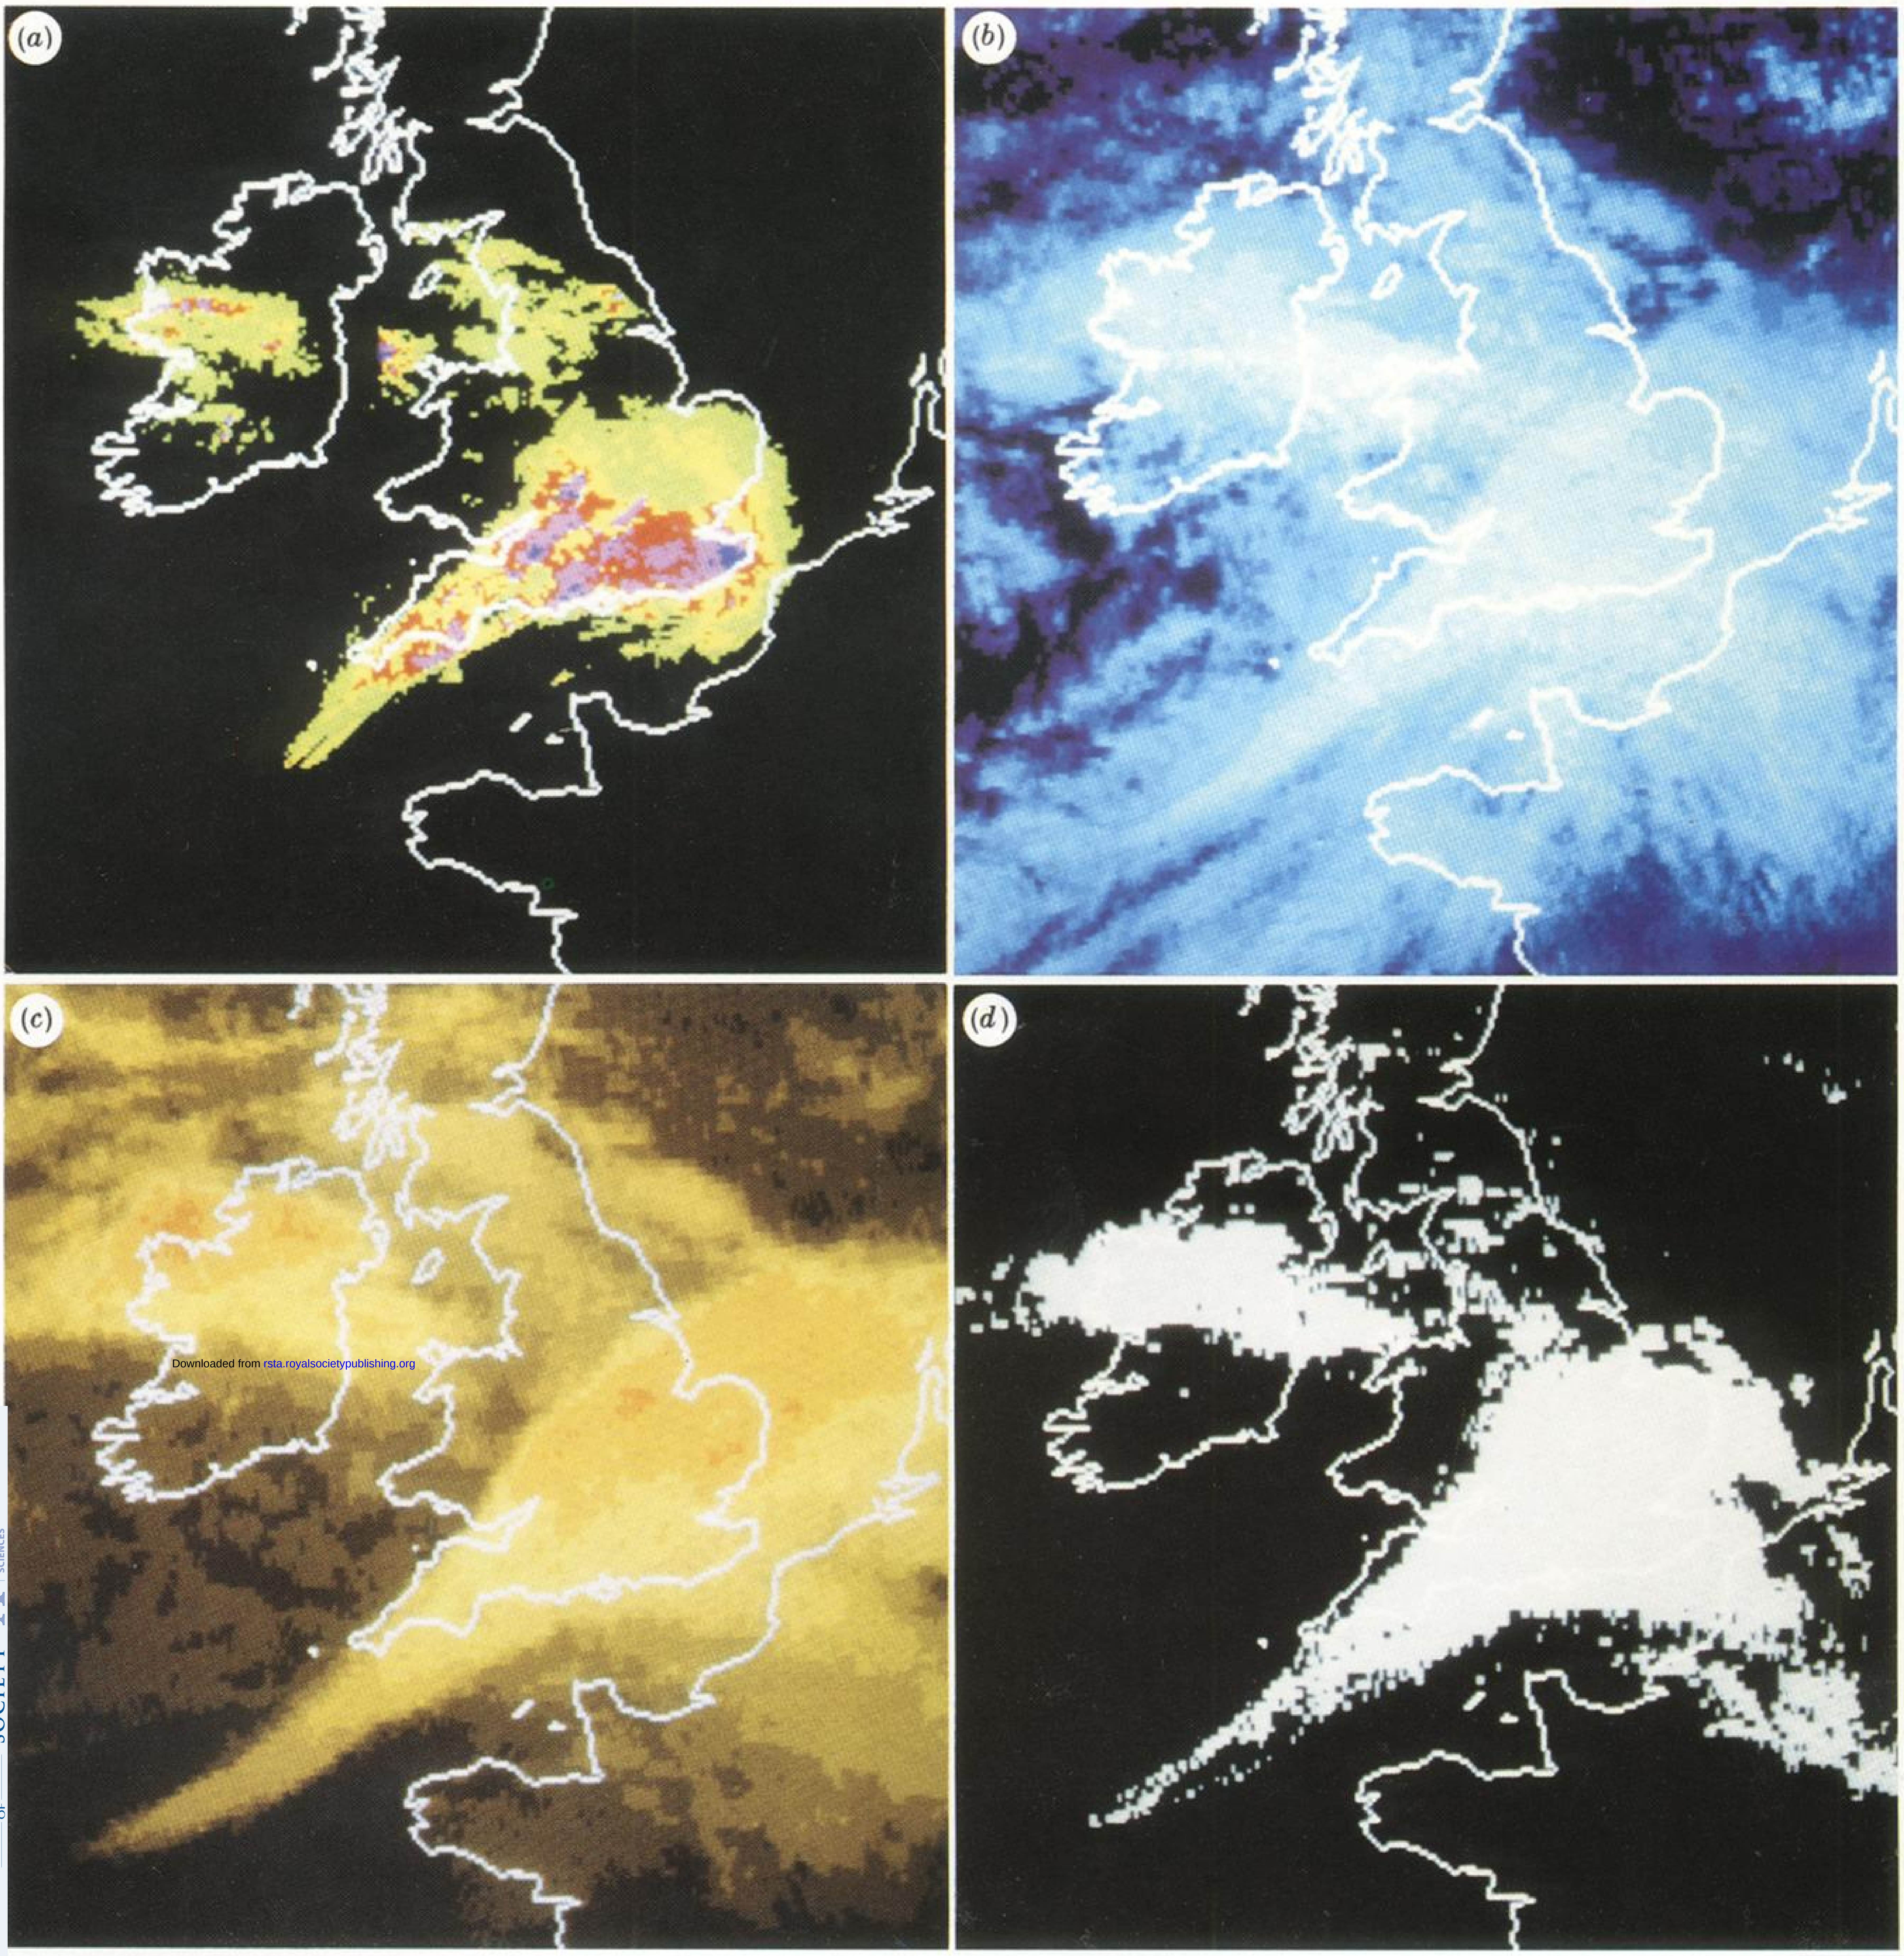

FIGURE 5. Derivation of an extended area rainfall field from *Meteosat* cloud images. (*a*) Radar composite for 13h00 GMT on 4 August 1985, showing observed rainfall rates after correction for errors in the FRONTIERS radar analysis stage. Green shows lightest rain (less than 1 mm h^{-1}), blue shows heaviest (8–16 mm h^{-1}). (*b*) *Meteosat* visible (0.4–1.1 μm) image for the same time as in (*a*). The whitest areas correspond to the brightest and therefore the thickest cloud. (*c*) *Meteosat* infrared (10.5–12.5 μm) image for the same time as in (*a*). Colours show the temperatures and therefore the heights of the cloud tops. The coldest cloud tops in the picture are indicated by the orange patches over eastern England and Northern Ireland. Very dark areas are warm cloud or surface. (*d*) Field of rainfall at 13h00 on 4 August 1985, inferred from the *Meteosat* visible and infrared images by correlating them with the radar rainfall pattern within the radar area. This field shows only the extent, not the intensity, of the rain.

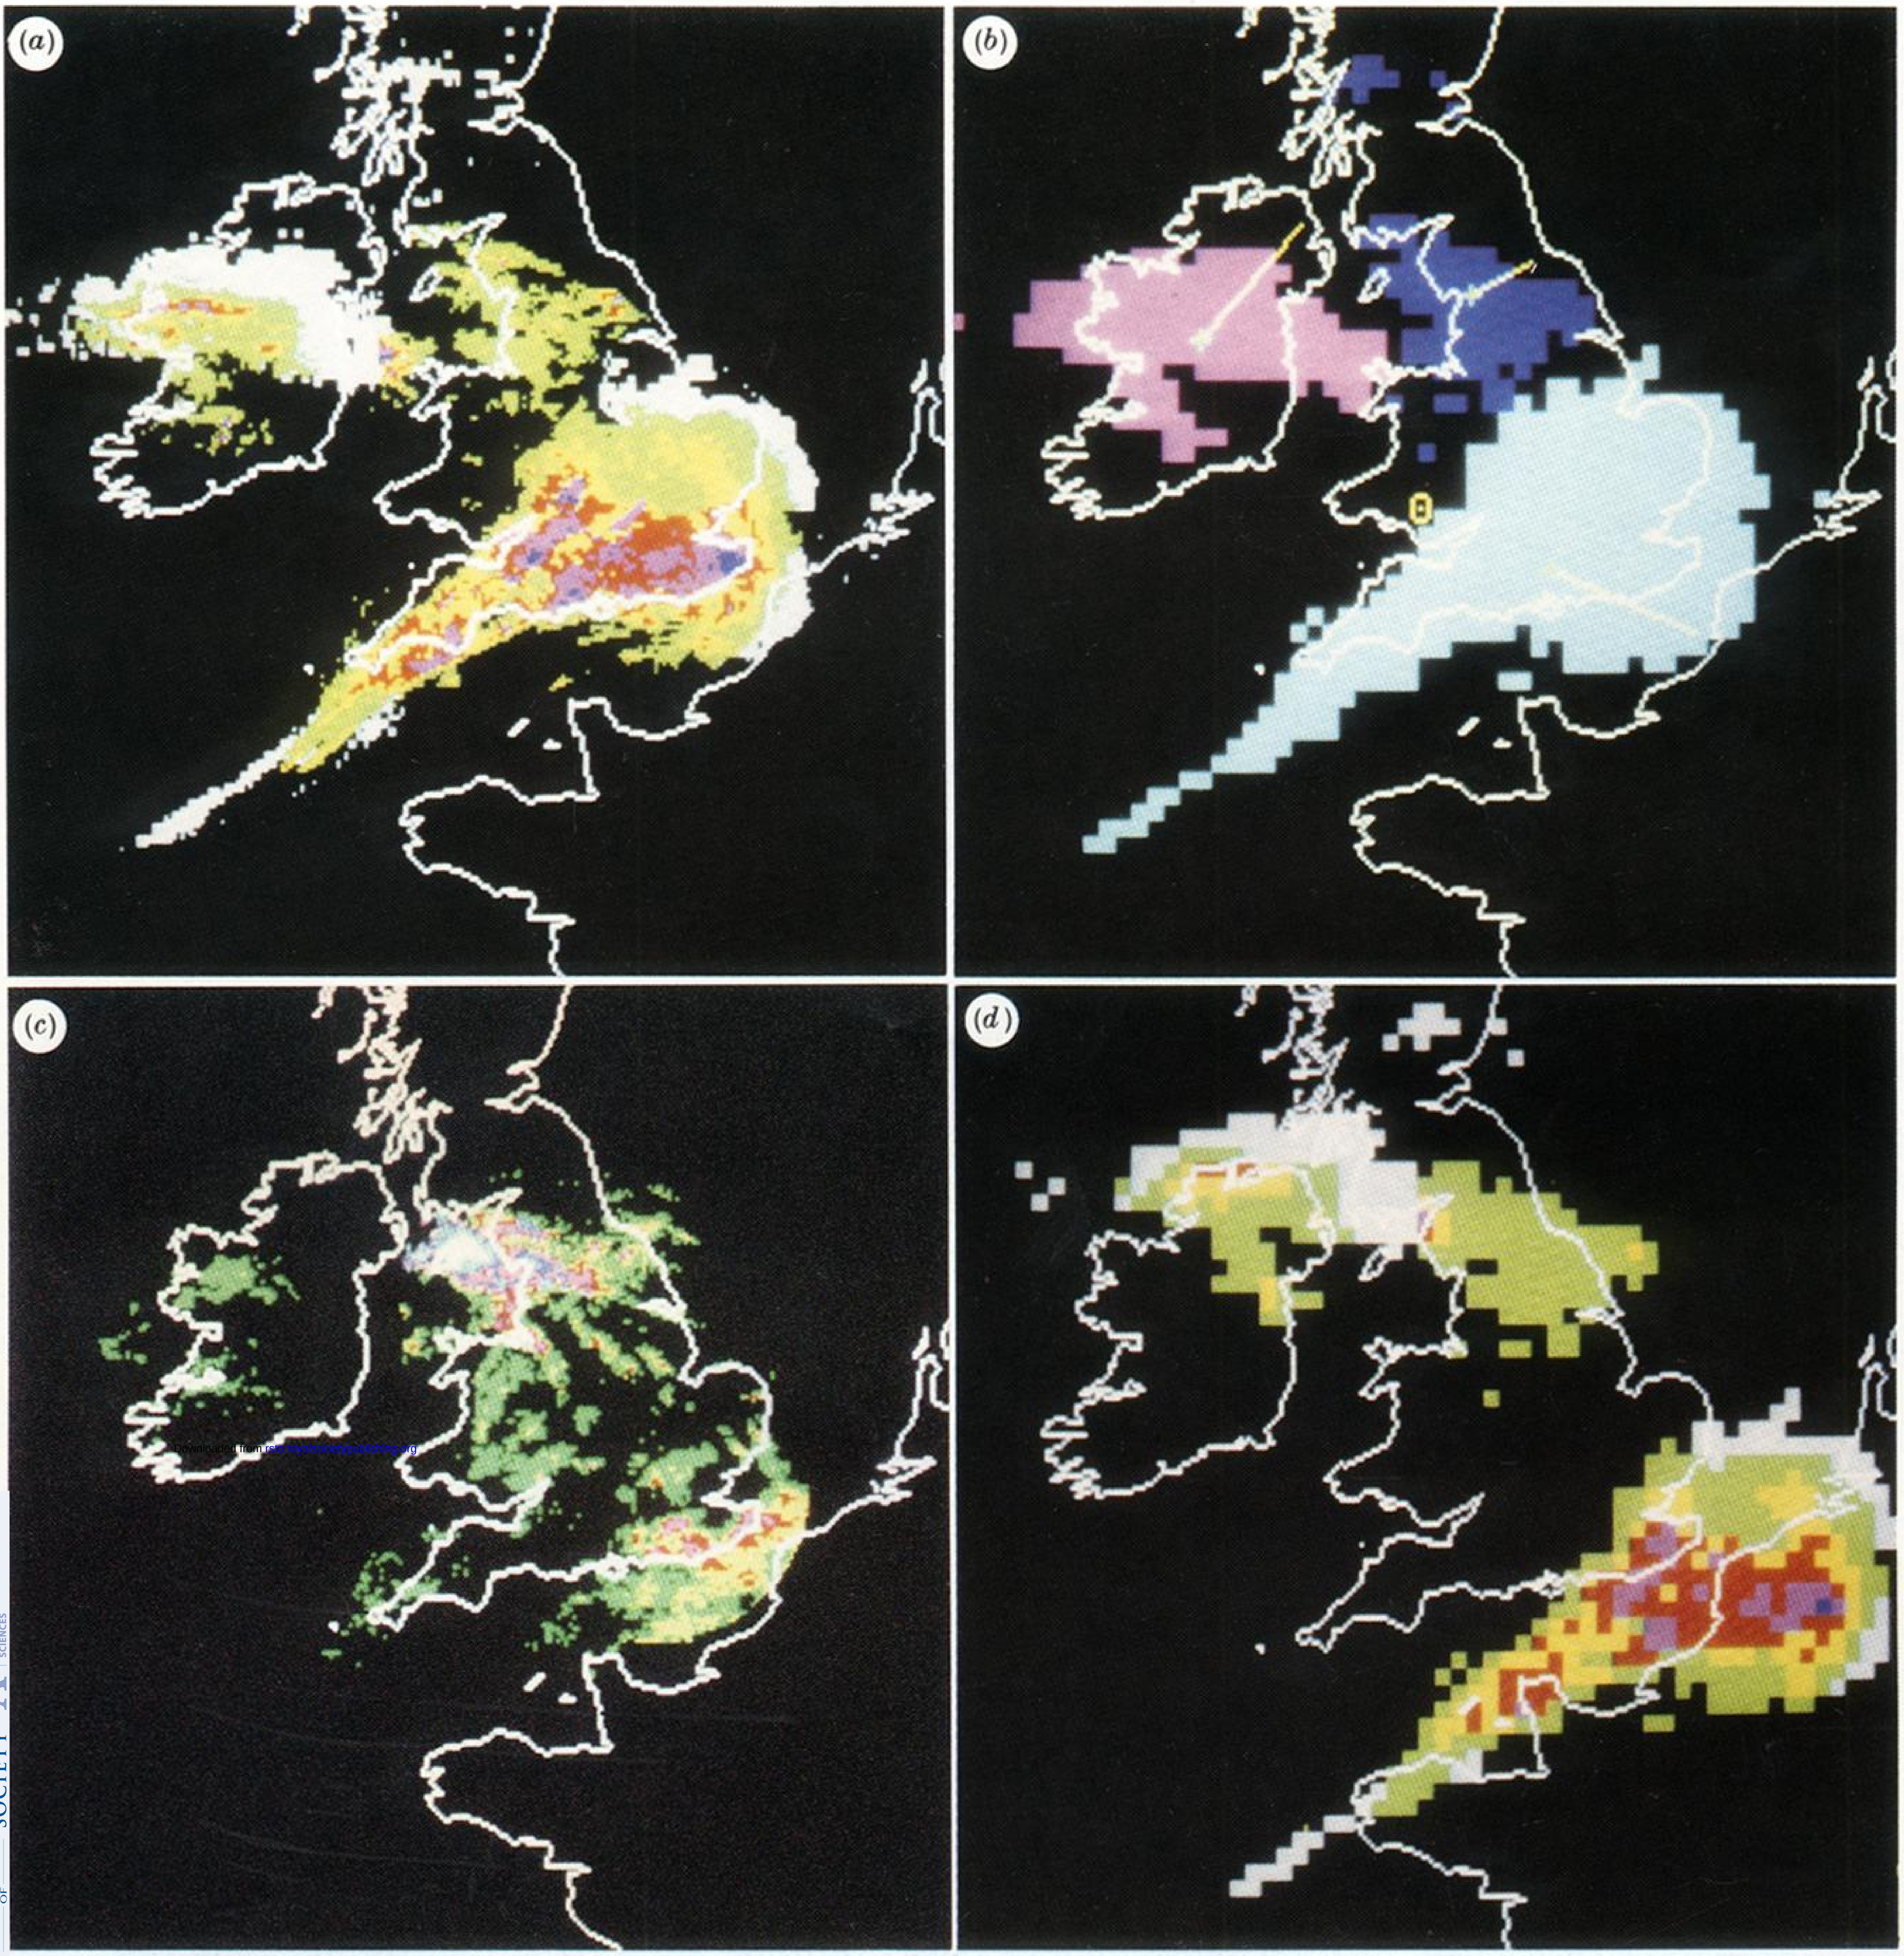

FIGURE 6. Generation of a rainfall forecast. (*a*) The radar–satellite combination rainfall field formed by adding the satellite-derived rainfall field of figure 5*d* to the quantitative radar rainfall field of figure 5*a*. The satellite-derived rainfall (shown in white) has been selectively trimmed by the FRONTIERS operator. (*b*) The rainfall pattern of figure 6*a* divided into clusters that will move independently in the forecast. Their velocities are indicated by the white and yellow vectors at their centroids. (*c*) The rainfall pattern 3 h later, at 16h00 as observed within the area of coverage of the radar network. (*d*) The 3 h extrapolation forecast for 16h00. The heavy rainfall in SE England is well forecast, as is the extent of the rain in northern England and the Scottish border region. The very heavy rain over the Irish Sea would have been missed almost entirely if the satellite data had not been available to bridge a gap in the radar coverage (figures 5*a* and 6*a*). The light rain over the Midlands and Wales developed *in situ* during the forecast period and so could not be forecast by extrapolation techniques. On this occasion the forecast actually gives better agreement with the surface observations in Northern Ireland than does figure 6*c* because the radar at Shannon (SW Ireland) is missing some rain in that area at 16h00 GMT.