3A.4

SHORT-TERM PREDICTION OF HIGH REFLECTIVITY CONTOURS FOR AVIATION SAFETY*

John C. Brasunas and Mark W. Merritt
Lincoln Laboratory
Massachusetts Institute of Technology
Lexington, MA 02173

1. INTRODUCTION

Airspace utilization and safety could benefit significantly from accurate, real-time, short-term predictions of hazardous weather regions (e.g., 5-30 minutes). For some hazards, such as heavy turbulence, the detection process itself is in an immature stage. No universally accepted algorithm exists for indicating the regions of current turbulence - let alone predicting it. For other hazards, such as hail and more particularly for heavy rain, the detection process is in a more mature state. In fact heavy rain may be unambiguously associated with high dBZ (reflectivity), if no ice phases are present. Hail is also associated with high reflectivities.

We have therefore chosen to place our initial emphasis on the prediction of reflectivity contours in the context of ATC (air traffic control) operations. For all of our prediction techniques, we begin by collecting fixed dBZ-level contours on a fixed-elevation scan by fixed-elevation scan basis, and then combining these elevation cell slices into volume cells as is done in the algorithm of Bjerkaas and Forsyth (1980). To these volume cells we attach translation vectors to make the desired prediction: at this time no provision is made for the growth or decay of reflectivity cells.

We generate our translation vectors using each of several algorithms which have already been described elsewhere. Firstly, we use the centroid-tracking approach of Bjerkaas and Forsyth (1980). This is the current tracker of choice in the NEXRAD (Next Generation Weather Radar) program. Secondly, we use tracking vectors of clusters of volume cells, as described by Crane (1979): much of this work was performed under the sponsorship of the Federal Aviation Administration (FAA). Thirdly, we generate translation vectors by cross-correlating low-altitude (0-4 cm) CAPPIs (constant-altitude plan position indicators): this correlation is done either for the entire storm, or for 30 km by 30 km segments of the storm. This approach has been motivated by the work of Rinehart and Garvey (1978), although we generally use a CAPPI of liquid water content. Fourthly, we use as a prediction the current, composite reflectivity map - our so-called status-quo prediction.

2. FALSE ALARMS AND FALSE SAFES

A dBZ contour may be scored in terms of its utility to FAA controllers and to the pilots themselves. The first requirement of a prediction is to indicate a hazardous region - failure to do so constitutes a "false safe". The second requirement is to minimize the region of excluded airspace - airspace falsely deemed dangerous constitutes a "false alarm". The importance of minimizing false alarms was stressed in some recent incidents in which LLWSAS (low-level wind shear alert system) alarms were apparently ignored due to excessive false alarms. Since there are both false alarms and false safes to consider, it is difficult to rank different predictions in terms of a scalar score, unless we know the cost-structure of false alarms and false safes. The ultimate choice of a prediction will probably be that method which minimizes the false safe rate, subject to the constraint that the rate of false alarms not exceed the level at which utilization of safe airspace materially declines. It would additionally be expected that the resulting false safe rate be lower than the current value.

We have decided on two basic approaches to the tallying of false alarms and false safes. The first, which we refer to as area-intersection, involves defining false alarms and false safes on a pixel-by-pixel basis, where a pixel typically has a size of 1 km by 1 km. As we now forecast only a single reflectivity level, typically 30 dBZ, a false alarm occurs when we forecast a pixel to be above 30 dBZ, but when the time comes it is actually below 30 dBZ. The area-intersection score is easily understood: if growth and decay are unimportant, it reflects our success in putting volume cells in the right locations. This is why we use it. However, its operational usefulness to aviation safety

*This work was sponsored by the Federal Aviation Administration. The United States Government assumes no liability for its contents or use thereof.

is not direct, and it is probably more applicable to hydrology, where it may be necessary to predict the rainfall accumulation in a watershed.

Our second approach is to tally false alarms and false safes on a flight-path by flight-path basis. In this approach, if a given flight-path intersects at least one hazardous region, but intersects not a single, predicted hazardous region, this is one false safe. A similar definition holds for false alarms. Figure 1 shows a hypothetical example. We have, so far, generated our flight-paths by a Monte-Carlo approach, so that almost each pixel has many straight paths going through it in all directions. A future approach will be to use actual airways, either en-route or terminal.

FLIGHT PATHS

FA - False Alarm
FS - False Safe
S - Success

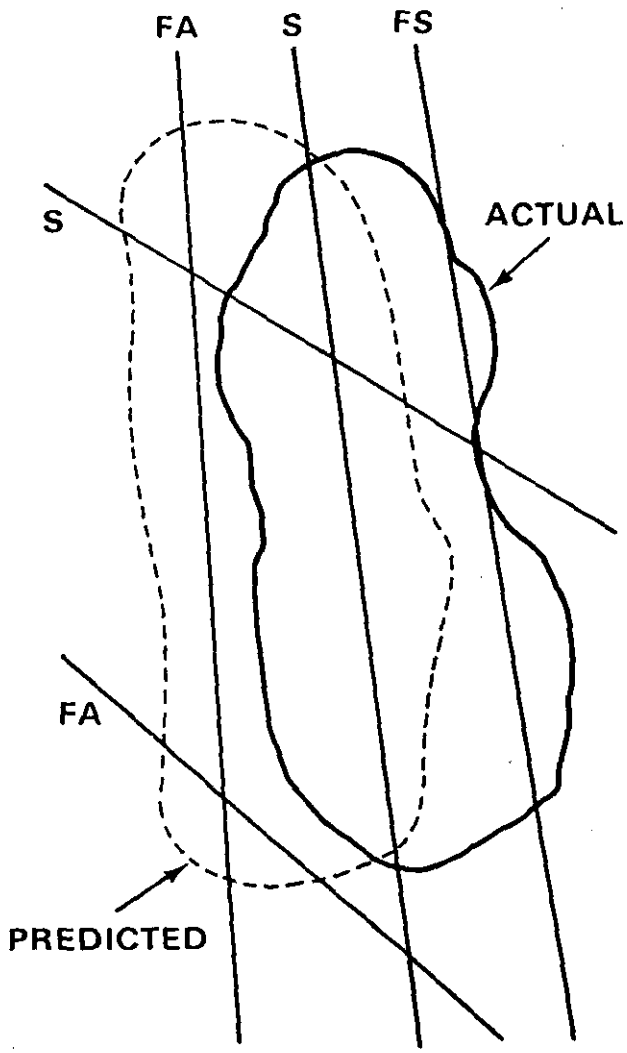

FIGURE 1

3. RANKING THE PREDICTION METHODS

Since there is no "typical" thunderstorm, we need to have data from a wide variety of locations and seasons in order to rank the various prediction methods. Therefore we have set up a procedure of converting raw data tapes to an intermediate format, our so-called common format tapes. To generate a common format tape we require the raw data be in a volume scan format - a collection of around six or more scans in azimuth, each one at a higher elevation angle. The volume scan update rate should be around five minutes. We have thus far reduced data from an MIT radar (S-band pulsed Doppler, pencil-beam radar with capability of estimating first three lags) and from the Norman and Cimarron radars at NSSL (the National Severe Storms Laboratory).

To score a prediction, we first generate a set of "truth" maps - composite reflectivity maps on a Cartesian grid with a bin size of 1 km by 1 km. For our work at a single reflectivity level, a bin is true if any of the raw radar cells at that x-y location, with altitude in the range of 0-4 km, exceeds the threshold. Typically we have worked at 30 dBZ because this is a favorable level for the centroid tracker. To make the prediction itself we take the volume cells, attach the appropriate translation vectors, and move the volume cells to the predicted location. We then make a composite reflectivity map, not from the raw radar cells but from the contoured volume cells. At this point we have two composite reflectivity maps, the truth and the prediction. False alarms and false safes are tallied in terms of area-intersection and flight-paths. Performance differences are due entirely to the translation vectors.

Our experience with the MIT and NSSL data sets has pointed out some difficulties in designing a set of algorithms to work automatically on data from a variety of sources. Particular problems occur when the volume scans cover a sector rather than the full 360 degrees. In this case the azimuthal limits may change, either within a volume scan or from volume scan to volume scan. Or there may be an additional azimuthal scan, perhaps with different limits. Another problem has been data runs insufficiently long to both initialize the tracker and to provide a truth map 30 minutes hence.

4. EVALUATION OF PARTICULAR STORMS

4.1 August 5, 1981 Storm at MIT

To date, most of our analysis has centered on this August storm in Boston - a squall-line thunderstorm with peak reflectivities exceeding 55 dBZ. Our data set consists of several hours of volume scans at 6 elevation angles, for the full 360 degrees, taken roughly every 6.5 minutes. The data have been subjected to clutter filtering, and all data within 30 km of the radar and below 1.5 degrees elevation are censored because of

the heavy clutter environment. Figure 2 shows the truth map for scan 17, at which time we began to make reflectivity predictions (30 dBZ). For Figure 2 and subsequent figures, the map center is at MIT. The map scale is ±128 km in the East-West and North-South directions. For the velocity maps, the cross-bars are at the vector tails. Figure 3 shows cross-correlation velocities obtained by cross-correlating CAPPIs from scans 16 and 17; figure 4 shows cluster velocities obtained from applying the Crane algorithm to scans 15, 16 and 17. Figure 5 shows a 32.5 minute prediction based on the volume cells and cross-correlation velocities valid for scan 17 - this extrapolation time corresponds to volume scan 22, whose truth map is shown in Figure 6.

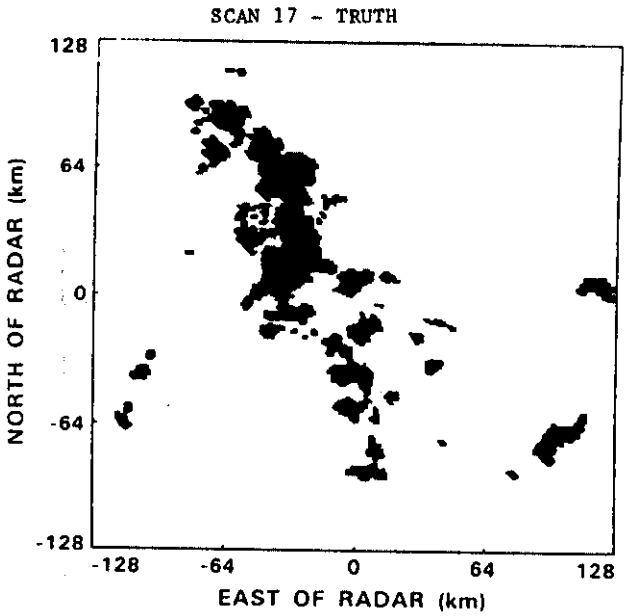

Fig. 2

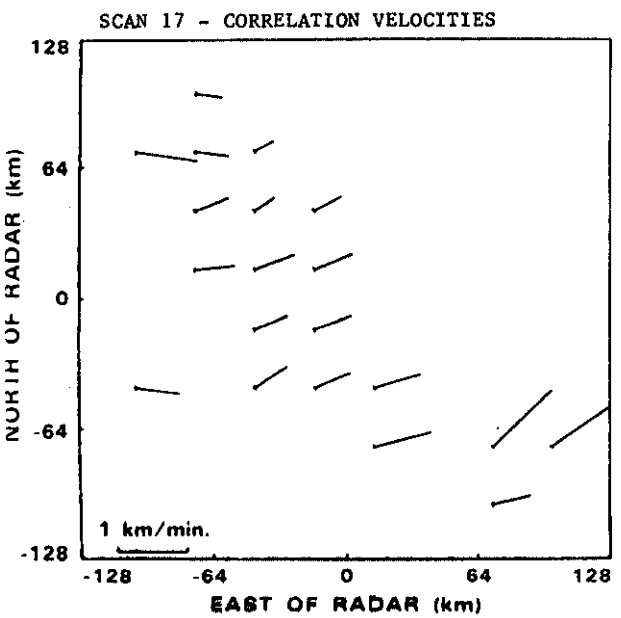

Fig. 3

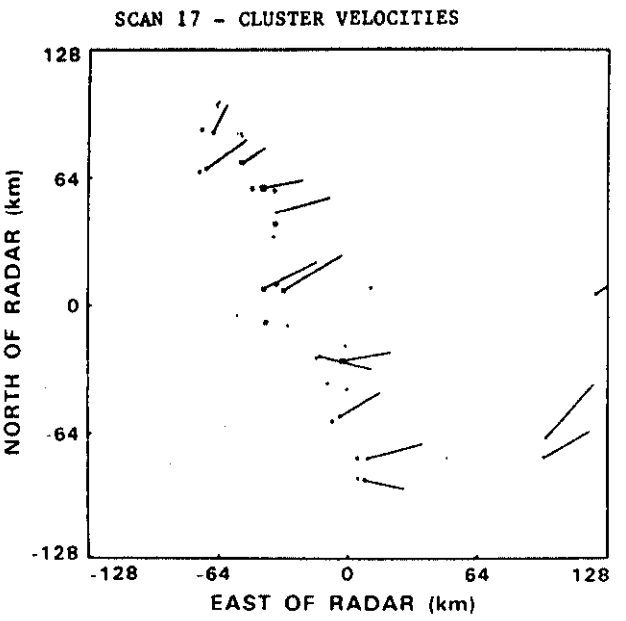

Fig. 4

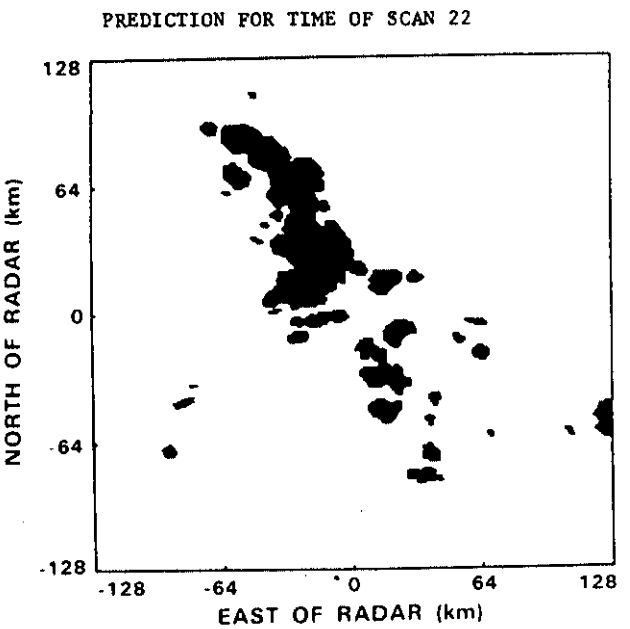

Fig. 5

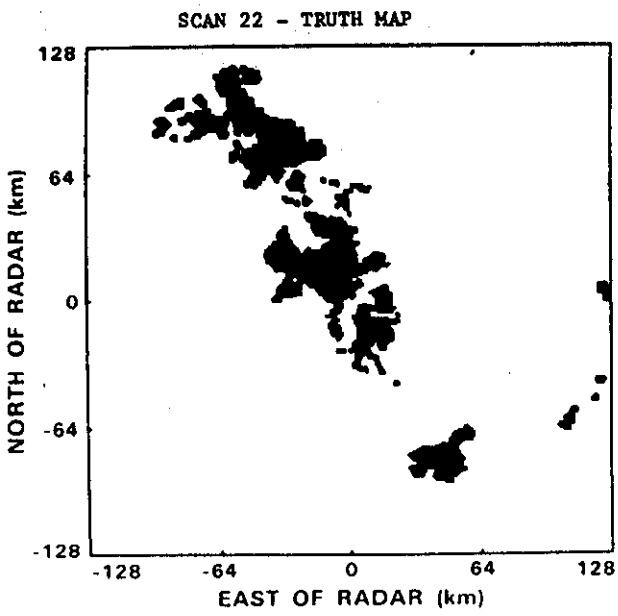

Fig. 6

There is quite clearly rough agreement between correlation and cluster velocities, although the cross-correlation is done on a spatial scale of 30 km by 30 km, whereas the cluster pertains to a scale of a few km. The prediction map in Figure 5 agrees fairly well with the truth map in Figure 6 in terms of location, but quite clearly there has been some growth and decay which we do not now attempt to predict.

Table 1 gives the false safe and false alarm results for the various trackers. NFS and NFA are respectively the number of false safes and the number of false alarms

Table 1
Statistics for a 32.5 Minute Prediction
(* indicates 27 minute)

	NFS	NFA
Area-intersection (±0.005)		
Status-Quo	.055	.065
Untracked Cells	.060	.055
Centroid Track	.060	.055
Single-Vector	.050	.045
Multi-Vector	.045	.045
Multi-Vector*	.040	.040
Crane*	.045	.040
Flight Paths (±0.01)		
Status-Quo	.06	.16
Untracked Cell	.06	.14
Centroid Tracking	.10	.12
Single-Vector	.08	.10
Multi-Vector	.12	.08
Multi-Vector*	.11	.08
Crane*	.16	.04

(normalized either by the total number of pixels or flight-paths). "Untracked cells" are the volume cells with zero velocity translation vectors attached. "Single vector" is where we cross-correlate the entire storm, whereas in "multi-vector" we cross-correlate individual 30 km by 30 km segments of the storm.

4.2 June 19, 1980 NSSL Storm

Data analysis has proceeded less far on the NSSL storm. It has been hampered, first of all, by the sector-scan format of the data. In any case, it is apparent that over a span of 15 minutes there has been sufficient dBZ decline so that the assumption of constant dBZ will lead to a very poor prediction, even for a threshold set as low as 25 dBZ. Whereas the predictions themselves will receive a poor score because of the unanticipated threshold crossing, the cross-correlation vectors are probably not greatly affected.

5. INTERIM CONCLUSIONS

From Table 1 we can draw several conclusions, which must be considered tentative because this pertains to only a single storm. Looking first at the area-intersection results, the status-quo predictor was almost the worst in terms of both false alarms and false safes. Centroid-tracking performed no better than stationary volume cells. The cross-correlation and Crane predictions both seemed to do somewhat better than the others. In terms of flight-paths, there is essentially a trade-off between false alarms and false safes, with the lowest number of false safes going with no tracking at all. There would seem to be a slight preference for single-vector correlation tracking if we assign equal costs to false safes and false alarms.

There has been a fair amount of success in predicting the future locations of dBZ contours, and this is undoubtedly due to the persistence of the synoptic-scale winds that provide the steering for these storms. Indeed, the predominant direction of Boston-area storms was seen to be about the same on August 11, 1981. On the other hand, analysis of NSSL data has shown the existence of dBZ trends which, coupled with a dBZ threshold for safe/dangerous decisions, can lead to very poor predictions unless we attempt to predict dBZ trends. Initial analysis of MIT data, however, indicates that dBZ trends are not very persistent, even on the scale of 10 minutes for a 30 km by 30 km patch of storm.

6. SUMMARY AND WORK IN PROGRESS

We have presented some contemporary tracking techniques for predicting hazardous regions, together with evaluation criteria related to short-term flight-path choice when in or near convective storms. Our evaluation provides an explicit tradeoff between failing to identify hazardous regions, and "crying wolf" so often that systems lose credibility.

These initial analyses have pointed out the necessity, first and foremost, of extending this study to more storms and to more sites, perhaps including the JAWS and NIMROD data. For a final product more dBZ levels have to be predicted, perhaps three. In addition to generating flight-paths through a Monte-Carlo technique, actual flight paths can be incorporated in the analysis. The persistence of wind direction and magnitude in the Boston area suggests that a reasonable baseline prediction against which all others might be compared would be a "climatological prediction", a translation of about 8 ms^{-1} in the direction of east-northeast.

A major enhancement to our approach will be to include dBZ changes. The dBZ trend itself does not appear to be terribly persistent. What we will probably need to do is look for features, be they convergences, updrafts, dBZ gradients, topography, etc., that can serve as indicators for growth and decay. Perhaps the most complete approach would be to evolve an actual model of the thunderstorm that is initialized by current observations, but we may lack easy access to certain key parameters like temperature and humidity profiles that would be of obvious use to such a model.

REFERENCES

1. Bjerkaas, C.L., and Forsyth, D.E., 1980: Operational Test of a Three-dimensional Echo Tracking Program. Preprints 19th Conf. Radar Meteor., Miami Beach, Amer. Meteor. Soc., 244-247.

2. Crane, R.K., 1979: Automatic Cell Detection and Tracking. IEEE Trans. Geosci, Elec., GE-17, 250-262.

3. Rinehart, R.E., and Garvey, E.T., 1978: Three-dimensional Storm Motion Detection by Conventional Weather Radar. Nature, 273, 287-289.