AUTOMATED GUST FRONT DETECTION USING KNOWLEDGE-BASED SIGNAL PROCESSING*

Richard L. Delanoy and Seth W. Troxel
Lincoln Laboratory, Massachusetts Institute of Technology

ABSTRACT

For reasons of aviation safety and airport operations efficiency, gust front detection and tracking is an important product of Doppler weather radars developed for use in airport terminal areas. Previous gust front algorithms, which have relied on the detection of one or two conspicuous signatures in Doppler radar imagery, have worked reasonably well in images generated by the high-resolution, pencil-beam Terminal Doppler Weather Radar (TDWR). The latest Airport Surveillance Radar, enhanced with a Wind Shear Processor (ASR-9 WSP), is being developed as a less expensive alternative weather radar. Although gust fronts are visible to human observers in ASR-9 WSP imagery, the lower sensitivity and less reliable Doppler measurements of this radar make automated gust front detection a much more challenging problem.

Using machine intelligence and knowledge-based signal processing techniques developed in the context of automatic target recognition, a Machine Intelligent Gust Front Algorithm (MIGFA) has been constructed that is radically different from the previous algorithms. Developed initially for use with ASR-9 WSP data, MIGFA substantially outperforms a state-of-the-art gust front detection algorithm based on earlier approaches. These results also indirectly suggest that MIGFA performance may be nearly as good as human performance. Preliminary results of an operational test period (2 months, approximately 15000 scans processed) during 1992 in Orlando, Florida are presented.

INTRODUCTION

A gust front is the leading edge of a cold air outflow from a thunderstorm. The outflow, which is deflected horizontally at the ground, may propagate many miles ahead of the generating thunderstorm. Gust fronts can have a significant impact on air terminal operations since they often produce pronounced changes in wind speed and direction, forcing a change of active runway and a rerouting of aircraft already in the terminal area. In addition, turbulence and wind shear along the front can be hazardous to aircraft. Reliable detection and forecasting of gust fronts would both improve air safety and reduce costly delays. The Federal Aviation Administration (FAA) has sponsored research and development of automated gust front detection algorithms to be included as critical components of a suite of hazardous weather detection capabilities for the Airport Surveillance Radar with Wind Shear Processor (ASR-9 WSP) systems and the Terminal Doppler Weather Radar (TDWR).

Gust fronts produce signatures that are observable to varying degrees in weather reflectivity and Doppler velocity data generated by these radars. In Doppler velocity images, gust fronts are recognizable as boundaries between converging velocities. In reflectivity images, gust fronts appear as thin lines of increased intensity, which occur as the result of insects, dust, and debris being lofted and concentrated at the leading edge of the front. Based on an approach developed nearly 10 years ago (Ref. [1]), existing automated gust front detection algorithms, such as the Gust Front Detection Algorithm (GFDA; Refs. [2, 3, 4]) and the Advanced Gust Front Algorithm (AGFA; Refs. [5, 6]), have achieved respectable levels of performance in TDWR imagery using only these signatures. However, many fronts or parts of fronts are missed due to the ambiguous or conditional nature of the observable signatures. For example, convergence signatures disappear as gust fronts become radially aligned and assume a bad Doppler viewing angle. Also, reflectivity thin lines can be obscured by low altitude precipitation.

The ASR-9 WSP provides a less expensive alternative to the TDWR as a terminal weather radar (Ref. [7]). Although not originally intended for weather imaging, this fan-beam Doppler radar generates images of sufficient quality that gust fronts can be identified and tracked. However, versions of AGFA adapted for ASR-9 WSP data have performed poorly. The lack of performance is due primarily to the reduced gain and lowered sensitivity inherent in the fan-beam design of the ASR-9. With lowered sensitivity, clear air velocity estimates are unreliable, making most convergence signatures invisible. Consequently, AGFA is forced to rely on only the thin-line portion of the algorithm. To make matters worse, the reduced sensitivity also makes faint thin-line signatures more fragmented and harder to resolve from background.

With both radar systems, humans detect gust fronts using perceptual skills that have been notoriously and surprisingly difficult to implement in computer vision systems. These skills include:

1. The use of motion as a signature for detection.
2. The effective use of spatial and temporal context.
3. The ability to maintain and assimilate weak, uncertain, ambiguous, and even contradictory evidence.
4. The ability to conditionally fuse information from various sources, reflecting knowledge that different signatures can have varying reliability that depends on situational context.
5. The use of knowledge of weather patterns and trends.

This paper introduces a new machine intelligence approach to detecting and tracking weather phenomena, in this case as applied to the development of a Machine Intelligent Gust Front Algorithm (MIGFA) for use with ASR-9 WSP imagery. MIGFA is a direct adaptation of the eXperimental Target Recognition System (XTRS), a general-purpose approach to object recognition that was initially developed in the context of automatic target recognition (Refs. [8, 9, 10, 11]).

The conventional wisdom in computer vision/object recognition research has been to use general image processing operations, ideally devoid of object- and context-dependent knowledge, at the initial stages of processing. Such operations might include edge detection, segmentation, cleaning, or optic flow estimations. And yet, the ideal is never really achieved in practice. For example, in order to effectively detect edges, some knowledge of the sensor and the expected scene contents are (implicitly) encoded in the form of thresholds or other such parameters. From the results of such general operations, image characteristics are extracted and represented symbolically. Machine intelligence is then applied, as if by definition, only on the symbolic representations at "higher" levels of processing.

In contrast, sensor-, object-, and context-dependent knowledge is applied in the earliest (image processing) levels of XTRS and MIGFA processing. Knowledge of the task is used in three ways. Knowledge is used to choose from a library those feature detectors that are *selectively indicative* of the object being sought. The selected set of feature detectors can differ, depending on environmental context. Knowledge is also

*The work described has been sponsored by the Federal Aviation Administration. The U.S. Government assumes no liability for its contents or use thereof.

0-7803-0934-0/93/$3.00 © 1993 IEEE

incorporated within feature detectors through the design of matched filters that are customized to the physical properties of the sensor, the environment, and the object to be detected. Finally, knowledge of the varying reliability of the selected feature detectors is used to guide data fusion.

KNOWLEDGE-BASED SIGNAL PROCESSING

Most feature detectors in MIGFA are constructed using a new technique of knowledge-based signal processing, called functional template correlation (FTC). FTC is a generalized matched filter incorporating aspects of fuzzy set theory (Ref. [10]). For comparison, standard 2-D cross correlation uses a kernel that is essentially a subimage consisting of expected image values. In contrast, the kernel of a functional template consists of a set of integers that each correspond to a unique scoring function. Each scoring function, given an image value as input, returns a score reflecting how well that image value matched expectations for a given location on the kernel. The results of all scoring functions within the functional template are then averaged and clipped to the continuous range $[0, 1]$.

By increasing or decreasing the interval over which affirming scores (i.e., > 0.5) are returned, scoring functions can encode varying degrees of uncertainty with regard to what image values are allowable. But in addition, knowledge of how a feature or object appears in sensor imagery can be encoded in scoring functions. And with various design strategies, the interfering effects of occlusion, distortion, noise, and clutter can be minimized. As a consequence, matched filters customized for specific applications using FTC are generally more robust than standard signal processing operations. The output of FTC is a map of values in the range $[0, 1]$, each of which reflects the degree of belief that the shape or object implicitly encoded in a functional template is present at that image location. In our ATR systems, FTC has been used primarily as a direct, one-step means of 3-D object detection and extraction. In MIGFA, FTC is used for edge detection, thin-line filtering, thin-line smoothing, shape analysis, and thinning of shapes.

Consider the simple matched filter shown in Fig. 1, which has been designed to detect gust fronts in reflectivity data. Gust fronts are observed as thin lines of moderately high reflectivity values (approximately 0 to 20 dBZ), with low reflectivity values (approximately -10 to 0 dBZ) ahead and behind the front. On the left is the template kernel consisting of integers, corresponding to the 2 scoring functions shown on the right. Elements of the kernel that do not have an index form guard regions in which image values are ignored and have no effect on match scores. Scoring function 1, corresponding to the flanking regions of low reflectivity, returns a maximal score of 1.0 for image values in the interval of -20 dBZ to -5 dBZ, a gradually decreasing score for image values in the interval -5 dBZ to 10 dBZ, and a score of -2.0 for image values larger than 10 dBZ. Scoring function 2, corresponding to the center of the kernel where moderately high reflectivity values are expected, returns maximal scores in the interval between 5 and 12.5 dBZ with gradually decreasing scores for both higher and lower image values. Note that while very low image values can generate scores of -1.0, a slower decline in score with a minimum score of 0.0 is evident for image values above the maximally scoring interval. The relatively greater tolerance to high reflectivity values reflects the belief that very high reflectivity gust fronts might exist and that gust fronts sometimes pass under high reflectivity storm cells.

INTEREST IMAGES

"Interest" is used as a medium for data fusion and for assimilating evidence at the pixel level (Ref. [9]). An interest image is a map of numeric values in the range $[0, 1]$, indicating the presence of some feature that is selectively indicative of an object being sought. Higher pixel values reflect greater belief that the intended feature is present at that location. Given the assumption that the output of any feature detector can be configured as an interest image, evidence from any number of registered sources of information can be easily combined using simple or arbitrarily complex rules of arithmetic or fuzzy logic. Clusters of high values in the combined interest image are then used

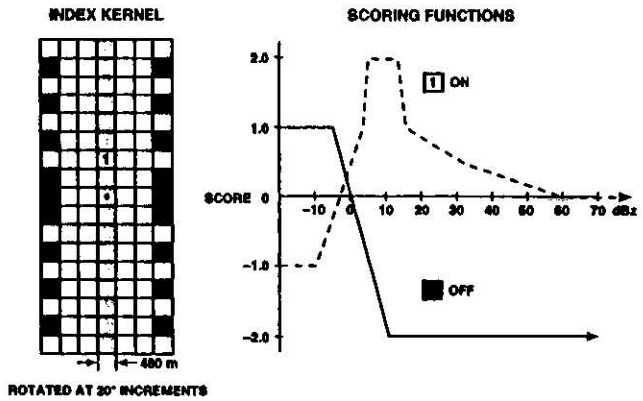

Figure 1: Example functional template for reflectivity thin line feature detection.

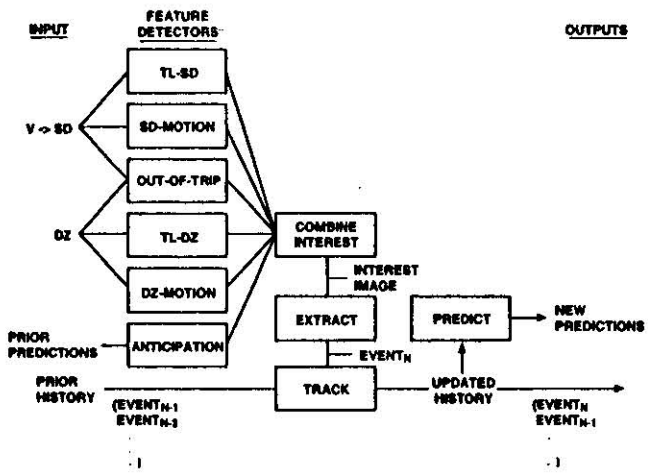

Figure 2: MIGFA block diagram.

to guide selective attention and serve as the input for object extraction. If done effectively, the combined interest image provides a better representation of object shape than is evident in any single sensory modality.

An individual feature detector may be reliable only under certain identifiable circumstances. By using knowledge of such circumstances and by allowing feature detectors to mutually support or compensate for each other, relatively good performance can be achieved using feature detectors that may individually be weakly or inconsistently discriminating.

ALGORITHM DESIGN

The system block diagram in Fig. 2 illustrates the configuration of the ASR-9 WSP version of MIGFA. Input images V (Doppler velocity image) and DZ (reflectivity image) are passed to five simple, independent feature detectors that use FTC. The resulting interest images are fused into a single combined interest image. From the combined interest image, a set of gust front points, collectively called an *event*, is extracted. This event is integrated with prior history in the TRACKING module, which establishes a basis for making predictions of subsequent behavior.

Input Images

Reflectivity and velocity data are converted from polar to Cartesian format, producing images DZ and V, respectively. As part of the conversion to Cartesian format, the data are subsampled so that the

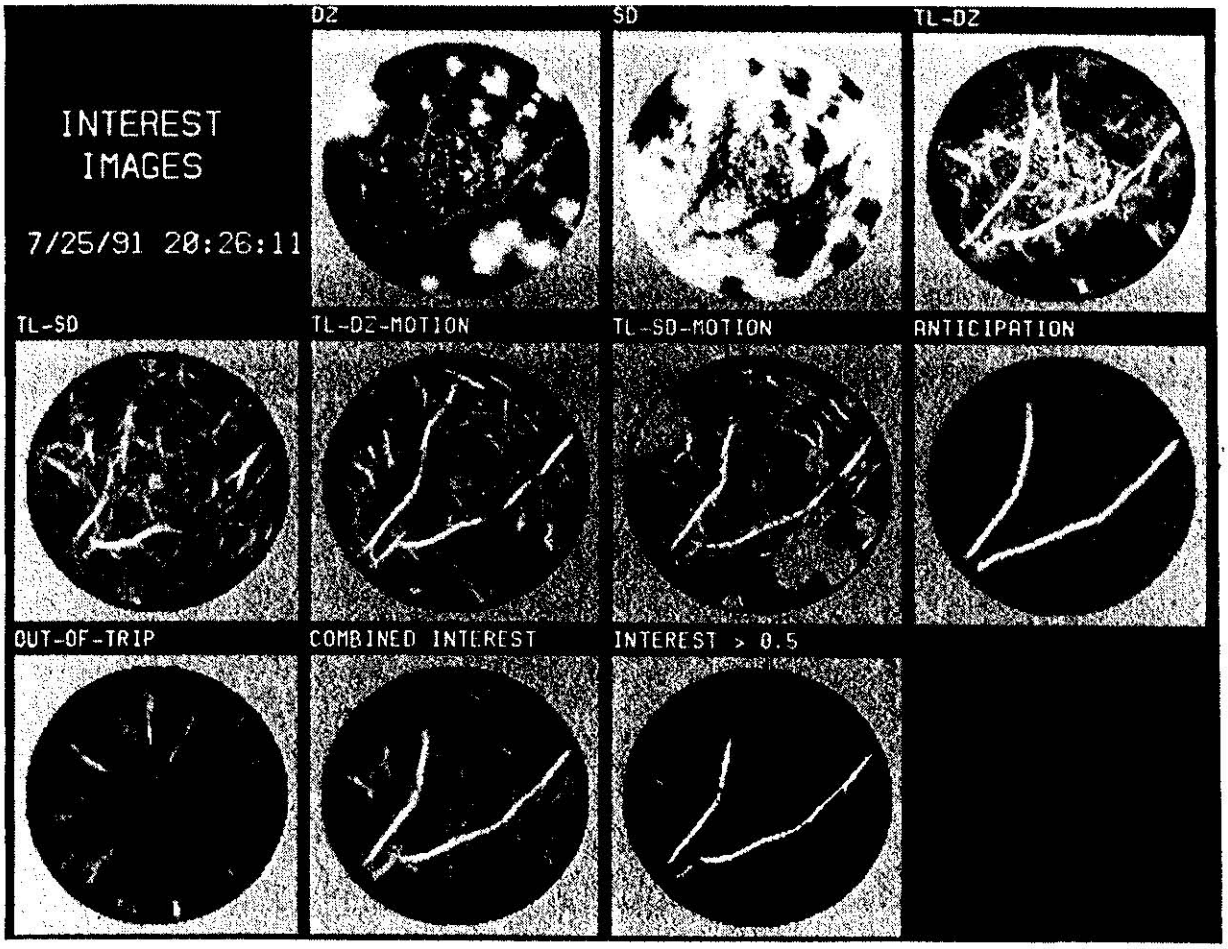

Figure 3: Combining multiple sources of interest.

pixel size is 480 meters per pixel.

Although convergence signatures are not reliably visible in ASR-9 data, the velocity data does contain relevant information. Velocity measurements within a gust front have higher signal-to-noise ratios, and consequently lower local Doppler variance, than the surrounding clear air measurements. Consequently, a map of velocity local standard deviation (SD) is computed from V. Gust fronts appear in SD as thin lines of low variance against a background of high variance.

Feature Detectors

The first feature detector referenced in Fig. 2, labelled TL-DZ, applies the thin line filter shown in Fig. 1 to the reflectivity image DZ. The output is an interest image shown in Fig. 3 highlighting the locations where thin lines were found in DZ. The second feature detector, DZ-MOTION, applies a similar thin line filter to the difference of two sequential DZ images. The thin line filter in DZ-MOTION uses the same kernel as TL-DZ, but scoring functions that reflect the effects of image differencing. Image differencing eliminates those thin lines that are not moving.

The feature detector OUT-OF-TRIP uses FTC to identify range-ambiguous echos of more distant weather. These range-ambiguous echos tend to appear as radially aligned thin lines and are a potential source of false alarms. Areas believed to reflect out-of-trip weather are given high interest values, which are then subtracted from corresponding areas of the other interest images.

The feature detectors TL-SD and SD-MOTION use thin line filters similar to those used in TL-DZ and DZ-MOTION. They differ only in that the scoring functions are suited for values of Doppler standard deviation instead of reflectivity.

Anticipation

The final feature detector, ANTICIPATION, provides a mechanism for spatially adjusting the detection sensitivity of MIGFA on the basis of knowledge of various environmental data, including current gust front trackings and dominant weather patterns.

The most important use of anticipation is as a replacement for *coasting*. Simply put, coasting is the continued tracking of a target on a radar screen for some time interval after the target signal falls below threshold. Coasting assumes that the loss of a target's signal is not due to a change of behavior (e.g., a change in velocity or perhaps the disappearance of the target). Gust fronts do change behavior, as in cases of collisions of gust fronts. Consequently, blindly coasting a signal after its loss is a potential source of false alarms. Anticipation provides an alternative by progressively increasing the sensitivity of the detection system, supporting weak evidence that would otherwise be below detection thresholds. Based on prior history of gust front behavior, a prediction is made of where the gust front is expected to be in the current scan. These predictions are used to create a band of elevated Interest values; not so high as to trigger a detection by themselves, but high enough to raise colocated weak signals above detection threshold.

Combining Interest

The four interest images TL-DZ, TL-SD, DZ-MOTION, and SD-MOTION are averaged together (missing values are ignored). The resulting averaged interest image and the ANTICIPATION interest image are combined as a weighted average; ANTICIPATION is given a weight of 0.25 while the average of the first four interest images is given a weight of 0.75. Finally, elements of the OUT-OF-TRIP interest

image are multiplied by 0.25 and subtracted from the elements of the weighted average. The resulting image is the combined INTEREST image. An example of how the outputs of multiple feature detectors are fused into a combined interest image is shown in Fig. 3.

Extraction

The goal of extraction in MIGFA is to identify the set of points (collectively called an *event*) that lie in *any* gust front. Certainly, some chains of points are spatially segregated or have different velocities. For purposes of reporting, such chains can be inferred to belong to separate gust fronts. However, there is no concerted attempt to label or track gust fronts as entities. Instead, individual points are tracked across time; the fact that a point belongs to one gust front or another is irrelevant to processing. Because different points can have variable velocities, MIGFA predictions are elastic.

The thin line shapes evident in the combined interest image are smoothed using a functional template whose shape increases the weights for the extremes of a thin line segment over that of the center. This operation serves to bridge gaps between thin line fragments. It also tends to suppress random interest values that are not aligned. A threshold of 0.5 (level of ambiguity) is applied to this smoothed interest image, setting all pixel values below threshold to 0. Note that this is the first discriminatory threshold applied in processing any one scan.

The resulting elongated shapes are thinned, resulting in chains of points. Chains less than a minimum length (6.25 km) are rejected. The remaining chains of points are extended from their end points along ridges of relatively high interest values until either the interest value of the next point falls below a minimum value (0.3) or the change in orientation from an initial end point to the next point exceeds some maximum (41 degrees).

At this point, the chains of points may form one or more disjoint complex networks, each potentially with internal closed loops and multiple end points. From this network of chain fragments is assembled the single most interesting (typically, but not necessarily, the longest) combined chain of points. Once the edges of this combined chain have been extracted, the process is repeated until either no more edges exist or the assembled combined chain has a summed interest value below a threshold (12.0). An example of extracted gust front points is shown in Fig. 4.

Tracking/Heuristics

Tracking is done by establishing point-by-point correspondence between successive scans. For each point in the current scan, a point in the previous scan is found that is nearby and that has a propagation velocity consistent with the point in the current scan. Once correspondence is established, a link is created from the point in the current scan to its corresponding point in the previous scan. Using this network of links, the complete history of any gust front point can be traced. After indexing is completed, each extracted chain of points is edited in order to smooth the computed propagation speeds and orientations over local segments of the chains.

Heuristics are then used to reduce the number of false alarms, making use of knowledge of how false alarms can be distinguished from real gust fronts. Examples of identifiable behaviors that are not associated with gust fronts are as follows:

1. The direction a point moves is inconsistent with the measured Doppler value.

2. A point that is approaching the radar site has a propagation velocity towards the radar site that is slower than the winds ahead of the front as measured by anemometer at the radar site.

The final stage of tracking is to make a binary decision whether or not to announce to the outside world that one or more gust fronts have been detected. Whether or not a particular chain is included in the report depends upon its summed interest score and the depths (number of events through which a point is traced) of constituent points. In chains with high summed interest scores, points with lower depths can be included.

	Gust Fronts		Gust Front Length	
	POD	PFA	PLD	PFD
Baseline (AGFA)	56.7	4.6	38.9	12.9
MIGFA	88.1	0.6	86.2	33.4

Table 1: AGFA and MIGFA performance on ASR-9 WSP data as scored against human interpretations.

	Gust Fronts		Gust Front Length	
	POD	PFA	PLD	PFD
Baseline (AGFA)	42.6	3.2	21.0	4.2
MIGFA	75.1	0.0	58.7	6.4

Table 2: AGFA and MIGFA performance on ASR-9 WSP data as scored against human interpretations of matching TDWR data.

Conversely, chains that have low summed interest scores are less likely to be gust fronts and are required to accumulate higher depths before being included in the announced gust front detections (see Fig. 4D and E).

Prediction

The current extracted event, indexed into the prior history, is used to make predictions of where the points having sufficient depth and interest are likely to be at some time in the future. Given the direction moved, the propagation speed, and the current coordinates of an extracted point, a new coordinate is computed for some time in the future. Fig. 4F shows the 10 and 20 minute predictions for two colliding gust fronts.

RESULTS

Table 1 compares performance of MIGFA against the previously constructed AGFA, which uses more conventional methods of signal processing and computer vision. A test set of ASR-9 WSP data collected in Orlando, Florida during field testing in 1991, contains 9 different gust front tracks through 15 hours (372 images). A human interpreter looking at the same data detected 280 instances of the 9 gust fronts tracked by the radar. Four figures of merit are shown for each of the two systems. The first one, the probability of detection (POD), is the number of detections made by each algorithm as a percent of human detected instances of gust fronts. The second number is the probability of false alarm (PFA). In addition to simply identifying fronts, the human interpreter delimited the length of each detected front. Detection quality was measuring by computing the degree of overlap between the gust front front as estimated by each algorithm and by the human interpreter. The third figure of merit indicates the length detected by each algorithm as a percent of the length delimited by the human interpreter (percent length detected or PLD). The fourth figure of merit indicates the amount of algorithm estimated gust front length that does not overlap with what the human interpreter could see (percent false length detected or PFD).

The first two columns indicate that MIGFA approximately doubled the number of fronts detected by AGFA, while decreasing the false alarm rate. Similarly, the PLD (column 3) reflects the improvement in detection rate. However, the increased PFD (from 12.9 % to 33.4 %) would suggest that MIGFA was doing a worse job of discriminating the extent of individual fronts. In order to better understand why MIGFA was extending fronts beyond what the human interpreter believed appropriate, we rescored AGFA and MIGFA results against human interpretations of TDWR data taken at the same time as the ASR-9 WSP data. In these rescored results, shown in Table 2, the POD and PLD remain relatively high for MIGFA. However, the PFD for MIGFA dropped from 33.4% using ASR-9 truth to 6.4% using TDWR truth. Given that the TDWR is a more sensitive radar, the difference between ASR-9 and TDWR truth probably reflects parts of actual gust fronts that were hard to see and consequently missed by the human interpreter of ASR-9 WSP data. These findings and an analysis of in-

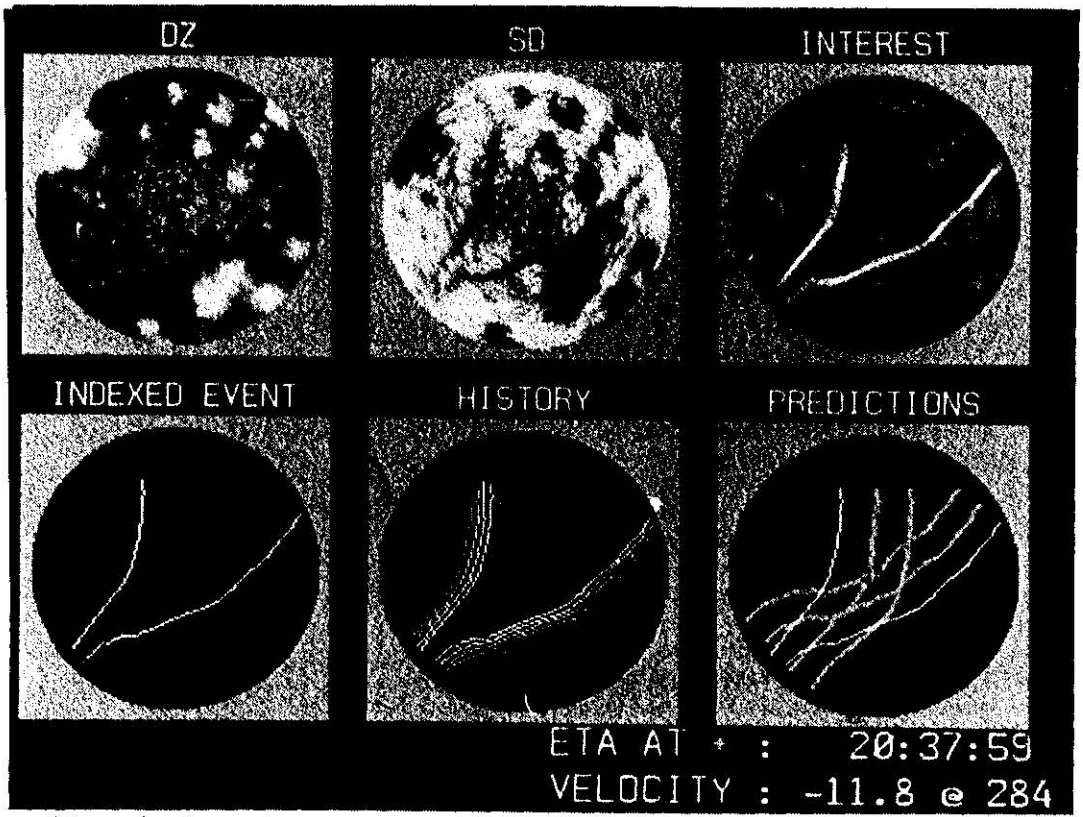

Figure 4: Summary of the results of processing a ASR-9 WSP scan containing two colliding gust fronts. (A) The reflectivity image DZ. (B) The image of local standard deviation SD. (C) The combined INTEREST image. (D) The set of extracted points called an EVENT; white points are those that have been tracked long enough to built up sufficient confidence for declaring them to be part of a gust front. (E) The gust front points shown in the context of prior HISTORY. (F) The 10 and 20 minute PREDICTIONS of gust front location.

	Gust Fronts		Gust Front Length	
	POD	PFA	PLD	PFD
MIGFA	75.4	1.8	80.8	21.1

Table 3: Results of MIGFA operational testing on ASR-9 WSP data collected in Orlando during August 1992.

dividual cases lead us to believe that the MIGFA-generated detections of gust fronts were more accurate than those generated by the human interpreter given the same ASR-9 WSP data.

This ASR-9 WSP version of MIGFA was deployed in Orlando for on-line operational testing during the summer of 1992. Table 3 shows the results for the period from 1 August to 20 September. In general, the results substantiate the off-line test results. Not surprisingly, the probability of detection (75%) and percent length detected (81%) were somewhat lower than those shown for the selected test set. Most of this difference can be explained by two problems. First, several gust fronts had reflectivity values at or below the sensitivity limits of the ASR-9. Of course, those fronts with reflectivity values below the ASR-9 limits were not detected by either MIGFA or the human interpreter. However, there were a few cases of marginal contrast in which the human could detect a gust front, but for which MIGFA never accumulated enough confidence to declare an alarm. Note, however, that the human interpreter had the opportunity to examine the sequence of radar images repeatedly and could use information from scans late in the sequence to confirm or deny the existence of the gust front in early scans. Not much can be done to overcome the sensitivity limits of the ASR-9. In most (but not all) cases, these gust fronts with marginal reflectivity levels were associated with weak wind shears. Since these weak fronts had minimal impact on airport operations, a failure to detect them should not be a significant liability.

The second problem was one of obscuration. In several cases, storm cells or out-of-trip weather were extensive enough to hide or fragment the thin line signatures so that some gust fronts were detected late, dropped early, or sometimes missed altogether. To improve performance in these cases, the ANTICIPATION feature detector has been modified. The interest values of anticipated gust front locations are now increased wherever potential obscuration is detected. By raising interest values high enough, detections are made even without supporting evidence from other sources (i.e., coasted).

Finally, there would seem to be an apparent problem in the relatively high PFD score (21.1%). However, as was shown in the offline test results described earlier, the PFD score is an overestimation. The high PFD score is almost entirely because MIGFA extended the ends of gust fronts beyond the signatures visible to the human interpreter. However, a case-by-case analysis using matching TDWR data indicates that most of the extensions contributing to the PFD are in fact reasonable extensions of real gust fronts. A more accurate estimation of the PFD will require the systematic rescoring of the results against TDWR truth.

CONCLUSIONS

Thin-line signatures in reflectivity images and convergence signatures in Doppler images are conceptually easy to define and to use in automated detection algorithms. And yet, several research groups have worked collectively for nearly 10 years to develop reliable automatic gust front algorithms. Despite the effort, no algorithm has demonstrated performance comparable to the ideal of human performance.

The problem is that automatic gust front detection is deceptively a much harder problem than simply detecting one or both of these signatures. In order for human interpreters to detect and track gust fronts, they make use of knowledge about the radar and the weather. They use spatial and temporal context. And, they have the ability to deal with uncertainty while assimilating ambiguous or even con-

tradictory evidence. The large performance gap between algorithmic and human gust front detection probably reflects the lack of these perceptual skills in previous algorithmic approaches. With the machine intelligence techniques developed as part of XTRS, MIGFA displays levels of performance that are competitive with human interpreters.

Given that reasonably good performance has been achieved despite the limitations posed by the ASR-9 characteristics, MIGFA performance ought to be very good on the higher resolution TDWR. A TDWR version of MIGFA has in fact already been assembled. It uses a different set of feature detectors in order to exploit the higher sensitivity reflectivity data and reliable Doppler convergence signatures. The control structure and the extraction, tracking and prediction modules are identical. Although a direct case-by-case comparison of AGFA and MIGFA has not yet been run for TDWR data, preliminary results indicate that MIGFA performance is again better than that for AGFA.

Finally, MIGFA is based on XTRS, a general system for automatic object detection and recognition that should be applicable to other meteorological detection problems. In conventional, hierarchical computer vision design, machine intelligence techniques are applied only in the higher (more abstract) levels manipulating symbolic representations of the data. In contrast, XTRS provides a framework for applying machine intelligence at the earliest stages of detection processing, manipulating raw pixel data. It is true that this approach will not solve the general vision problem of understanding a complex scene, i.e., understanding the relationships of several instances of hundreds of possible objects in an unconstrained contextual environment. But for problems in which the goal is to detect one particular type of object and discriminate it from an understood background, XTRS provides an effective, straightforward means.

REFERENCES

[1] H. Uyeda and D. Zrnic. Automated detection of gust fronts. *J. Atmos. Oceanic Tech.*, 3:36, 1986.

[2] A. Witt and S. Smith. Development and testing of the gust front algorithm. Technical Report DOT/FAA/PS-87/4, FAA, 1987.

[3] S. D. Smith, A. Witt, M. Eilts, D. Klingle-Wilson, S. Olson, and J. P. Sanford. Gust front detection algorithm for the Terminal Doppler Weather Radar part I: Current status. In *Proceedings of the 5th International Conference on the Aviation Weather System*, pages 31-34, Anaheim, CA, January 1989.

[4] G. Hermes, A. Witt, S. Smith, D. Klingle-Wilson, D. Morris, G. Stumpf, and M. Eilts. The gust front detection and wind shift algorithms for the Terminal Doppler Weather Radar system. *Journal Atmos. Oceanic Tech.*, (in press), 1992.

[5] M. Eilts, S. Olson, G. Stumpf, L. Hermes, A. Abrevaya, J.Culbert, K. Thomas, K. Hondl, and D. Klingle-Wilson. An improved gust front detection algorithm for the TDWR. In *Proceedings of the 4th International Conference on the Aviation Weather System*, pages J37-J42, Paris, France, June 1991.

[6] M.W. Merritt, D. Klingle-Wilson, and S.D. Campbell. Wind shear detection with pencil-beam radars. *Lincoln Laboratory Journal*, 2(3):483-510, 1989.

[7] M. Weber and T. Noyes. Wind shear detection with airport surveillance radars. *Lincoln Laboratory Journal*, 2(3):511-525, Fall 1989.

[8] Jacques G. Verly, Richard L. Delanoy, and Dan E. Dudgeon. Machine intelligence technology for automatic target recognition. *Lincoln Laboratory Journal*, 2(2):277-311, Summer 1989.

[9] Richard L. Delanoy, Jacques G. Verly, and Dan E. Dudgeon. Pixel-level fusion using interest images. In *Proceedings of the 4th National Symposium on Sensor Fusion*, Orlando, Florida, April 1991.

[10] Richard L. Delanoy, Jacques G. Verly, and Dan E. Dudgeon. Functional templates and their application to 3-D object recognition. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, San Francisco, California, March 1992.

[11] Richard L. Delanoy, Jacques G. Verly, and Bryan Williams. Region-based target recognition from laser radar imagery using appearance models. Technical Report TR-933, MIT Lincoln Laboratory, Lexington, MA, July 1992.

BIOGRAPHIES

Machine Intelligence Technology and Weather Sensing Groups
Lincoln Laboratory, Massachusetts Institute of Technology
244 Wood Street, Lexington, MA 02173-9108
Telephone: (617) 981-2545, Email: rld@ll.mit.edu

Richard Delanoy is a staff member of the Machine Intelligence Technology Group at Lincoln Laboratory. His degrees include a Ph.D. in neuroscience from the University of Florida in 1979 and a subsequent M.S. in computer science from the University of Virginia in 1987. After working as a software engineer with General Electric's Factory Automation Division, he joined Lincoln Laboratory in 1987, where he develops automatic object detection and recognition systems.

Seth W. Troxel received a B.S. in meteorology from San Jose State University, CA in 1983 and went on to work as a meteorologist in the Atmospheric Lidar Group of the NOAA Wave Propagation Laboratory, Boulder CO. In 1987 he joined Lincoln Laboratory as a meteorologist and software engineer with the Weather Sensing Group.