10.4

MACHINE INTELLIGENT GUST FRONT DETECTION FOR THE INTEGRATED TERMINAL WEATHER SYSTEM (ITWS)*

Seth W. Troxel and Richard L. Delanoy†

MIT Lincoln Laboratory
Lexington, Massachusetts

1. INTRODUCTION

The Integrated Terminal Weather System (ITWS), currently in development by the FAA, will produce a fully–automated, integrated terminal weather information system to improve the safety, efficiency and capacity of terminal area aviation operations. The ITWS will acquire data from FAA and National Weather Service sensors as well as from aircraft in flight in the terminal area. The ITWS will provide products to Air Traffic personnel that are immediately usable without further meteorological interpretation. These products include current terminal area weather and short–term (0–30 minute) predictions of significant weather phenomena. The Terminal Doppler Weather Radar (TDWR) will serve as a principle sensor providing data to a number of the ITWS algorithms.

One component of the ITWS will be an algorithm for detecting gust fronts and wind shifts. A gust front is the leading edge of a cold air outflow from a thunderstorm. The outflow, which is deflected at the ground, may propagate many miles ahead of the generating thunderstorm, and may persist as an outflow boundary long after the original storm has dissipated. 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 in active runway configuration and rerouting of aircraft within in the terminal airspace. In addition, wind shear, turbulence, and cross–winds along the frontal boundary pose significant safety hazards to departing and landing aircraft. Reliable detection and forecasting of gust fronts and wind shifts will both improve air safety and reduce costly delays.

Lincoln Laboratory has developed an Initial Operational Capability (IOC) Machine Intelligent Gust Front Algorithm (MIGFA) for the ITWS which currently utilizes TDWR and LLWAS or ASOS anemometer data and makes use of new techniques of knowledge–based signal processing originally developed in the context of automatic target recognition [Verly, 1989]. Extensions to the IOC to incorporate additional sensor or product data available under the ITWS (e.g., NEXRAD, terminal winds) are currently under development. MIGFA was first developed for the Airport Surveillance Radar with Wind Shear Processor (ASR–9 WSP). Its design and performance have been documented in previous reports by the authors [Delanoy 1993a]. This paper focuses on the design of the more recently developed TDWR MIGFA and its extension and adaptation to the ITWS (a more detailed description of the TDWR MIGFA can be found in Troxel [1994]). An overview of the signal processing techniques used for detection and tracking is presented, as well as a brief discussion of the wind analysis methods used to arrive at the wind shift and wind shear estimates. Quantitative performance analyses using data collected during recent field testing in Orlando, FL and Memphis, TN are presented. Test results show that MIGFA substantially outperforms the gust front detection algorithm used in current TDWR systems [Hermes, 1993] (MIGFA is currently under consideration as an upgrade option for TDWR).

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

†*Corresponding author address:* Seth Troxel, 244 Wood Street, MIT Lincoln Laboratory, Lexington, MA 02173–9108

2. GUST FRONT SIGNATURES

Gust fronts produce signatures that are observable to varying degrees in weather reflectivity and Doppler velocity data generated by Doppler weather radars. One such signature is the reflectivity thin line echo. This low reflectivity echo is thought to be produced by a concentration of scatterers (dust, insects, rain droplets) along the leading edge of the outflow, or by changes in refractive index across the front. Because reflectivities within the thin line are low (–5 to 20 dBZ), the thin line echo may vanish in regions where return signal is below system noise, thereby making recognition and detection more difficult. In addition, gust fronts embedded within precipitating storm regions are unlikely to produce discernable thin lines. Hence, reflectivity thin lines cannot be relied upon as the sole basis for detection.

A second signature is the velocity convergence signature, found along the boundary between the leading edge of the outflow and the ambient winds ahead of the gust front. It is identified by a sharp decrease in Doppler velocity when examined along a single radar radial. Like thin line signatures, convergence signatures are not 100% reliable. Convergence signatures often diminish in extent or vanish altogether when gust fronts become radially aligned (no radial wind component). Furthermore, clear air regions ahead of or behind some gust fronts are sometimes too "clean" to produce sufficient signal–to–noise ratios for accurate Doppler measurements on both sides of the gust front, resulting in fragmented convergence signatures.

Another useful gust front signature is motion. When sequential radar scans are compared, convergence and thin line signatures of a gust front will move conspicuously in a direction perpendicular to the orientation of the convergence boundary and reflectivity thin line, while other features in the background scene (storm cells, ground clutter) are relatively stationary over the short time scales being considered.

3. LOW–LEVEL MACHINE INTELLIGENCE

In MIGFA, sensor–, object–, and context–dependent knowledge is applied in the earliest (image processing) levels of processing through the use of two techniques: functional template correlation, and data level fusion.

Functional template correlation (FTC) [Delanoy, 1992] is a generalized matched filter that permits the construction of customized, knowledge–based image processing operations. The output of FTC is a map of match scores, each reflecting the degree of belief that the shape or object implicitly encoded in a functional template is present at that image location. A detailed description of FTC, along with examples of templates constructed for gust front detection can be found in Delanoy [1993a] and Troxel [1994].

Knowledge of the varying reliability of the selected feature detectors is used to guide data fusion. Such conditional data fusion is simplified by using "interest" as a common denominator (Delanoy 1993b). An interest image is a map of numeric values in the range [0,1], in which higher pixel values reflect greater confidence that an intended feature is present at that location. Clusters of high interest values are then used to guide selective attention and serve as the input for object extraction. Because interest values are dimensionless, evidence from any number of registered sources of information can be easily combined using simple or arbitrarily complex rules of arithmetic or fuzzy logic.

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. If done effectively, the combined interest image provides a better representation of object shape than is evident in any single sensory modality.

4. ALGORITHM DESIGN

4.1. Overview

The system block diagram in Figure 1 illustrates the current configuration of the IOC version of the ITWS MIGFA. The design is nearly identical to the ASR version of MIGFA, except for the number and type of feature detectors employed. In preparation for processing, input images DZ (reflectivity) and V (Doppler velocity) from the received TDWR scan are converted from polar (400 range bins X 360 radials) to Cartesian (260 X 260) representations with 480 m pixel resolution. A map of radial shear (radial velocity change over a 1 km distance) is derived from input image V and serves as a third input image called DV. The input images are then passed to multiple simple independent feature detectors that attempt to localize those features which are selectively indicative of gust fronts. The outputs of each of these feature detectors, most of which are based on some application of FTC, are expressed as interest images that specify evidence indicating where and with what confidence a gust front may be present. The different interest images are fused to form a combined interest image, providing an overall map of evidence indicating the locations of possible gust fronts.

From the combined interest image, fronts are extracted as chains of points. The chains extracted from a radar scan, collectively called an event, are integrated with prior events by establishing a point–to–point correspondence. Heuristics are then applied to reject those chain points which have an apparent motion that is improbable and to edit properties of individual points that are inconsistent with their neighbors. The updated history is used to make flexible, point–by–point predictions of where points along the front will be located as some future time. Such predictions are also used in the processing of subsequent images,

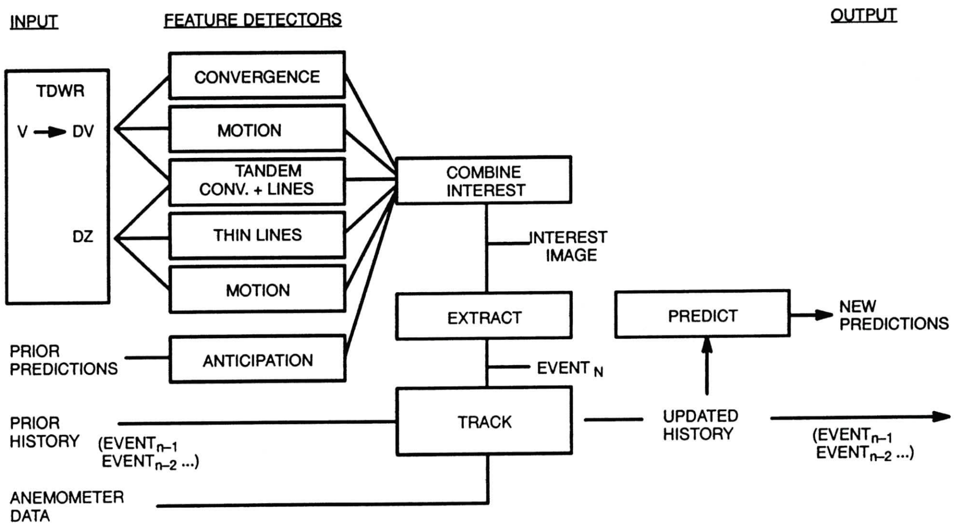

Figure 1. Block diagram of the Machine Intelligent Gust Front Algorithm (MIGFA) for the ITWS

specifically in the feature detector called ANTICIPATION (described later).

For each extracted gust front chain, Doppler velocities in the vicinity of the segment of the front nearest the airport are analyzed in conjunction with LLWAS or ASOS wind sensor data to produce estimates of the wind speed and direction behind the front (i.e., the wind shift forecast) and the velocity change (ΔV) across the front.

4.2. Feature Detectors

TDWR input data is processed by the MIGFA using approximately ten different FTC–based feature detectors (a similar set of feature detectors has been designed for use with NEXRAD data). Figure 2 shows several interest images produced by applications of FTC to the input images that are displayed in gray–scale along the top row of the figure (in the velocity image V, white pixels indicate strong winds directed away from the radar, while black pixels indicate strong winds directed toward the radar). The figure also shows the combined interest image resulting from combination of the individual interest images. Homogenous, mid–level gray regions denote mask areas where specific feature detectors are prevented from expressing an opinion regarding the presence of a gust front, deferring instead to the evidence generated by other feature detectors. In each of the interest images, white pixels indicate locations of maximum confirming interest, while black pixels indicate locations of maximum disconfirming or negative interest.

The first interest image in the second row (TDWR–TL–DZ–CONV) is generated by a tandem feature detector that looks for thin lines in the DZ image that are coincident and aligned with velocity convergence in the DV image. Since obscuration prevents detection of thin line echoes inside storm cells, the TDWR–TL–DZ–CONV detector is prevented from generating opinions in these areas.

The TDWR–DZ–CONV–MOTION detector is similar to TDWR–TL–DZ–CONV except that it looks for tandem motion of reflectivity thin lines and velocity convergence lines. Motion detectors are based on simple differencing. The DZ image from the previous scan (received approximately 5 minutes earlier) is subtracted from the DZ image from the current scan. In the dif-

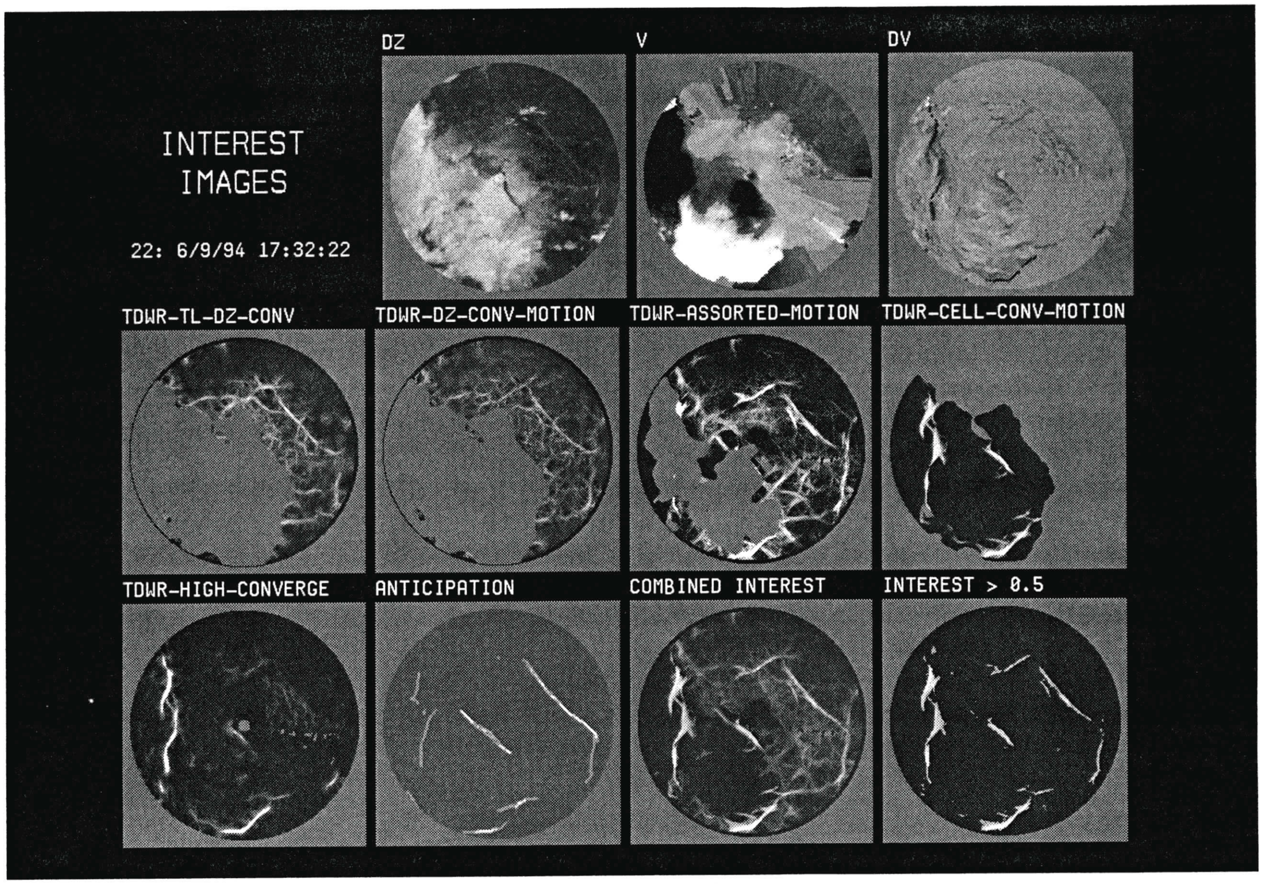

Figure 2. Combining of interest images. Input DZ (reflectivity), V (velocity), and DV (velocity change) images are shown along with output interest images from some of the feature detectors. The last two frames in the lower right show the images resulting from combining the various interest images.

ferenced DZ image, gust fronts appear as white lines (positive values at the front's position in the current scan) that are trailed by parallel dark lines (negative values at the front's position in the previous scan). The functional template returns maximal scores where thin lines of positive values occur.

The interest image labelled TDWR–ASSORTED–MOTION represents the combined (maximum) evidence from a number of single feature detectors. This constitutes a comparatively liberal detector that helps to offset the relatively conservative opinions produced by the tandem detectors. TDWR–CELL–CONVERGE and TDWR–CELL–CONV–MOTION look for static and moving velocity convergence boundaries respectively within storm regions. Note that all non–storm regions have been masked so that these detectors offer no opinions outside of storm regions.

Finally, the ANTICIPATION feature detector provides a mechanism for spatially adjusting the detection sensitivity of MIGFA on the basis of knowledge of various environmental data including the prior history of the gust fronts being tracked and dominant weather patterns. In particular, anticipation is used as a replacement for *coasting*, which is the continued tracking of an object (e.g., a gust front) for some time interval after the object's signal falls below threshold. Coasting works only when the target being tracked maintains a consistent velocity. But in reality, the reason the object's signal falls below threshold is often because the object's behavior changed. In contrast, anticipation works by creating bands of high interest values where the object is expected to be in the current scan. Anticipation is set not so high as to trigger a detection by itself (i.e., coasting), but high enough to raise collocated weak signals above detection threshold.

5. WIND SHIFT AND WIND SHEAR ESTIMATION

In order to generate wind shift forecasts and wind shear hazard reports, wind field analyses need to be conducted on multiple spatial scales and at different locations; for runway planning purposes, wind shift forecasts are most valuable when they indicate what the persistent winds are at some distance behind the front, while meaningful wind shear hazard estimates require an analysis that captures the gradient in the immediate vicinity of the front. We found that using a single estimation technique worked well for many cases, but produced less than satisfactory results with different combinations of viewing geometry and data quality problems.

Instead, MIGFA currently utilizes a multi–algorithmic consensus approach to gust front wind estimation that fits nicely into the ITWS framework. Wind analyses are performed in several regions around each gust front, and the results are used in construction of the wind shift forecasts and wind shear hazard reports. The analyses results rely on a consensus of estimation techniques including:

- prior history,
- linear least squares fit of Doppler values in a region 10–minutes behind the part of the front closest to the airport,
- mean of viewing angle–adjusted Doppler values within the front, and
- LLWAS or ASOS anemometer measurements.

Using a simple rule set, each estimate is assigned a weight according to the reliability of that estimate in the given situational context, and the weighted average is computed. For each gust front, consensus estimates of winds behind, winds ahead, and winds within the gust front are used to derive the final wind shift forecast and wind shear hazard value assigned to the front.

6. RESULTS

During the summer of 1993, MIGFA was operated in the Lincoln Laboratory TDWR testbed during real–time testing in Orlando, Florida. MIGFA ran for 7 hours a day (minimum), 7 days a week, during most of the summer. During much of the same period, the current production TDWR algorithm (GF88) was also run, providing an opportunity for comparing the performance of the new algorithm against an existing standard. Performance of both algorithms was assessed by automated scoring of algorithm detections against a truth database generated by visual inspection of each input image processed by the two algorithms.

To generate the truth database, an expert analyst had access to Doppler velocity and reflectivity images for an entire sequence of TDWR scans, which could be viewed separately or in sequence as a movie. For each scan, the analyst entered a list of coordinates marking the gust front end points along with an intermediate sampling of points in between. For categorization of results, the estimated maximum wind shear in the convergence zone was also stored (Table 1).

Table 1. Gust front strength categories

Category	Radial Velocity Difference ΔV (m/s)
Weak	$5 \le \Delta V < 10$
Moderate	$10 \le \Delta V < 15$
Strong	$15 \le \Delta V < 25$
Severe	$\Delta V \ge 25$

An automated scoring procedure, described in detail by Klingle–Wilson [1992] compares computed gust front detections against coordinates contained in the truth database. Briefly described, the scoring algorithm connects the sequence of coordinates defining the limits of the gust front and expands the collection of connected line segments into a 5–km wide region that is called a truth box. The scoring algorithm measures detection performance by two metrics. The first measure is a crude event–based statistic that counts a detection as valid if any part of the detection overlaps any part of a truth box. A detection is counted as false if it falls completely outside of any truth boxes.

An overall probability of detection (POD) is computed by dividing the number of successfully detected fronts by the number of fronts identified by the human analyst. The probability of false alarm (PFA) is the number of false detections divided by the total number of detections (both valid and false). Detection quality was assessed by comparing the length of the front as estimated by each algorithm against the length identified by the human analyst. The percent length detected (PLD) is expressed as a ratio of the length detected to the length delimited by the human analyst. The percent of false length detected (PFD) reflects the fraction of total detection length that was not verified by truth.

Table 2 compares automated scoring results categorized by strength for MIGFA and GF88 for the Orlando 1993 test set. Results were computed from a substantial database comprised of 230 hours of data collected on 30 different days during the test period. As can be seen from the table, MIGFA significantly outperformed the current TDWR gust front algorithm. MIGFA correctly detected and tracked over 70% of all gust fronts identified by human analysts, compared to approximately 30% for the existing algorithm. MIGFA did an even better job of detecting the overall length of gust fronts. MIGFA detected 66% of the total length of all gust fronts, representing a four–fold improvement over GF88 performance in this category.

Note also that with both metrics, MIGFA decreased the false alarm rate. Of the nearly 2750 scans processed, false alarms occurred on only 54 scans with a resulting false alarm probability of only 0.03. The percent of false length detected by MIGFA (7%) is significantly less than that for GF88 (17%). The majority of MIGFA's false detection length came directly from sporadic false detections (predominantly from leading edges of storm regions and thin, weak rain echoes) that were not associated with real gust front events. A small fraction (5–10%) of the PFD was due to inappropriate extension of some fronts beyond limits identified by the analyst. An equally small fraction of the PFD represents situations where the analyst was uncertain about the presence of a gust front, but decided not to generate truth for it; MIGFA may have been inappropriately penalized in some of these cases.

A significant fraction of the false alarms (and false detection length) issued by the GF88 algorithm were the result of coasting prior detections to maintain tracking during overhead passage (when velocity convergence signatures often vanish). As discussed earlier, this is problematic since gust fronts do not always maintain their characteristics over the coasting period. They may speed up, slow down, or dissipate altogether. Since MIGFA does not use coasting, it does not suffer from this problem.

MIGFA's substantial improvement in detection performance over the existing detection algorithm can be understood by considering that the current TDWR algorithm utilizes traditional 1–D processing of radial Doppler velocity convergence signatures [Uyeda, 1986] as its primary means of detection. Without spatial context and additional information from other sources of evidence such as thin lines and motion, the GF88 algorithm is at a distinct disadvantage. Although some recent experimental algorithms have attempted to make additional use of thin line recognition [Eilts, 1991], modest improvements in detection probability have been accompanied by undesirably high false alarm probabilities. Once again, the traditional processing methods limit achievable performance gains.

After examining results from the 1993 operations, several improvements to MIGFA were made to further reduce false alarms, improve detection of gust fronts embedded within storm cells, and improve the accuracy of wind shift and wind shear hazard estimates. These improvements were included in the IOC ITWS MIGFA that operated during the 1994 ITWS Demonstration and Validation tests in Memphis, TN. Preliminary results from a limited data set consisting of 22 gust fronts that impacted the Memphis airport during June and July have been compiled. MIGFA detected 16 of the 17 gust fronts that were operationally significant (i.e., $\Delta V > 5$ m/s). The missed detection was a case where initial formation of a weak (5.5 m/s) gust front occurred on the airport with no advance indication. Only one on–airport detection issued by MIGFA was determined to be false (due in part to data contamination from second trip weather) – the associated ΔV of 4.4 m/s was well below the alert threshold of 15 knots (7.72 m/s) and consequently did not adversely impact airport operations.

Wind shift forecasts and wind shear hazard reports were compared against LLWAS centerfield anemometer data for the same 23 gust fronts. As can be seen in Table 3, MIGFA did a good job of characterizing the wind shear hazard and forecasting the wind shift speed behind each front. Wind direction forecasts, with a mean difference of more than 35 degrees, appear to be less accurate. However, it should be noted that the limited data set contains a number of complex events including gravity waves, rapidly evolving gust fronts, and dissipating fronts. Furthermore, when wind direction values obtained by inspection of TDWR Doppler images are compared against corresponding LLWAS centerfield measurements, a mean difference of 25 degrees and a standard deviation of 31 degrees results. This suggests that much of the discrepancies between MIGFA and LLWAS wind direction estimates can be attributed to differences between wind observations obtained by the two sensors.

7. SUMMARY AND FUTURE WORK

Automated gust front detection and tracking, together with wind shear warnings and forecasts of wind shifts, serve important traffic safety and route planning needs for ATC. In this paper, we have described the Machine Intelligent Gust Front Algorithm (MIGFA). MIGFA uses maps of evidence called interest images, which can be fused at the pixel level, to produce a combined interest image that represents a single pixel–map of evidence of where gust fronts are likely to be. Individual interest images are constructed using functional template correlation (FTC), a generalized matched filter that facilitates knowledge–based image processing operations. From these combined maps of evidence, MIGFA is able to identify and track the most probable locations of gust fronts.

Table 2. Performance comparison between the current TDWR gust front algorithm (GF88) and MIGFA for Orlando gust fronts.

ORLANDO, 1993	TDWR GF88				MIGFA			
	EVENT		LENGTH		EVENT		LENGTH	
STRENGTH	POD	PFA	PLD	PFD	POD	PFA	PLD	PFD
WEAK	0.28	–	0.10	–	0.70	–	0.57	–
MODERATE	0.65	–	0.31	–	0.84	–	0.66	–
ALL	0.31	0.08	0.11	0.17	0.71	0.03	0.57	0.07

Table 3. MIGFA report error for wind shear (ΔV) and wind shift (speed and direction) estimates when compared against LLWAS centerfield wind measurements (Memphis, 1994 data).

	MEAN ERROR	STD. DEV.
ΔV	1.5 m/s	1.2 m/s
SPEED	2.2 m/s	2.7 m/s
DIRECTION	36.0 deg	27.2 deg

Feature detectors used in MIGFA take advantage of three generally recognized radar gust front signatures: reflectivity thin lines, radial convergence, and motion. MIGFA's use of data fusion to combine evidence from the various feature detectors provides a more robust means of recognizing gust fronts than has been possible with the more traditional 1–D radial–based signal processing techniques utilized by the current TDWR gust front detection algorithm.

Results from real–time operational testing of MIGFA at Orlando, FL in 1993 were presented. They show that MIGFA provided substantial performance gains when compared against the existing TDWR gust front detection algorithm. The TDWR–based version of MIGFA has also been shown to be suitable as the basis for an ITWS MIGFA. Preliminary results from recent tests conducted during the 1994 ITWS DemVal are encouraging. The IOC ITWS MIGFA detected 16 of the 17 operationally significant gust fronts that impacted the Memphis airport during June and July.

A number of near–term improvements are under development that will serve to transition the IOC ITWS MIGFA to a more full–fledged ITWS MIGFA and further improve MIGFA's performance. One potential improvement is the use of 2–km gridded wind analyses available from the ITWS terminal winds algorithm. These wind data would serve as an additional source of wind information to help refine MIGFA's wind shift and wind shear reports.

Another area of development is geared toward expanding coverage by assimilating data from multiple, overlapping TDWR or NEXRAD data sources (multiple TDWR's will exist at 8 major cities). Of course, in order to construct a version of MIGFA that can combine interest from multiple radars, one precondition is the development of feature detectors tuned for each radar source (the rest of MIGFA is identical for any Doppler weather radar). To that end, a set of NEXRAD feature detectors have been developed and are undergoing further refinement and testing.

8. REFERENCES

Delanoy, R. L., J. G. Verly, and D. E. Dudgeon, 1992: "Functional templates and their application to 3–D object recognition," *Proceedings of the 1992 International Conference of Acoustics, Speech, and Signal Processing*, San Francisco, CA, March 23–26, 1992.

Delanoy, R. L., and S. W. Troxel, 1993a: "Machine intelligent gust front detection," *Lincoln Laboratory Journal*, Vol. 6, No. 1, 187–211, 1993.

Delanoy, R. L., J. G. Verly, and D. E. Dudgeon, 1993b: "Pixel–level fusion using interest images," *Technical Report TR–979*, MIT Lincoln Laboratory, Lexington, MA, May 1993.

Eilts, M., S. Olson, G. Stumpf, L. Hermes, A. Abrevaya, J. Culbert, K. Thomas, K. Hondl, and D. Klingle–Wilson, 1991: "An improved gust front detection algorithm for the TDWR," *Proc. 4th Intl. Conf. on the Aviation Weather System*, Paris, June 1991, p. J37.

Hermes, L., A. Witt, S. Smith, D. Klingle–Wilson, D. Morris, G. Stumpf, and M. Eilts, 1993: "The gust front detection and wind shift algorithms for the Terminal Doppler Weather Radar system," *J. Atmos. Oceanic Tech.*, Vol. 10, No. 5, 693–709, October, 1993.

Klingle–Wilson, D. L., M. F. Donovan, S. H. Olson, and F. W. Wilson, 1992: "A comparison of the performance of two gust front detection algorithms using a length–based scoring technique," *Project Report ATC–185*, MIT Lincoln Laboratory, Lexington, MA, May 1992.

Troxel, S., and R. L. Delanoy, 1994: "Machine intelligent approach to automated gust front detection for Doppler weather radars", *SPIE Proceedings – Sensing, Imaging, and Vision for Control and Guidance of Aerospace Vehicles*, Vol. 2220, 182–192, Orlando, FL, April 4–5, 1994.

Uyeda, H., and D. S. Zrnic', 1986: "Automated detection of gust fronts," J. Atmos. Oceanic Tech. **3**, 36, 1986.

Verly, J. G., R. L. Delanoy, and D. E. Dudgeon, 1989: "Machine intelligence technology for automatic target recognition," *Lincoln Laboratory Journal*, Vol. 2, No. 2, 277–311, 1989.