A MACHINE INTELLIGENT GUST FRONT ALGORITHM FOR DOPPLER WEATHER RADARS

Richard L. Delanoy and Seth W. Troxel

Lincoln Laboratory, Massachusetts Institute of Technology
Lexington, MA 02173-9108 USA

1. INTRODUCTION

Gust fronts generated by thunderstorms can seriously affect the safety and efficiency of airport operations. Lincoln Laboratory, under contract with the Federal Aviation Administration (FAA), has had a significant role in the development of two Doppler radar systems that are capable of detecting low-altitude wind shears, including gust fronts, in the airport terminal control area. These systems are the the latest generation Airport Surveillance Radar, enhanced with a Wind Shear Processor (ASR-9 WSP) and the Terminal Doppler Weather Radar (TDWR).

Gust fronts produce signatures that are observable to varying degrees in reflectivity and Doppler velocity data generated by these radars. In Doppler velocity images, gust fronts are recognizable as zones of converging winds. In reflectivity images, gust fronts appear as thin lines of increased intensity, which occur as the result of rain, dust, insects, or debris being lofted and concentrated at the leading edge of the front. An existing automated gust front detection and forecasting algorithm, developed principally for TDWR data and called in this paper the Advanced Gust Front Algorithm (AGFA), has achieved respectable levels of performance using these data (Eilts, 1991 and Merritt, 1989). With clear, unambiguous radar signatures AGFA performs reasonably well. The challenge is in constructing an algorithm that can detect marginally detectable, ambiguous cases without incurring unacceptable false alarm rates.

Several sources of ambiguity exist. For example, gust front thin-line signatures can be obscured by large areas of precipitation. Velocity convergence signatures can vanish when gust front orientations result in bad Doppler viewing angles. Gust front signatures can also be mimicked by other natural phenomena, such as flocks of birds, clouds of dust stirred up at construction sites, elongated low-intensity rain echoes, and ground clutter. Finally, gust fronts can have very low radar cross-section densities, sometimes below the sensitivity of the radar system, making detection difficult.

The ASR-9 WSP provides a less expensive alternative to the TDWR as a terminal weather radar (Weber, 1989). 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 primary reason for the lack of performance is 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. Consequently, few wind convergence signatures are visible, forcing AGFA to rely on only the thin-line portion of its algorithm. Unfortunately, the reduced sensitivity also makes faint thin-line signatures more fragmented and harder to resolve from background.

The radars themselves are sufficient for the task of gust front detection, since experienced human observers can detect and track gust fronts in images generated by these radars. And yet, sufficiently high detection rates with few false alarms has been an elusive goal for developers of automated gust front detection algorithms. Skilled human interpreters rely upon spatial and temporal contextual information and assimilate weak, uncertain, ambiguous, and even contradictory evidence. Humans are also adept at conditionally fusing information from various sources, reflecting knowledge that different signatures can have varying reliability that depends on situational context. In contrast, such traits of perceptual intelligence have been notoriously and surprisingly difficult to implement in computer vision systems.

2. ALGORITHM DESCRIPTION

A general-purpose approach to object recognition, which was originally developed in the context of automatic target recognition (ATR), has been incorporated in a Machine Intelligent Gust Front Algorithm (MIGFA). Use of the term "machine intelligence" in particular reflects the use of two new techniques of knowledge-based signal processing.

The first technique, functional template correlation (FTC) (Delanoy, 1992) is a generalized matched filter incorporating aspects of fuzzy set theory. For comparison, standard 2-D cross correlation relies upon 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, resulting in a score in the 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. With various design strategies, the interfering effects of occlusion, distortion, noise, and clutter can be minimized. As a consequence, matched filters

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

customized for specific applications using FTC are generally more robust than traditional signal processing operations. The output of FTC is a map of values in the range [0,1], each of which reflects a 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 matching, and homotopic thinning of shapes.

The second major tool is the use of "interest" as a medium for data fusion and for assimilating evidence at the pixel level (Delanoy, 1991). An interest image is a map of evidence (values in the range [0,1]) for some feature that is selectively indicative of an object being sought (note that the output of FTC can be an interest image). Higher pixel values reflect greater confidence 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 to guide selective attention and serve as the input for object extraction. In practice, we often use several weakly or inconsistently discriminating feature detectors to mutually support or compensate for each other, resulting in relatively robust performance.

The system block diagram in Fig. 1 illustrates the configuration of the ASR-9 WSP version of the system. Input images V (Doppler velocity image) and DZ (reflectivity image) are passed to six simple, independent feature detectors that use FTC.

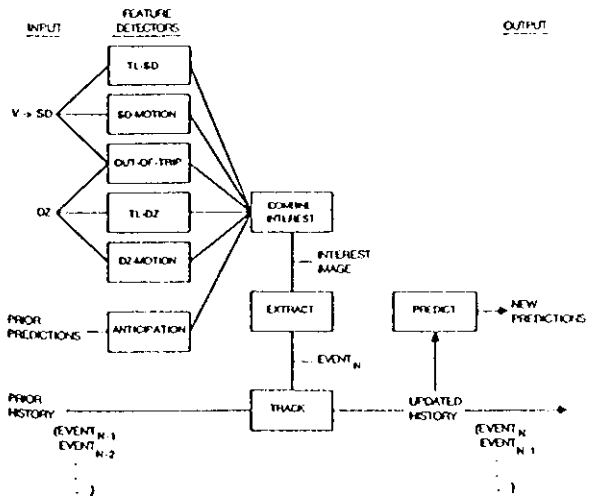

Figure 1. MIGFA block diagram.

The first one, marked TL-SD, uses an FTC filter for thin-lines in a map of local Doppler variance, called SD (standard deviation)[1]. A second feature detector, labelled SD-MOTION, applies a thin line filter to the difference of two sequential reflectivity images. Since differencing suppresses those signals that are stationary, the SD-MOTION detector tends to highlight only those thin lines that are moving. The third feature detector, labelled OUT-OF-TRIP, uses FTC to identify range ambiguous echoes of more distant weather. Areas believed to reflect out-of-trip weather are given low interest values. The next two feature detectors, TL-DZ and DZ-MOTION perform the same kind of thin-line analysis as is done for TL-SD and SD-MOTION. The final feature detector, ANTICIPATION, will be discussed below. The outputs of the several feature detectors are combined as a weighted average to form a combined interest image.

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 the prior history by establishing point-to-point correspondence. Heuristics are applied at this point to reject chain points that have apparent motion that is improbable. The updated history is used to make predictions of where points along the front will be located at some future time. Such predictions are used in the processing of subsequent images, specifically in the feature detector called ANTICIPATION, which places high interest values wherever fronts are expected to be and by so doing selectively sensitizing the system to detecting gust fronts at specific locations. Anticipation is tuned so that it will not by itself automatically trigger a detection, but when averaged with other interest images will support weak evidence that would otherwise be insufficient to trigger a detection.

A second version of the gust front algorithm has been constructed for TDWR data, differing only in the set of feature detectors used. Differences in the TDWR-detector set primarily reflect the greater dependence on Doppler data for finding areas of convergence.

3. RESULTS

A test set of ASR-9 WSP data collected in Orlando, Florida during AGFA field testing in 1991, contains 9 different gust fronts in a set of 450 images (15 hours). A human interpreter looking at the same ASR-9 WSP data detected 280 instances of the 9 gust fronts tracked by the radar. Four figures of merit are shown for each of the two algorithms. The probability of detection is the number of detections made by each algorithm as a percent of human-detected instances of gust fronts. The probability of false alarm (PFA) is the number of false alarms divided by the total number of algorithm detections. In addition to simply identifying fronts, the human interpreter delimited the length of each detected front. Detection quality was further assessed by comparing the length of the front as estimated by each algorithm against

1. Because measurements within gust front thin line have higher signal to noise ratios than the surrounding clear air measurements, gust fronts show up in SD as thin lines of lower signal variance.

that indicated by the human interpreter. The percent length detected (PLD) is the length detected expressed as a percent of the length delimited by the human interpreter. The percent of false length detected (PFD) reflects situations where the detected gust front lengths extended beyond what the human interpreter could see.

Table 1 compares performance of MIGFA against the previously constructed AGFA, which uses more conventional methods of signal processing and computer vision. The first two columns indicate that MIGFA, relative to AGFA, substantially increased the POD while decreasing the PFA. Similarly, the PLD reflects the improvement in detection rate. However, the increased PFD (from 13% to 33%) 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. We assume that truth derived from the more reliable TDWR data is more accurate than that for ASR-9 data. These results, shown in Table 2, confirm the general trend of the first 3 figures of merit shown in Table 1. However, now the PFD is quite low, essentially the same as that for AGFA. Consequently, these results, along with analyses of individual cases, leads us to believe that the MIGFA detected fronts were in fact more accurate than detections made by the human interpreter given the same ASR-9 WSP data.

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
AGFA	56.7	4.6	38.9	12.9
MIGFA	88.1	0.6	86.2	33.4

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

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

This same version of MIGFA was installed on an ASR-9 WSP in Orlando, Florida for operational testing during the summer of 1992, the results of which are shown in Table 3. During the period from 1 August to 20 September, MIGFA correctly detected and tracked approximately 75% of all gust fronts identified by human interpreters examining ASR-9 WSP data. TDWR and anemometer data were also used for verification. Those gust fronts that were missed either had reflectivity levels near the sensitivity limits of the ASR-9 or were obscured by storm cells along the front. The false alarm rate was under 2%. Although TDWR-based truth has not yet been compiled for these data, an analysis of individual cases again indicates that the relatively high PFD (21%) consists largely of believable extensions of gust fronts that were not identified by the human interpreter.

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

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

4. STATUS

MIGFA represents a substantial improvement in performance over previous efforts and is the prime candidate for deployment in production ASR-9 WSP systems. The TDWR version of MIGFA is scheduled for testing in summer 1993. Adaptations of the techniques used in MIGFA are currently being used, or are being considered, for other weather detection problems, including microburst prediction and sensor fusion.

REFERENCES

Delanoy, R.L., J.G. Verly, and D.E. Dudgeon, 1991: Pixel-Level Fusion Using Interest Images, *Proceedings of the 4th National Symposium on Sensor Fusion*, Orlando, FL, April 2-4, 1991.

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.

Eilts, M.D, 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, *Preprints, 4th International Conference on Aviation Weather Systems*, June 24-26, 1991, Paris.

Merritt, M.W., D. Klingle-Wilson, and S.D. Campbell, 1989: Wind Shear Detection with Pencil-Beam Radars, *Lincoln Laboratory Journal*, Vol. 2, No.3, 483-510.

Weber, M.E. and T.A. Noyes, 1989: Wind Shear Detection with Airport Surveillance Radars, *Lincoln Laboratory Journal*, Vol. 2, No. 3, 511-526.