DOT/FAA/NR-92/1

Project Report
ATC-185

A Comparison of the Performance of Two Gust Front Detection Algorithms Using a Length-Based Scoring Technique

D. L. Klingle-Wilson
M. F. Donovan
S. H. Olson
F. W. Wilson

8 May 1992

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS

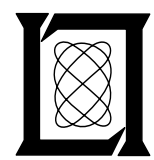

Prepared for the Federal Aviation Administration,
Washington, D.C. 20591

This document is available to the public through
the National Technical Information Service,
Springfield, VA 22161

This document is disseminated under the sponsorship of the Department of Transportation in the interest of information exchange. The United States Government assumes no liability for its contents or use thereof.

TECHNICAL REPORT STANDARD TITLE PAGE

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
ATC-185	DOT/FAA/NR-92/1	

4. Title and Subtitle	5. Report Date
A Comparison of the Performance of Two Gust Front Detection Algorithms Using a Length-Based Scoring Technique	8 May 1992
	6. Performing Organization Code

7. Author(s)	8. Performing Organization Report No.
D.L. Klingle-Wilson, M.F. Donovan, S.H. Olson, F.W. Wilson	ATC-185

9. Performing Organization Name and Address	10. Work Unit No. (TRAIS)
Lincoln Laboratory, MIT P.O. Box 73 Lexington, MA 02173-9108	11. Contract or Grant No. DTFA-01-89-Z-02033

12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
Department of Transportation Federal Aviation Administration Systems Research and Development Service Washington, DC 20591	Project Report
	14. Sponsoring Agency Code

15. Supplementary Notes

This report is based on studies performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. The work was sponsored by the Department of the Air Force under Contract F19628-90-C-0002.

16. Abstract

The Terminal Doppler Weather Radar (TDWR) Gust Front Algorithm provides, as products, estimates of the current locations of gust fronts, their future locations, the wind speed and direction behind the gust fronts, and the wind shear hazard to landing or departing aircraft. These products are used by air traffic controllers and supervisors to warn pilots of potentially hazardous wind shears during take-off and landing and to plan runway reconfigurations.

Until recently, an event-based scoring scheme was used to evaluate the performance of the algorithm. With the event-based scoring scheme, if any part of a gust front was detected, a valid detection was declared. Unfortunately, this scheme gave no indication of how much of the gust front length was detected; nor could the probabilities be easily related to the probability of issuing a wind shear alert for a specific approach or departure path which was being impacted by a gust front. To make the scoring metric more nearly reflect the operational use of the product, a new length-based scoring scheme was devised. This scheme computes the length of the gust front detected by the algorithm. When computed over a large number of gust fronts, this length-based scoring scheme yields the probability that any part of the gust front will be detected. As improvements to the algorithm increase the length detected, the probability of detecting any part of a gust front increases. In particular, an improved algorithm means an increased probability of correctly issuing wind shear alerts for the runways impacted by a gust front, and length-based scoring is a more accurate technique for assessing this probability of detection.

This paper describes the length-based scoring scheme and compares it with event-based scoring of the algorithm's gust front detection and forecast performance. The comparison of the scoring methods shows that recent enhancements to the gust front algorithm provide a substantial, positive impact on performance.

17. Key Words		18. Distribution Statement
gust front gust front algorithm wind shear	performance assessment air traffic control Terminal Doppler Weather Radar	This document is available to the public through the National Technical Information Service, Springfield, VA 22161.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	46	

ABSTRACT

The Terminal Doppler Weather Radar (TDWR) Gust Front Algorithm provides as products, estimates of the current locations of gust fronts, their future locations, the wind speed and direction behind the gust fronts, and the wind shear hazard to landing or departing aircraft. These products are used by air traffic controllers and supervisors to warn pilots of potentially hazardous wind shears during take–off and landing and to plan runway reconfigurations.

Until recently, an event–based scoring scheme was used to evaluate the performance of the algorithm. With the event–based scoring scheme, if **any part** of a gust front was detected, a valid detection was declared. Unfortunately, this scheme gave no indication of how **much** of the gust front length was detected; nor could the probabilities be easily related to the probability of issuing a wind shear alert for a specific approach or departure path which was being impacted by a gust front. To make the scoring metric more nearly reflect the operational use of the product, a new length–based scoring scheme was devised. This scheme computes the length of the gust front detected by the algorithm. When computed over a large number of gust fronts, this length–based scoring scheme yields the probability that any part of the gust front will be detected. As improvements to the algorithm increase the length detected, the probability of detecting any part of a gust front increases. In particular, an improved algorithm means an increased probability of correctly issuing wind shear alerts for the runways impacted by a gust front, and length–based scoring is a more accurate technique for assessing this probability of detection.

This paper describes the length–based scoring scheme and compares it with event–based scoring of the algorithm's gust front detection and forecast performance. The comparison of the scoring methods shows that recent enhancements to the gust front algorithm provide a substantial, positive impact on performance.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

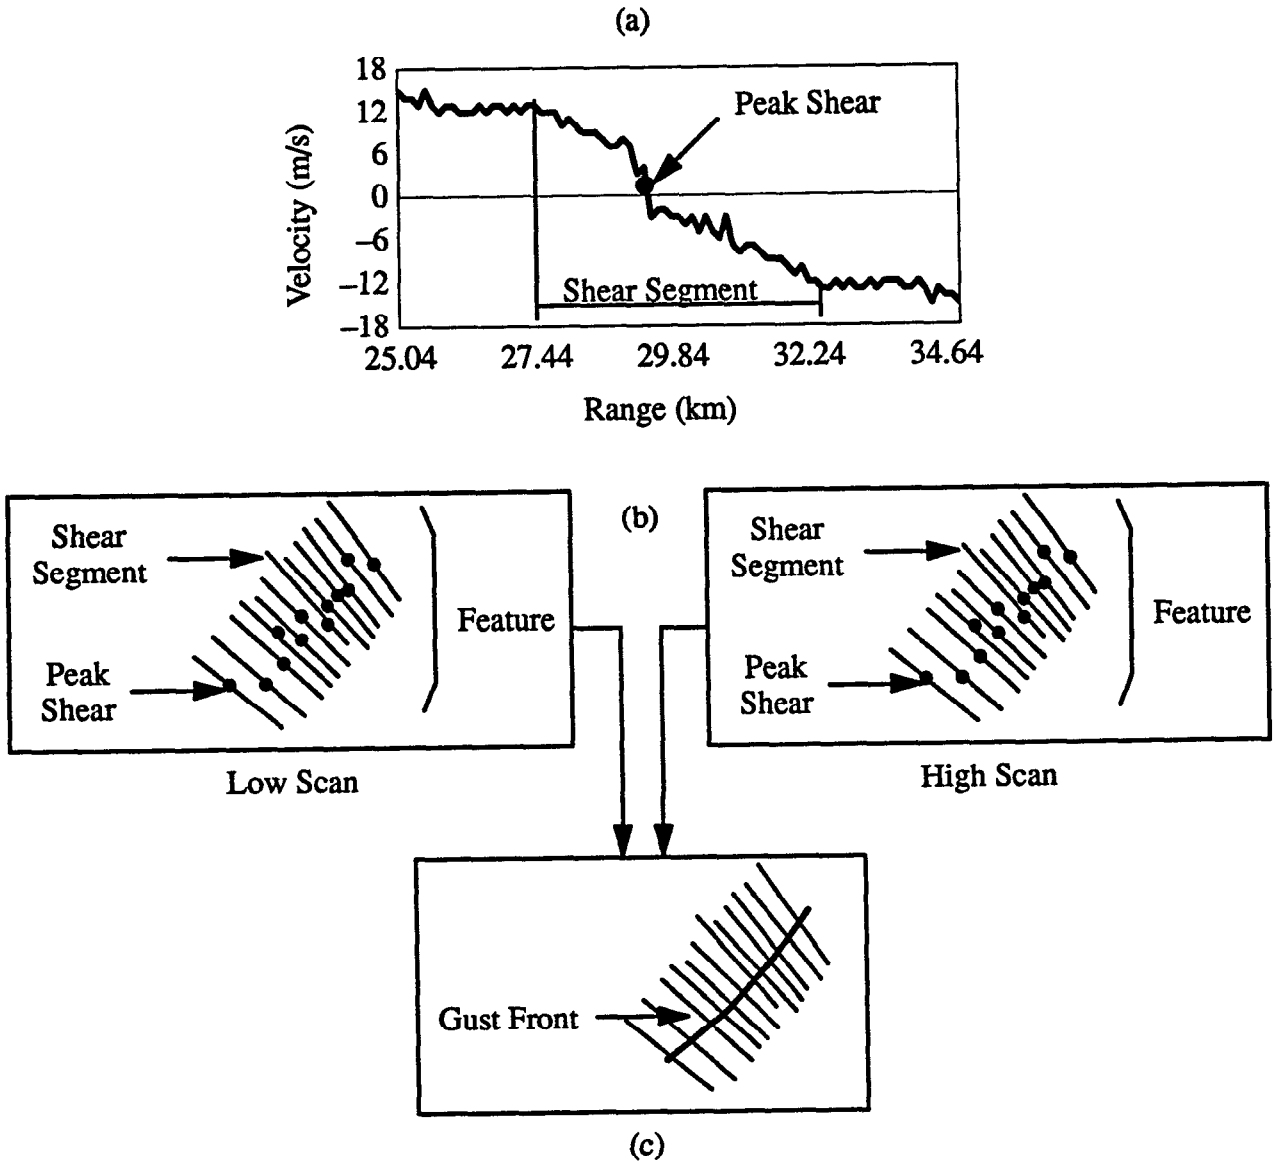

Figure 1. Diagram illustrating gust front detection. (a) Radial velocity data are searched for segments of decreasing Doppler velocity, which indicate radial convergence. The dots represent the locations of the peak shears along each segment. (b) Segments from the low and high scans are associated into features, and (c) features are grouped together based on spatial proximity into gust fronts.

To produce a forecast of the gust front location, the motion of the gust front with time must be established, which requires detections of the same gust front in two consecutive volume scans (Figure 2). The centroid of each detection on each tilt is computed. IGFA calculates the distance between the centroids on consecutive tilts and, if this distance passes a threshold, the detections are identified as the same gust front.

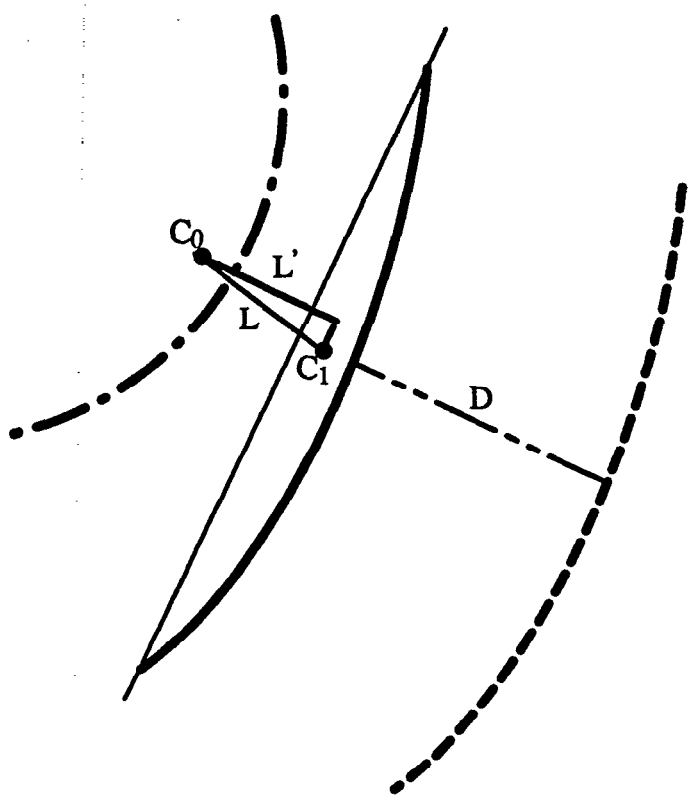

▬ • ▬	Previous gust front detection (at time t_0)
▬▬▬	Current gust front detection (at time t_1)
C_0	Centroid of gust front at time t_0
C_1	Centroid of gust front at time t_1
L	Centroid–to–centroid distance
———	Gust front orientation (at time t_1)
L'	Component of L perpendicular to gust front orientation
D	Distance gust front is forecast to move in 10 minutes
▬ ▬ ▬	Predicted 10–minute gust front location (for time t_1 + 10)

Figure 2. Diagram illustrating gust front prediction. For each detection, the orientation and centroid location of the detection are computed.

The gust front forecast is based upon its estimated propagation speed and direction. Gust fronts tend to propagate perpendicular to their orientations, so the component of the centroid–to–centroid vector perpendicular to the orientation is computed. The magnitude of the perpendicular component is divided by the time difference between consecutive detections to derive propagation speed. Propagation direction is parallel to the perpendicular–component vector. The forecast is simply the current

detection moved along the propagation direction by a distance defined by the propagation speed multiplied by the desired forecast time (typically 10 and 20 minutes).

IGFA attempts to estimate the wind speed and wind direction ahead of and behind the gust front. The algorithm uses data from the 0.5° tilt and assumes a uniform horizontal wind within specified spatial sectors. Estimates of the wind components are obtained by regressing the smoothed Doppler velocities within each sector onto sine and cosine functions, and minimizing the sum of the squared errors between the measurements and the fitted values (Smith 1986).

2.2 ADVANCED GUST FRONT ALGORITHM

In addition to the radial convergence detected by IGFA, AGFA makes use of additional signatures in the Doppler velocity and reflectivity data to detect gust fronts. Since a single–Doppler radar is only capable of resolving the component of velocities along the radar beam, velocity features which have components perpendicular to the beam are not easily observed. If shears are aligned across an azimuth, they often can be observed as an azimuthal variation of the wind field rather than as a radial variation. One component of AGFA attempts to use the information contained in azimuthal variations in Doppler velocity to augment estimates of radial convergence. This is referred to as azimuthal–shear detection.

Reflectivity thin lines are often associated with gust fronts. Unlike Doppler velocities, reflectivity is invariant with viewing direction and a reflectivity thin line can be identified independent of the viewing angle. Detecting reflectivity thin lines provides information on the location of gust fronts, especially when the gust fronts are oriented so that radial convergence is not readily observed.

In addition to radial–convergence detection described in Section 2.1, AGFA can use either reflectivity thin lines, azimuthal shears, or both to improve detections. A description of one technique for detecting azimuthal shears is presented by Eilts *et al.* (1991). The version of AGFA used in this study incorporates only reflectivity thin line detection, which is detailed below, and is referred to hereafter as AGFA–TL.

Using reflectivity thin line features as a means of identifying gust fronts presents four difficulties: 1) not all gust fronts have thin lines; 2) the reflectivities in the thin line signatures are usually not much larger than the background reflectivities and the linear patterns are difficult for the algorithm to identify, even when they are apparent to a human; 3) the appearance of a thin line does not indicate the strength of the convergent boundary that causes it and 4) some meteorological phenomena (*e.g.*, cloud streets) and radar data artifacts (*e.g.*, range folding) that are not gust fronts are associated with reflectivity thin lines.

The reflectivity thin line feature detection algorithm initially subjects the reflectivity data to a median filtering technique. The basis of the thin line feature detection algorithm is the use of multi–

thresholding and shape analysis to isolate thin lines in the reflectivity field. To find segments of a thin line along radials, the algorithm searches along the radial, finding runs of reflectivity values that are above one or more thresholds. These thresholds are currently set at 0, 2.5, 5, 7.5, 10, 12.5, and 15 dBZ. All threshold values are processed in parallel, thus any given reflectivity datum may be a part of several segments at once. This leads to the situation where a single "hump" is found to contain segments at several threshold levels (Figure 3).

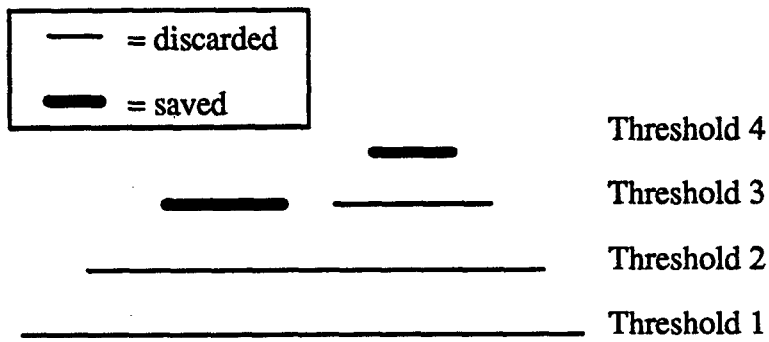

Figure 3. Overlapping segments generated at four threshold levels.

Segments are constrained to have a minimum (1–1.5 km) and a maximum (about 4 km) length. This helps reduce incorrect associations. The algorithm has the ability to skip over a radial that contains no nearby segments. It is not unusual for even a strong thin line to have a few individual radials where the algorithm fails to detect a valid segment. This ability to skip a segment is helpful in reducing the fragmentation of the detected thin line.

After segments are grouped together, reflectivity thin line features are constructed. A feature is described by a sequence of points representing the peak reflectivity point of each segment in the feature. Properties of the features such as length, area, maximum dBZ, minimum dBZ, and average dBZ are computed. These are used later to help discriminate against false features caused by data artifacts such as range folding and velocity folding.

Because of the past success of the radial convergence feature detection algorithm, especially the low false alarm ratio, it is used as the starting point for the multiple feature association. On a given tilt, an attempt is made to associate features from the combined shear (if present) and reflectivity thin line algorithms with features from the radial convergence algorithm. These features are joined using an endpoint proximity check. If either endpoint of the two features are within 5 km of each other they are joined. To ensure that the features are a good match their orientation must also differ by less than 30 degrees. All possible combinations of features on a tilt are joined together. Only features that have one radial convergence feature as part of the combination are considered candidate gust fronts and are used for further association checks.

In order to minimize the number of false alarms, features from the two elevation scans are required to be vertically associated for a gust front detection to be declared. Feature endpoints are com-

pared with the peak shear locations of all features on the other tilt. If the endpoint of one feature is within 2 km of any point on another feature on the opposite tilt, the two features become vertically associated. All possible combinations of features from the two tilts are then put together to determine the total gust front.

The gust front detection that results from the merging of the various features is not a smooth curve. The smooth curve representing the gust front is generated by fitting a least–squares polynomial (in x,y) to the peak shear locations in the features that have been vertically associated.

If a gust front is detected on two consecutive scans, an attempt is made to establish time continuity between the pair of detections. AGFA uses 1 km sections along the front to determine the association of old and new fronts. If a sufficient percentage of old front sections are within some critical distance (along a perpendicular line) of the new front, the old and new fronts are time associated. A secondary technique, used in cases where the other technique fails, bases time association on the overall gust front orientation angle and the distance between midpoints. This technique is designed to time–associate front pairs which have similar shapes.

The propagation speed of the gust front is determined by averaging the distance along 1 km sections between the current gust front and its time–associated partner from the previous scan. Those distances which lie more than 2 standard deviations from the mean are rejected, and a new average is calculated with the remaining values. The polynomial representation of the current front is then propagated forward, in the direction of the average perpendicular vector between the two fronts, the average speed determined from the average distance. Forecasts are generally produced for 10– and 20–minute periods.

The wind shift and wind shear estimates are computed using the same technique as IGFA.

3. DETECTION SCORING

3.1 EVENT–BASED SCHEME FOR SCORING DETECTIONS

The basic statistics used in the event–based scoring technique are the Probability of Detection (POD) and Probability of False Alarm (PFA). POD is the total number of detections divided by the total number of events. A valid detection is declared if **any part** of an algorithmic declaration overlaps a truth box. The PFA is defined as the number of false detections divided by the total number of detections (true plus false), where a false detection is declared if there was no truth box overlapping a detection.

POD is the probability that some part of an event is detected, but it does not indicate how well an event is detected. Figure 4 provides an example of detections made by two different algorithms.

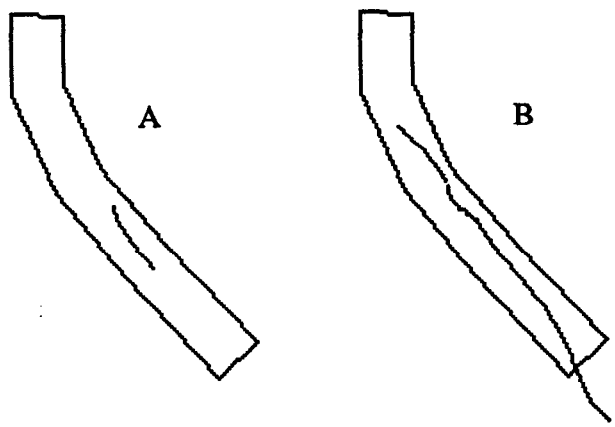

Figure 4. Example of valid detections using the event–based scoring scheme. The rectangular shape is a gust front truth box and the solid line represents a detection.

Clearly, the detection in Figure 4B is better than the detection in Figure 4A. However, both algorithms are credited with a detection and the resulting POD is the same for both algorithms. A measure of the "goodness" of the detection is provided by the Percent of Length Detected (%L). To compute %L, a truth box is divided lengthwise into smaller bins whose widths (typically 1 kilometer) are user–specified (Figure 5). Percent of Length Detected is the number of bins "hit" by a detection divided by the total number of bins in the truth box, expressed as a percentage. Thus, %L for the event depicted in Figure 5 is about 60 percent.

3.2 LENGTH–BASED SCHEME FOR SCORING DETECTIONS

The goal of the recent work on the gust front algorithm has been to improve detections by increasing the detected length of the gust front. The event–based scoring technique does not test how

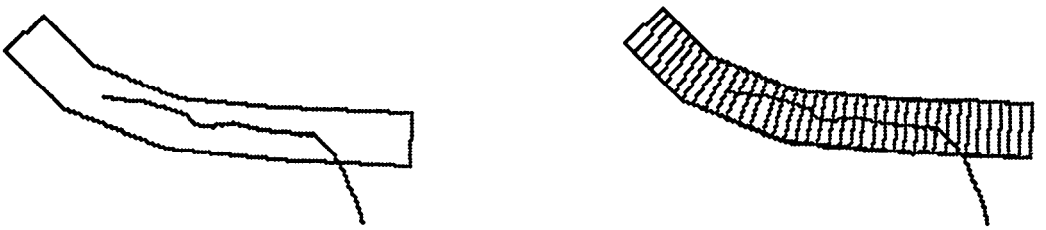

Figure 5. Truth box (A) and truth box subdivided into bins (B). The solid line represents a detection.

well that goal is met. The length–based scheme assumes that the probability of an event occurring over the area scanned by the radar is evenly distributed. Then, the probability that a gust front passing through a specific location is detected at that location is given by the Probability of Local Detection (PLD) or:

$$PLD = \frac{\text{length of event detected}}{\text{total length of events}}$$

For Figure 5B, the PLD is about 60 percent.

Another flaw in event–based scoring is that the algorithm is not penalized for portions of detections that fall outside of truth boxes, which may be considered false alarms. For example, assume that the portion of the detection that lies outside the truth box is a result of data quality problems. If that portion of the detection crosses an airport with no resulting wind shear or wind shift, it is considered a false alarm by the air traffic user. However, this is not considered a false alarm in the event–based scoring scheme. For the length–based technique, the probability that a specific location on a gust front detection is not a location on an actual gust front is the Probability of False Detection (PFD). This is equal to the length of detections that fall outside the truth box divided by the total length of detections or:

$$PFD = \frac{\text{length of detections not associated with event}}{\text{total length of detections}}$$

For Figure 5B, the PFD is about 30 percent.

3.3 OVER–THE–AIRPORT SCORING

Until recently, scoring of those gust fronts that impacted the airport was performed manually. Events were compared to detections to determine if a gust front that passed over the airport was detected on the airport. This process was very time–consuming and labor–intensive, which resulted in a very small number of events that could be used to generate performance statistics. Automation of the over–the–airport scoring task was implemented with the length–based scoring scheme. Now, only those portions of gust fronts that fall within a user–specified range of the radar (*e.g.*, 30 kilometers) are considered eligible to impact the airport and are scored. In this scheme, the over–the–airport score is no longer tied to a specific airport but pertains to any airport located within 30 kilometers of the radar. In addition to the work–load advantages of automated scoring, this approach incorporates a larger number of gust fronts into the estimation of the detection probabilities.

3.4 PERFORMANCE STATISTICS FOR DETECTIONS

Table 1 provides an example of the differences between the event–based and length–based scoring. The site–adaptable parameters for the algorithms are listed in Appendix A, Table A–1. The values from which the probabilities in Table 1 are derived are given in Table A–2. Gust front strength is determined by the average peak change in Doppler velocity (ΔV) across the gust front. The strength of a gust front is defined as "weak" for 5 m/s $\leq \Delta V < 10$ m/s; "moderate" for 10 m/s $\leq \Delta V < 15$ m/s; "strong" for 15 m/s $\leq \Delta V < 25$ m/s; and "severe" for $\Delta V \geq 25$ m/s. PLD is less than POD for both IGFA (0.30 versus 0.50) and AGFA–TL (0.41 versus 0.57) algorithms. The POD for both algorithms is roughly the same (0.50 for IGFA and 0.57 for AGFA–TL), which indicates that both algorithms detected some part of the same number of events. PLD shows that the AGFA–TL algorithm (0.41) shows that AGFA–TL detected a greater portion of the lengths of the events than the IGFA algorithm (0.30), an improvement of about 37 percent.

PLD and PFD as functions of location are provided in Table 2. Table A–3 provides the values from which these probabilities are computed. Except for severe Kansas City gust fronts (where the performance was the same for both algorithms), the AGFA–TL algorithm detected gust fronts better than the IGFA algorithm. The PFD for the IGFA and AGFA–TL algorithms (all locations) is 0.17 and 0.21, respectively, which indicates that about 20 percent of the lengths of detections by both algorithms are considered false (*i.e.*, they do not overlap truth). The increase in performance of AGFA–TL over IGFA also is summarized in Table 2. The greatest improvement in detection capability (PLD) as a function of location is seen in Orlando. The greatest improvement as a function of location is associated with weak gust fronts.

PLD for severe Orlando gust fronts is less than other locations for both algorithms. This statistic is based on two observations of the same gust front, each 63 km long. As is often the case, this gust front triggered the development of a storm cell. Since events are classified according to peak strength, the very strong outflow from the storm caused these events to be classified as severe. For nearly 75 percent of their lengths, these events are weak or moderate. The algorithms detected the

stronger portions of the events. AGFA–TL detected more of the weak portion than IGFA, with an increase in PLD over IGFA of 92 percent.

TABLE 1.
POD* and PLD statistics for all gust front strengths
for the IGFA and AGFA–TL algorithms.

	POD	PLD
IGFA		
Weak	0.30	0.15
Moderate	0.74	0.40
Strong	0.87	0.45
Severe	1.00	0.58
All Strengths	0.50	0.30
AGFA–TL		
Weak	0.38	0.25
Moderate	0.79	0.54
Strong	0.90	0.54
Severe	1.00	0.60
All Strengths	0.57	0.41
Percent Increase (AGFA–TL over IGFA)		
Weak	27%	67%
Moderate	7%	35%
Strong	3%	20%
Severe	0%	3%
All Strengths	14%	37%

*POD is the number of events detected divided by the total number of events. PLD is the length of events detected divided by the total length of all events.

TABLE 2.
PLD and PFD* statistics for all gust front strengths and locations for the IGFA and AGFA–TL algorithms.

	Weak	Moderate	Strong	Severe	All Strengths	
	PLD					PFD
IGFA						
Denver	0.15	0.43	0.48	0.61	0.29	0.11
Kansas City	0.21	0.39	0.45	0.62	0.37	0.27
Orlando	0.13	0.36	0.38	0.25	0.25	0.10
All Locations	0.15	0.40	0.45	0.58	0.30	0.17
AGFA–TL						
Denver	0.24	0.55	0.60	0.64	0.40	0.19
Kansas City	0.25	0.46	0.51	0.62	0.43	0.27
Orlando	0.26	0.57	0.53	0.48	0.41	0.16
All Locations	0.25	0.54	0.54	0.60	0.41	0.21
Percent Increase (AGFA–TL over IGFA)						
Denver	60%	28%	25%	5%	38%	
Kansas City	19%	18%	13%	0%	16%	
Orlando	100%	58%	39%	92%	64%	
All Locations	67%	35%	20%	3%	37%	

*PFD is the length of false detections divided by total length of all detections. The Percent Increase of PLD of AGFA–TL over IGFA is given.

Table 3 provides over–the–airport performance statistics for the IGFA and AGFA–TL algorithms. The values used to compute these results are given in Table A–4. These data show that AGFA–TL detects gust fronts passing over the airport better than IGFA. The greatest improvement is associated with weak gust fronts and with Orlando gust fronts. PFD increased for Orlando and Denver but remained the same for Kansas City, resulting in an overall increase in PFD over IGFA.

A comparison of Table 2 and Table 3 shows that the over–the–airport performance for both algorithms (within 30 kilometers) is better than the performance of the algorithm within 60 kilometers. This suggests that the algorithms detect close gust fronts better than distant gust fronts.

TABLE 3.
Over–the–Airport PLD and PFD statistics for all gust front strengths and locations for the IGFA and AGFA–TL algorithms.

	Weak	Moderate	Strong	Severe	All Strengths	
	PLD					PFD
IGFA						
Denver	0.26	0.54	0.61	0.61	0.41	0.13
Kansas City	0.20	0.52	0.55	0.83	0.46	0.38
Orlando	0.19	0.49	0.54	0.39	0.35	0.12
All Locations	0.23	0.52	0.57	0.63	0.40	0.19
AGFA–TL						
Denver	0.35	0.62	0.69	0.64	0.49	0.17
Kansas City	0.24	0.60	0.61	0.84	0.53	0.35
Orlando	0.29	0.68	0.72	0.64	0.49	0.19
All Locations	0.33	0.64	0.65	0.73	0.50	0.21
Percent Increase (AGFA–TL over IGFA)						
Denver	35%	15%	13%	5%	20%	
Kansas City	20%	15%	11%	1%	15%	
Orlando	53%	39%	33%	64%	40%	
All Locations	43%	23%	12%	16%	25%	

4. FORECAST SCORING

4.1 MANUAL SCORING TECHNIQUE

Manual (event–based) scoring of forecasts was performed by comparing a forecast issued at time T to the truth at the time for which the forecast was valid (time T+F). For example, a 10–minute forecast (F = 10) issued at 12:10 (T) was compared to the truth at 12:20 (T+F). To determine if a forecast was a candidate for scoring, the detection at time T had to overlap a truth box at time T. This eliminated scoring of forecasts from false detections that were persistent enough to generate a forecast. If the forecast issued at time T overlapped a truth box at time T+F, a valid forecast was declared. If there was no truth box associated with the forecast, a false forecast was declared. If the forecast was associated with a truth box but the two did not overlap, a miss was declared.

Figure 6 provides an example of the manual forecast scoring technique. In this example, four detections were declared and forecasts were issued at time T, but only three of the detections were candidates for scoring (those associated with GF1, GF2, and GF3). The fourth detection was a false detection, and although a forecast was issued, the forecast was not scored. At time T+F, the forecast

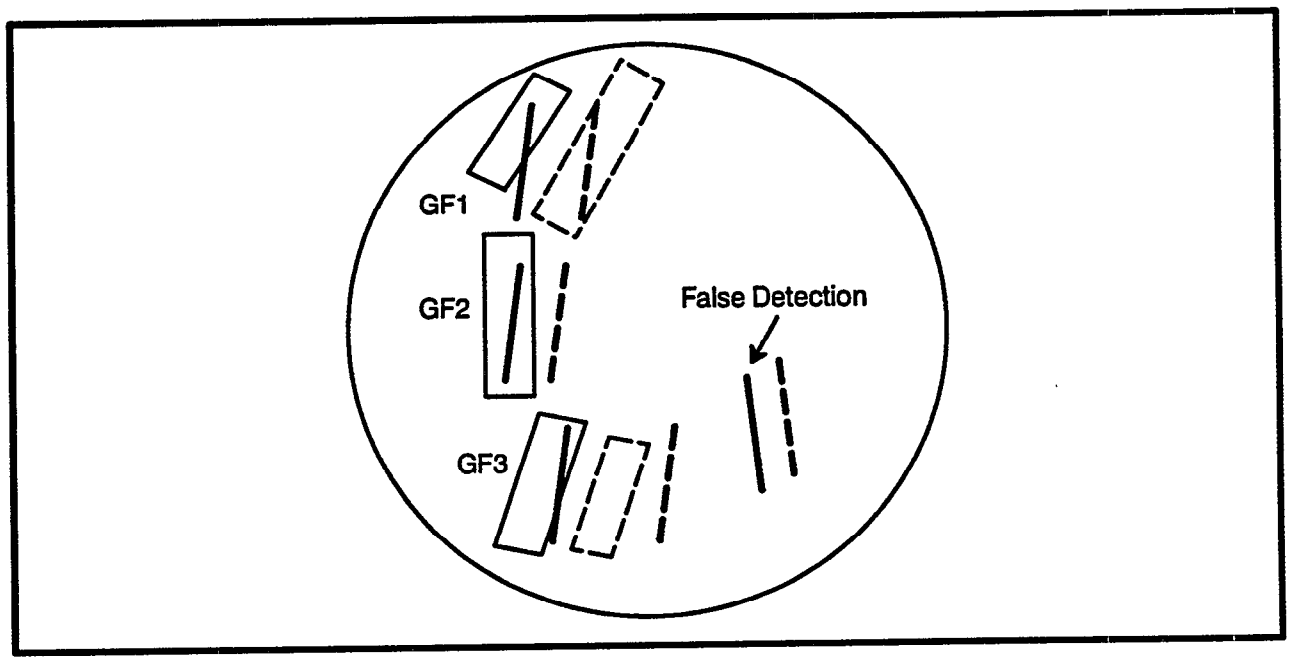

Figure 6. Illustration of the manual forecast scoring procedure. Solid lines are detections and dashed lines are forecasts made at time T for time T+F. Solid boxes are truth at time T, dashed boxes are the truth at time T+F. GF1, GF2, and GF3 are event identifiers. The forecasts for GF1 and GF3 are eligible for evaluation because the generating gust front still exists. The forecast for GF2 is false because GF2 has dissipated. The forecast generated by the false detection is not considered in scoring.

and event GF1 overlapped, so a valid forecast was declared. GF2 had dissipated, so the forecast associated with this event was a false forecast. Since the forecast for GF3 did not overlap a truth box, a miss was declared.

The metrics for manually evaluating the forecast function are the Probability of a Correct Forecast (PCF) and Probability of False Forecast (PFF). These are given by:

$$PCF = \frac{\text{number of valid forecasts}}{\text{number of events forecasted}}$$

$$PFF = \frac{\text{number of false forecasts}}{\text{number of events forecasted + number of false forecasts}}$$

Missed forecasts, which represent a timing error and are generally caused by an incorrect estimate of the propagation speed, are not represented in these statistics. However, these missed forecasts were analyzed to determine the extent of the timing error.

Gust front forecasts were scored manually for data collected during the TDWR operational demonstrations of 1988 through 1990. The scores of the 10– and 20–minute forecasts for 1988 (Denver), 1989 (Kansas City), and 1990 (Orlando) are presented by Bernella (1991). Overall, PCF was better than 0.95 for the 10–minute forecasts and better than 0.75 for the 20–minute forecasts. PFF for the 10– and 20–minute forecasts averaged near 0.10 and 0.18, respectively. The high PCF values show that, when generated, forecasts were very accurate.

4.2 LENGTH–BASED SCHEME FOR SCORING FORECASTS

To properly evaluate the forecast function, it is necessary to know how well gust fronts are forecasted and how well forecasts verify. For gust fronts that impact an airport, it is important to assess if the user (ATC supervisor) received sufficient notification. Alternatively, if a gust front is forecasted to impact an airport, it must be assessed whether or not a gust front did indeed impact the airport.

Figure 7 illustrates how one determines the probability that events at time T were forecasted at time T–F. Each truth box at time T is compared to forecasts generated at time T–F. If an overlap occurs, a hit is declared. If the truth does not overlap with a forecast, a miss is declared. If a forecast does not overlap a truth box, a false forecast is declared. In this figure, all events present at time T are compared to forecasts issued at time T–F. The forecasts issued at time T–F that overlap an event at time T are correct.

As with detection scoring, it is desirable to use a length–based scoring scheme for scoring forecasts. This provides a clear indication of the impact of algorithm changes on algorithm performance

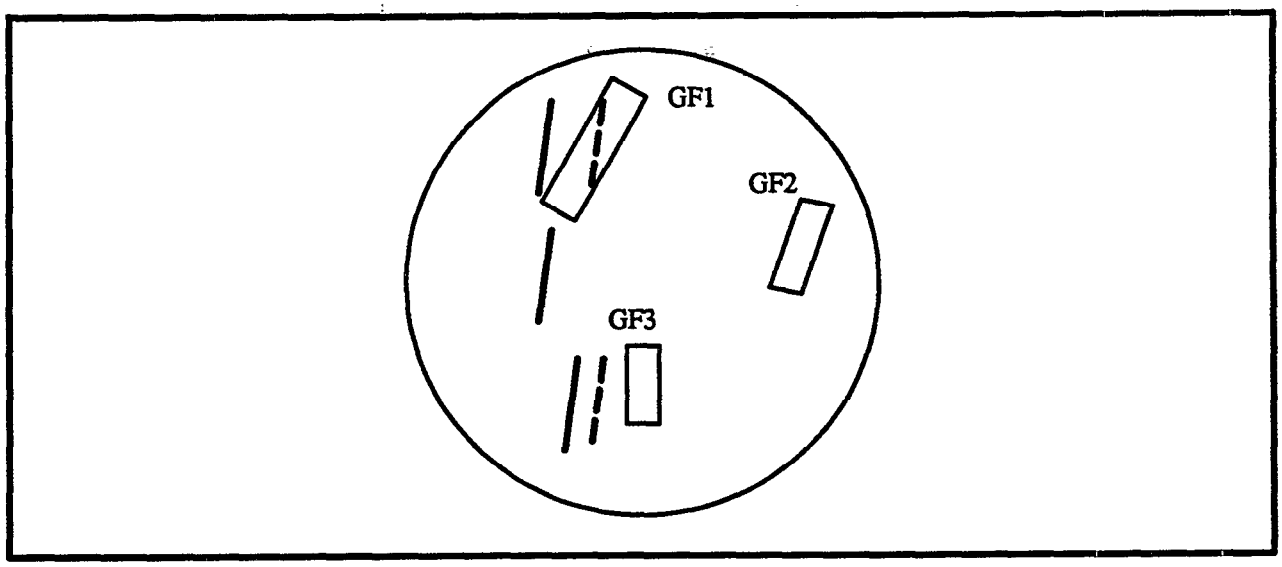

Figure 7. Illustration of the automated forecast scoring procedure. Solid lines are detections and dashed lines are forecasts made at time T. Forecasts are valid for time T+F. Solid boxes are truth at time T+F. GF1, GF2, and GF3 are event identifiers. The forecast for GF1 is valid because it overlaps a truth box. The forecast near GF3 is false because it does not overlap the truth box. GF2 and GF3 are missed events.

and consistency throughout the performance assessment. Length–based scoring for forecasts is performed in a manner analogous to detection scoring. The Correct Forecast Probability (CFP) and the False Forecast Probability (FFP) are given by:

$$CFP = \frac{\text{length of events overlapped by forecast}}{\text{total length of events}}$$

$$FFP = \frac{\text{length of forecasts not overlapped by event}}{\text{total length of forecasts}}$$

These statistics, computed over a large number of events and forecasts, provide an estimate of how well the algorithm performs the forecasting function.

Another metric of interest is the Probability of generating a Forecast (PF), an event–based metric. There are two ways to compute this metric. The first is to divide the number of forecasts by the number of **detections**. This measure of forecasting ability is denoted PFld, for the Probability of generating a Forecast given a detection. PFld is of interest to algorithm developers because the algorithms must detect a gust front on two consecutive volume scans and correctly associate those detections in order to generate a forecast. PFld estimates how well the algorithms perform that function. The second method for computing PF is to divide the number of forecasts by the number of **events**, which is denoted PFle, for PF given an event. This is of interest to the product users who wish to

know how well the algorithm forecasts events that impact the airport, regardless of whether or not the event was detected.

4.3 PERFORMANCE STATISTICS FOR FORECASTING

PF|d for forecasts from the IGFA and AGFA algorithms are shown in Table 4. These data are further categorized according to the strength of the gust fronts at the time the forecast verified, where PF|d is the number of correct forecasts divided by the number of detections of a given strength. This is interpreted as the probability of generating a forecast for a gust front of a given strength given a detection. PF|d for false forecasts is the number of false forecasts divided by the total number of detections (all strengths). This is the probability of generating a false forecast given a detection. PF|d for all forecasts is the total number of forecasts divided by the total number of detections. The values used to compute these statistics are given in Table A–5. The "All Forecasts" column indicates that given a detection, the probability that a forecast will be generated by IGFA and AGFA–TL is 0.81 and 0.79, respectively. This implies that AGFA–TL does not produce forecasts as well as IGFA. It is important to note that for moderate, strong, and severe gust fronts, AGFA–TL is consistently equal to or better than IGFA. In addition, AGFA–TL generates a smaller percentage of false forecasts than IGFA. It is evident from Table A–5 that AGFA–TL generates more forecasts than IGFA, but it also declares more detections. The increase in detections slightly out–paces the increase in forecasts. Thus, overall PF|d for AGFA–TL is slightly less than for IGFA.

Table 4.
PF|d* from the IGFA and AGFA algorithms.

	Weak	Moderate	Strong	Severe	False Forecasts	All Forecasts
IGFA						
Denver	0.71	0.65	0.74	1.00	0.07	0.76
Kansas City	0.40	0.49	0.72	0.63	0.30	0.85
Orlando	0.66	0.78	0.75	1.00	0.12	0.85
All Locations	0.63	0.64	0.73	0.68	0.15	0.81
AGFA–TL						
Denver	0.65	0.70	0.79	1.00	0.06	0.75
Kansas City	0.40	0.55	0.79	0.74	0.23	0.83
Orlando	0.69	0.77	0.81	1.00	0.10	0.85
All Locations	0.62	0.69	0.79	0.77	0.11	0.79

*PF|d is the number of forecasts divided by the number of detections.

PFIe for IGFA and AGFA–TL is given in Table 5 and Table A–6. For the various strength categories, PFIe is the number of correct forecasts divided by the number of events of a given strength. This is interpreted as the probability of generating a forecast for a gust front of a given strength given an event of that strength. PFIe for false forecasts is the number of false forecasts divided by the total number of events (all strengths) and is the probability of generating a false forecast given an event of any strength. PFIe for all forecasts is the total number of forecasts (correct and false) divided by the total number of events. PFIe is greater for AGFA–TL than for IGFA for all strengths and the percentage of false AGFA–TL forecasts is lower. The percentage improvement of AGFA–TL over IGFA is also shown. A negative percentage in the False Forecast column indicates a performance increase (i.e., fewer false forecasts are issued). The greatest increase in PFIe as a function of location is associated with Orlando. The greatest increase as a function of strength is associated with severe gust fronts (although these statistics are based upon small numbers). As shown in Table A–5, the number of events is constant throughout the evaluation. Thus, the improvement in forecast ability of AGFA–TL over IGFA (15 percent) is better shown by PFIe.

Table 5.
PFIe from the IGFA and AGFA algorithms.

	Weak	Moderate	Strong	Severe	False Forecasts	All Forecasts
IGFA						
Denver	0.22	0.49	0.68	1.00	0.07	0.37
Kansas City	0.14	0.35	0.57	0.63	0.17	0.49
Orlando	0.15	0.56	0.71	1.00	0.05	0.38
All Locations	0.19	0.48	0.63	0.68	0.07	0.40
AGFA–TL						
Denver	0.26	0.56	0.75	1.00	0.03	0.42
Kansas City	0.15	0.42	0.66	0.74	0.14	0.50
Orlando	0.25	0.62	0.76	1.00	0.05	0.46
All Locations	0.23	0.54	0.71	0.77	0.06	0.45
Percent Increase (AGFA–TL over IGFA)						
Denver	18%	14%	10%	0%	–57%	14%
Kansas City	7%	20%	16%	17%	–18%	2%
Orlando	67%	11%	7%	0%	0%	21%
All Locations	21%	13%	13%	13%	–14%	13%

*PFIe is the number of forecasts divided by the number of events.

CFP and FFP statistics for the 10–minute forecasts generated by the IGFA and AGFA–TL algorithms are presented in Table 6 and Table A–7. In general, the AGFA–TL algorithm out–performs the IGFA algorithm, especially for weak and moderate events. The CFP increases for all locations (with the greatest improvement seen in Orlando) and all strengths (with the greatest improvement associated with moderate gust fronts). On the other hand, FFP increases slightly for Denver and Orlando and decreases for Kansas City. For AGFA–TL (all strengths and locations), the probability that a gust front is correctly forecast is 0.27 and the probability that a forecast will not verify is 0.36.

TABLE 6.
CFP* and FFP statistics for the 10–minute forecasts
for the IGFA and AGFA–TL algorithms.

	Weak	Moderate	Strong	Severe	All Strengths	
	CFP					FFP
IGFA						
Denver	0.10	0.19	0.25	0.03	0.15	0.34
Kansas City	0.06	0.20	0.29	0.28	0.20	0.46
Orlando	0.09	0.26	0.21	0.45	0.17	0.26
All Locations	0.09	0.21	0.26	0.29	0.17	0.36
AGFA–TL						
Denver	0.16	0.37	0.37	0.03	0.26	0.37
Kansas City	0.08	0.25	0.39	0.30	0.25	0.42
Orlando	0.17	0.41	0.40	0.67	0.30	0.30
All Locations	0.15	0.36	0.38	0.33	0.27	0.36
Percent Increase (AGFA–TL over IGFA)						
Denver	60%	95%	48%	0%	73%	
Kansas City	33%	25%	34%	7%	25%	
Orlando	89%	58%	90%	49%	76%	
All Locations	67%	71%	46%	14%	59%	

*CFP is the length of events overlapped by forecasts divided by total length of all events. FFP is the length of forecasts not overlapped by events divided by the total length of all forecasts. The Percent Increase of the performance metrics of AGFA–TL over IGFA is given.

Table 7 and Table A–8 provide the performance statistics for the 20–minute forecasts. Again, AGFA–TL performed better than IGFA in all strength and location categories, with the greatest improvement for Orlando gust fronts and weak gust fronts. The 20–minute forecast performance of both algorithms is poorer than the 10–minute forecast. This is expected since forecasts are based solely on gust front propagation and do not take into account gust front evolution. Therefore, the longer range the forecast, the less accurate it is likely to be. To improve 20–minute forecasts, it may be necessary to incorporate an acceleration term into the propagation estimate.

TABLE 7.
CFP and FFP statistics for the 20–minute forecasts for the IGFA and AGFA–TL algorithms.

	Weak	Moderate	Strong	Severe	All Strengths	
	CFP					FFP
IGFA						
Denver	0.09	0.20	0.15	0.0	0.14	0.53
Kansas City	0.03	0.14	0.21	0.14	0.13	0.63
Orlando	0.07	0.19	0.15	0.34	0.13	0.43
All Locations	0.07	0.19	0.18	0.16	0.13	0.54
AGFA–TL						
Denver	0.12	0.29	0.24	0.0	0.20	0.54
Kansas City	0.04	0.18	0.26	0.15	0.17	0.59
Orlando	0.13	0.29	0.23	0.40	0.21	0.48
All Locations	0.11	0.27	0.25	0.18	0.20	0.54
Percent Increase (AGFA–TL over IGFA)						
Denver	33%	45%	60%	0%	43%	
Kansas City	33%	29%	24%	7%	31%	
Orlando	86%	53%	53%	18%	62%	
All Locations	57%	42%	39%	13%	54%	

4.4 OPERATIONAL ASSESSMENT OF FORECAST FUNCTION

Although it has been shown that AGFA–TL forecasts better than IGFA, the actual numbers are quite low. The IGFA algorithm was used operationally in the ATC towers at Kansas City International Airport and Orlando International Airport. Surveys pertaining to the usefulness of the various TDWR products were given to air traffic controller and supervisors following those demonstrations. The evaluators were asked to rate the usefulness of the gust front products on a scale of +3 (very good) to –3 (very poor). The users rated the gust front product as "fairly good" (+1).

This apparent discrepancy between the automated and user performance assessments may arise because the automated scoring technique uses a more stringent criteria for success. Operationally, forecasts serve as a "heads up" warning to ATC supervisors. Forecast information is used to coordinate runway configurations (and possible changes) with the various air traffic managers. Typically, no action is taken until the wind shift associated with the gust front is confirmed (e.g., by the presence of blowing dust or a wind change at an outlying anemometer). Errors in time–of–arrival or location are not as important to the users as the information that a wind shift is approaching. So, although the automated assessment may indicate bad forecast performance by both algorithms, the ATC user considers the product useful. The strength of the automated scoring methodology lies in its ability to provide a baseline against which changes to the algorithm can be evaluated easily.

5. SUMMARY

The gust front detection algorithm is an integral part of the TDWR system. Its purpose is to provide warnings of potentially hazardous wind shears to pilots of landing and departing aircraft and provide information of impending wind shifts to ATC supervisors. Scoring gust front algorithm products is important for evaluating how well the algorithm performs. The ATC user is primarily concerned with three things: Was the algorithm reliable at detecting gust front–related wind shears impacting the approach and departure paths? Was a forecast received for a gust front that impacted the airport? When a gust front was forecasted to impact the airport, was the forecasted location and/or time of arrival accurate (i.e., did the forecast verify)?

In the past, an event–based scoring scheme was used to assess algorithm performance. For this scheme, a valid detection is declared if any part of an algorithmic declaration overlaps a truth box, where a false detection is declared if no truth box overlaps a detection. The algorithm's performance is not penalized for portions of detections that fall outside of truth boxes and that cross an airport with no wind shear or wind shift. These instances may be considered false alarms by the air traffic user. In addition, the event–based scoring technique cannot easily be related to the probability of issuing an alert for a specific runway.

A length–based scoring technique was implemented to provide greater precision in evaluating the algorithm's ability to detect gust fronts. This scoring technique compares the length of the detection (or forecast) to the length of the ground truth. The length–based scheme assumes that all areas within a given range from the radar are of equal interest when considered over the ensemble of all TDWR sites. Based on this, the probability of a gust front passing through a specific location being detected at that location is given by the PLD (length of event detected divided by total length of event). The probability that a specific location on a gust front detection is not a location on an actual gust front is the PFD (length of detections not associated with event divided by total length of detections). When computed over a large number of gust fronts, length–based scoring yields the probability that any part of the gust front will be detected or forecasted.

An assessment of the performance of the initial–deployment TDWR gust front algorithm (IGFA) and an enhanced version (AGFA–TL) were presented. It was shown that AGFA–TL detects more of the gust front length than IGFA, although the probability of false detection also increases. The improvement of AGFA–TL over IGFA is about 40 percent.

The ability of the gust front algorithm to create good forecasts is of great interest to the ATC user. It was shown that on an event basis, AGFA–TL detects more gust fronts than IGFA. However, the AGFA–TL forecasts keep pace with the increase in detections. As a result, the probability of generating a forecast, given a detection, is the same for AGFA–TL as IGFA. In other words, once a detection is made, there is a high probability that the algorithm will generate a forecast. On the other hand, AGFA–TL is more likely to forecast the portion of the gust front that impacts the airport since it detects a larger portion of the gust front length.

Comparisons of the IGFA and AGFA–TL gust front algorithms for both detections and forecasts show that the length–based scheme, which more closely relates to the operational usage of the product, shows that AGFA–TL provides significant improvement over IGFA. As improvements to the algorithm increase the length detected, the probability of detecting any part of a gust front increases. This increases the likelihood that the wind shear warnings will be appropriately issued.

Length–based scoring may not be appropriate for assessing the operational usefulness of the forecasts. ATC controllers and supervisors have indicated that the forecasts are very useful for planning and coordinating runway changes, even though the automated scoring results do not necessarily support this assessment. This reason for the discrepancy is that automated scoring does not reflect how the product is used. However, automated scoring does provide a framework within which changes to the algorithm easily can be assessed.

LIST OF ABBREVIATIONS

AGFA	Advanced Gust Front Algorithm
AGFA–TL	Advanced Gust Front Algorithm using Thin Line only
ATC	Air Traffic Control
CFP	Correct Forecast Probability
FAA	Federal Aviation Administration
FFP	False Forecast Probability
IGFA	Improved Gust Front Algorithm
PCF	Probability of a Correct Forecast
PF	Probability of generating a Forecast
PFA	Probability of False Alarm
PFD	Probability of False Detection
PFF	Probability of a False Forecast
PLD	Probability of Local Detection
POD	Probability of Detection
TDWR	Terminal Doppler Weather Radar
%L	Percent of Length Detected

REFERENCES

D. M. Bernella, "Terminal Doppler weather radar operational test and evaluation Orlando 1990," MIT Lincoln Laboratory, Lexington, MA, ATC–179 DOT/FAA/NR–91/2, (April 1991).

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," *Preprints of the Fourth International Conference on the Aviation Weather System,* Paris, France, June 1991, American Meteorological Society, Boston, MA, (1991) pp. J37–J42.

J. Evans and D. Turnbull, "Development of an automated windshear detection system using Doppler weather radar," *Proceedings of the IEEE,* vol. 77(11), (1989) pp. 1661–1673.

S.D. Smith, "Sectorized Uniform Wind Algorithms,: NEXRAD Joint Systems Program Office Report, National Severe Storms Laboratory (1986).

S.D. Smith, A. Witt, M. Eilts, D. Klingle–Wilson, S. Olson, J. Sanford, "Gust front detection algorithm for the terminal Doppler weather radar: part I, current status," *Preprints of Third International Conference on the Aviation Weather System,* Anaheim, CA, January 1989, American Meteorological Society, Boston, MA, (1989) pp. 31–34.

H. Uyeda and D.S. Zrnic, "Automatic Detection of Gust Fronts," J. Atmos. and Oceanic Tech. 3, (1986) pp. 36–50.

APPENDIX A: ALGORITHM STATISTICS

Table A–1.
Algorithm Parameters

PARAMETER	VALUE for IGFA	VALUE for AGFA–TL
Velocity Difference: Lower Tilt	7.0 m/s	7.0 m/s
Velocity Difference: Upper Tilt	5.0 m/s	5.0 m/s
Peak Shear: Lower Tilt	2.0 m/s	2.0 m/s
Peak Shear: Upper Tilt	2.0 m/s	2.0 m/s
Azimuth Overlap	3.3°	3.3°
Range Overlap	2.0 km	2.0 km
Number of Segments	5	5
Feature Length	5.0 km	5.0 km
Feature Distance	5.0 km	5.0 km
Combined Length	5.0 km	5.0 km
Front Length	10.0 km	10.0 km
Auto–plot Length	15.0 km	15.0 km
Overhead Tracking Range Threshold	10.0 km	10.0 km
Overhead Tracking Speed Threshold	4.0 m/s	4.0 m/s
Number of Coasts Threshold	12	10

Table A–2.
Base data for POD and PLD statistics for all gust front strengths
for the IGFA and AGFA–TL algorithms.

	POD	PLD
IGFA		
Weak	351/1191 (0.30)	5695/37135 (0.15)
Moderate	444/601 (0.74)	12315/30574 (0.40)
Strong	205/237 (0.87)	6322/14156 (0.45)
Severe	22/22 (1.00)	678/1165 (0.58)
All Strengths	1022/2051 (0.50)	25010/83030 (0.30)
AGFA–TL		
Weak	452/1191 (0.38)	9268/37135 (0.25)
Moderate	473/601 (0.79)	16418/30574 (0.54)
Strong	212/237 (0.90)	7685/14156 (0.54)
Severe	22/22 (1.00)	700/1165 (0.60)
All Strengths	1159/2051 (0.57)	34071/83030 (0.41)

Table A–3.
Base data for PLD and PFD statistics for all gust front strengths and locations for the IGFA and AGFA–TL algorithms.

	Weak	Moderate	Strong	Severe	All Strengths	
	PLD					PFD
IGFA						
Denver	3272/21712 (0.15)	6713/15532 (0.43)	2396/5020 (0.48)	19/31 (0.61)	12400/42295 (0.29)	1353/12458 (0.11)
Kansas City	1100/5145 (0.21)	2485/6410 (0.39)	2986/6659 (0.45)	628/1008 (0.62)	7199/19222 (0.37)	2910/10638 (0.27)
Orlando	1324/10278 (0.13)	3117/8632 (0.36)	940/2477 (0.38)	31/126 (0.25)	5412/21513 (0.25)	591/6004 (0.10)
All Locations	5695/37135 (0.15)	12315/30574 (0.40)	6322/14156 (0.45)	678/1165 (0.58)	25010/83030 (0.30)	4964/30089 (0.17)
AGFA–TL						
Denver	5318/21712 (0.24)	8535/15532 (0.55)	3002/5020 (0.60)	20/31 (0.64)	16875/42295 (0.40)	4121/21271 (0.19)
Kansas City	1297/5145 (0.25)	2983/6410 (0.46)	3371/6659 (0.51)	620/1008 (0.62)	8271/19222 (0.43)	3347/12334 (0.27)
Orlando	2653/10278 (0.26)	4900/8632 (0.57)	1312/2477 (0.53)	60/126 (0.48)	8925/21513 (0.41)	1902/12214 (0.16)
All Locations	9268/37135 (0.25)	16418/30574 (0.54)	7685/14156 (0.54)	700/1165 (0.60)	34071/83030 (0.41)	9370/45819 (0.21)

Table A–4.
Base data for Over–the–Airport PLD and PFD statistics for all gust front strengths and locations for the IGFA and AGFA–TL algorithms.

	Weak	Moderate	Strong	Severe	All Strengths	
	PLD					PFD
IGFA						
Denver	2338/9158 (0.26)	3992/7361 (0.54)	1224/2000 (0.61)	19/31 (0.61)	7573/18550 (0.41)	1189/9211 (0.13)
Kansas City	276/1335 (0.20)	807/1537 (0.52)	1503/2747 (0.55)	83/101 (0.83)	2669/5720 (0.46)	1806/4749 (0.38)
Orlando	867/4525 (0.19)	1836/3787 (0.49)	503/927 (0.54)	31/80 (0.39)	3237/9319 (0.35)	499/4024 (0.12)
All Locations	3480/15018 (0.23)	6636/12685 (0.52)	3230/5674 (0.57)	133/212 (0.63)	13479/33589 (0.40)	3489/17978 (0.19)
AGFA–TL						
Denver	3267/9158 (0.35)	4546/7361 (0.62)	1384/2000 (0.69)	20/31 (0.64)	9217/18550 (0.49)	2124/12494 (0.17)
Kansas City	335/1339 (0.24)	995/1652 (0.60)	1849/3051 (0.61)	84/101 (0.84)	3263/6143 (0.53)	1995/5694 (0.35)
Orlando	1318/4525 (0.29)	2563/3787 (0.68)	666/927 (0.72)	51/80 (0.64)	4598/9319 (0.49)	1294/6754 (0.19)
All Locations	4919/15018 (0.33)	8062/12685 (0.64)	3696/5674 (0.65)	155/212 (0.73)	16832/33589 (0.50)	5196/24423 (0.21)

Table A–5.
Base data for PF|d from the IGFA and AGFA algorithms.

	Correct Forecasts				False Forecasts	All Forecasts
	Weak	Moderate	Strong	Severe		
IGFA						
Denver	144/204 (0.71)	147/227 (0.65)	60/81 (0.74)	1/1 (1.00)	38/513 (0.07)	390/513 (0.76)
Kansas City	31/77 (0.40)	49/101 (0.49)	63/88 (0.72)	12/19 (0.63)	86/285 (0.30)	241/285 (0.85)
Orlando	46/70 (0.66)	90/116 (0.78)	27/36 (0.75)	2/2 (1.00)	26/224 (0.12)	191/224 (0.85)
All Locations	221/351 (0.63)	286/444 (0.64)	150/205 (0.73)	15/22 (0.68)	151/1022 (0.15)	823/1022 (0.81)
AGFA–TL						
Denver	170/260 (0.65)	168/240 (0.70)	66/84 (0.79)	1/1 (1.00)	33/585 (0.06)	438/585 (0.75)
Kansas City	34/84 (0.40)	58/105 (0.55)	73/92 (0.79)	14/19 (0.74)	69/300 (0.23)	248/300 (0.83)
Orlando	75/108 (0.69)	99/128 (0.77)	29/36 (0.81)	2/2 (1.00)	27/274 (0.10)	232/274 (0.85)
All Locations	279/452 (0.62)	325/473 (0.69)	168/212 (0.79)	17/22 (0.77)	129/1159 (0.11)	918/1159 (0.79)

Table A–6.
Base data for PF|e from the IGFA and AGFA algorithms.

	Correct Forecasts				False Forecasts	All Forecasts
	Weak	Moderate	Strong	Severe		
IGFA						
Denver	144/659 (0.22)	147/302 (0.49)	60/88 (0.68)	1/1 (1.00)	30/1050 (0.07)	390/1050 (0.37)
Kansas City	21/227 (0.14)	49/139 (0.35)	63/111 (0.57)	12/19 (0.63)	86/496 (0.17)	241/496 (0.49)
Orlando	46/305 (0.15)	90/160 (0.56)	27/38 (0.71)	2/2 (1.00)	26/505 (0.05)	191/505 (0.38)
All Locations	221/1191 (0.19)	286/601 (0.48)	150/237 (0.63)	15/22 (0.68)	151/2051 (0.07)	823/2051 (0.40)
AGFA–TL						
Denver	170/659 (0.26)	168/302 (0.56)	66/88 (0.75)	1/1 (1.00)	33/1050 (0.03)	438/1050 (0.42)
Kansas City	34/227 (0.15)	58/139 (0.42)	73/111 (0.66)	14/19 (0.74)	69/496 (0.14)	248/496 (0.50)
Orlando	75/305 (0.25)	99/160 (0.62)	29/38 (0.76)	2/2 (1.00)	27/505 (0.05)	232/505 (0.46)
All Locations	279/1191 (0.23)	325/601 (0.54)	168/237 (0.71)	17/22 (0.77)	129/2051 (0.06)	918/2051 (0.45)

Table A–7.
Base data for CFP and FFP statistics for the 10–minute forecasts for the IGFA and AGFA–TL algorithms.

	Weak	Moderate	Strong	Severe	All Strengths	
	CFP					FFP
IGFA						
Denver	2099/20883 (0.10)	2856/15115 (0.19)	1188/4807 (0.25)	1/31 (0.03)	6144/40836 (0.15)	3205/9416 (0.34)
Kansas City	288/4757 (0.06)	1249/6073 (0.20)	1791/6134 (0.29)	277/1008 (0.28)	3605/17972 (0.20)	3271/7139 (0.46)
Orlando	845/9893 (0.09)	2243/8590 (0.26)	508/2394 (0.21)	57/126 (0.45)	3653/21003 (0.17)	1317/5079 (0.26)
All Locations	3232/35533 (0.09)	6344/29778 (0.21)	3487/13335 (0.26)	335/1165 (0.29)	13398/79811 (0.17)	8129/22401 (0.36)
AGFA–TL						
Denver	3326/20883 (0.16)	5661/15115 (0.37)	1790/4807 (0.37)	1/31 (0.03)	10778/40836 (0.26)	6200/16694 (0.37)
Kansas City	379/4757 (0.08)	1536/6073 (0.25)	2358/6134 (0.39)	296/1008 (0.30)	4569/17972 (0.25)	3515/8477 (0.42)
Orlando	1670/9893 (0.17)	3521/8590 (0.41)	965/2394 (0.40)	84/126 (0.67)	6240/21003 (0.30)	3120/10226 (0.30)
All Locations	5375/35533 (0.15)	10718/29778 (0.36)	5113/13335 (0.38)	381/1165 (0.33)	21587/79811 (0.27)	12835/35437 (0.36)

Table A–8.
Base data for CFP and FFP statistics for the 20–minute forecasts for the IGFA and AGFA–TL algorithms.

	Weak	Moderate	Strong	Severe	All Strengths	
	CFP					FFP
IGFA						
Denver	1753/20558 (0.09)	3169/15362 (0.20)	706/4792 (0.15)	0/31 (0.0)	5628/40743 (0.14)	4971/9365 (0.53)
Kansas City	115/4560 (0.03)	802/5722 (0.14)	1319/6259 (0.21)	138/976 (0.14)	2374/17517 (0.13)	4314/6893 (0.63)
Orlando	660/9985 (0.07)	1609/8419 (0.19)	356/2394 (0.15)	42/126 (0.34)	2667/20924 (0.13)	2129/4906 (0.43)
All Locations	2527/35103 (0.07)	5579/29508 (0.19)	2381/13445 (0.18)	180/1133 (0.16)	10667/79184 (0.13)	11914/21950 (0.54)
AGFA–TL						
Denver	2534/20558 (0.12)	4519/15362 (0.29)	1161/4792 (0.24)	0/31 (0.0)	8214/40743 (0.20)	9031/16671 (0.54)
Kansas City	182/4560 (0.04)	1021/5722 (0.18)	1632/6259 (0.26)	150/976 (0.15)	2985/17517 (0.17)	4750/8024 (0.59)
Orlando	1353/9985 (0.13)	2491/8419 (0.29)	542/2394 (0.23)	51/126 (0.40)	4437/20924 (0.21)	4734/9913 (0.48)
All Locations	4069/35103 (0.11)	8031/29503 (0.27)	3335/13445 (0.25)	201/1133 (0.18)	15636/79184 (0.20)	18515/34608 (0.54)