

---

# Automated Storm Tracking for Terminal Air Traffic Control

Edward S. Chornoboy, Anne M. Matlin, and John P. Morgan

■ Good estimates of storm motion are essential to improved air traffic control operations during times of inclement weather. Automating such a service is a challenge, however, because meteorological phenomena exist as complex distributed systems that exhibit motion across a wide spectrum of scales. Even when viewed from a fixed perspective, these evolving dynamic systems can test the extent of our definition of motion, as well as any attempt at automated tracking of this motion. Image-based motion detection and processing appear to provide the best route toward robust performance of an automated tracking system.

ON APRIL 14, 1993, AMERICAN AIRLINES FLIGHT 102 was unable to hold the runway while landing at Dallas–Fort Worth International Airport. In the resulting accident there were many injuries—two of them serious—and the plane (a DC-10) was irreparably damaged. It was raining at the airport that morning, and numerous thunderstorms were occurring throughout the area. The darkness of the early hour, the fatigue of the flight crew after an all-night flight, and the bad weather were all suspected causes of the accident. Although the National Transportation Safety Board officially concluded that the stormy weather was not a contributing factor to the crash (despite high cross winds from a severe storm passing over the airport, the aircraft was able to touch down on the runway [1]), the weather clearly did play an important role in the events of that day.

The crew of Flight 102 had access to a variety of weather information that morning, including their own radar. Their information sources included American Airlines flight dispatch, Fort Worth Air Route Traffic Control Center (ARTCC), Dallas–Fort Worth approach control, and the Automatic Terminal Information Service (ATIS). To these sources were added the verbal information that came from controllers observing their radar scopes, as well as information reported by other pilots (known as *pireps*). What is interesting about this particular accident is that the

official transcripts of the pilot-controller conversation before the landing show that both the pilot and the controller were aware of the severe weather conditions, yet neither could anticipate the extreme dynamics of the weather. Table 1 gives the transcript of the conversation between the pilot and controller approximately ten minutes before the aircraft landed. Neither of them knew of new storm cells intensifying in the region, nor were they aware of the forty-knot motion that would rapidly change route conditions by placing such dramatic weather in the path of the plane. The four radar images in Figure 1 illustrate the rapid development of the severe weather conditions during a thirty-minute period before and during the landing of Flight 102.

## A High-Payoff Service

Understanding the relationship between weather evolution and cases such as Flight 102 can lead to greater flight safety during arrival and departure. The evidence is already clear that terminal-area storm tracking will lead to a real and tangible payoff for air traffic control, namely, the improved management of foul-weather resources, which is easily linked to significant cost savings. Controllers and controller supervisors—those who oversee and specify aircraft spacing—need to plan ahead to maintain a steady (and maximal) arrival stream of aircraft. Abrupt unanticipated chang-

**Table 1. Pilot-Controller Excerpts from Flight 102 on April 14, 1993**

"On April 14, 1993, about 0659:43 Central Daylight time, American Airlines Flight 102 (AAL102) . . . departed the right side of Runway 17 left, following landing at Dallas-Fort Worth International Airport, Texas. . . ." [1]

In general, it is the controller's role to advise the pilot of the thunderstorm situation and, if possible, suggest a flight path relative to the weather. The pilot makes the final decision on whether to follow the suggested path, using all the information available to him or her [2]. Here is a transcript of the conversation between the captain of Flight 102 and the Dallas-Fort Worth tower controller.

(captain) . . . "One one zero three OK uh, how's it look on your scope for gettin' in there?" [0649:34]

(controller) "Well uh, I show an area of weather at, at fifteen miles either side of DFW airport proceeding uh, straight north uh fifteen miles on uh, each side uh, for about thirty miles."

(captain) . . . "OK uh, and is this stuff moving?" [0650:18]

(controller) "Uh, does not appear to be moving uh, much if any . . . turn right . . . and join the runway one seven left localizer."

(captain) "Uh, I don't think we're goin' to be able to do that that's uh, that's a pretty big red area on our scope . . . that's about what we're looking at. Uh, we're gonna have to, just go out I guess and wait around to see what's goin' on here."

(controller) ". . . Eight miles south of you's a heavy DC-8 at three thousand joining uh, the final's uh, reporting a smooth ride at three." [0650:46]

(captain) "Oh, OK, eight miles south of us? . . . OK uh, we'll head down that way then. . . ." [0650:59]

es, such as the onset of bad weather or the loss of an arrival gate, for example, result in unpleasant and costly delays while planes are rerouted and arrival queues are reestablished.

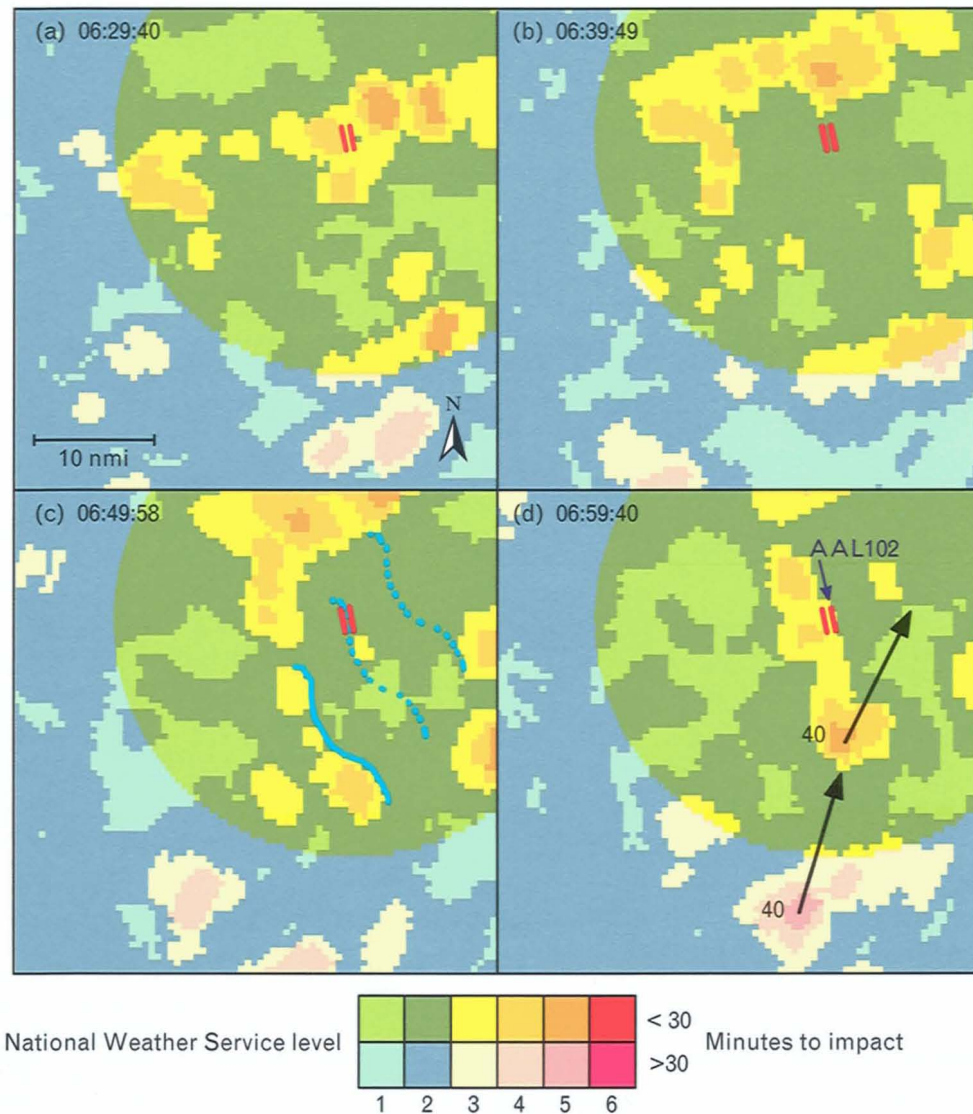
At Lincoln Laboratory, our accumulated experience in the development of the Terminal Doppler Weather Radar (TDWR) [3, 4], the Airport Surveillance Radar Wind Shear Processor (ASR-WSP) [5], and the Integrated Terminal Weather System (ITWS) [6] confirms the utility of (and the controllers' desire for) basic services such as precipitation mapping and storm-motion tracking. With regard to the day-to-day execution of the controller's duties, these two services produce significant payoffs in the effort to reduce weather-related delay [7]. Under ITWS, the importance of storm-motion tracking can be underscored still further because this tracking is used in turn by many other ITWS product algorithms, such as storm-cell information, microburst prediction, anomalous propagation editing, and pilot data link. For more information on ITWS products, see the article by James Evans and Elizabeth Ducot, entitled

"The Integrated Terminal Weather System (ITWS)," and the article by Marilyn Wolfson et al., entitled "Automated Microburst Wind-Shear Prediction," in this issue.

### **Air Traffic Control Products**

The weather-information needs of controllers can be simply stated: they want someone to provide them with unambiguous weather products that are free of any need for interpretation and coordination and that expedite air traffic control decision making [8]. Generally speaking, controllers and pilots need to know where bad weather and its hazards will be and when these hazards will occur; that is, they require predictions that, among other things, account for storm motion. Predicting this motion is a difficult task because, as Figure 1 illustrates, weather events are complex phenomena. Figure 1(a) shows weather cells thirty minutes before Flight 102 landed, and Figure 1(b) shows weather cells twenty minutes before the landing; both of these figures give little apparent evidence of the storm to come. In Figure 1(c), which occurs ten





**FIGURE 1.** Thirty-minute perspective on storm growth, decay, and motion. These radar images span the thirty-minute interval prior to the landing of American Airlines flight 102 (AAL102) at Dallas–Fort Worth Airport on 14 April 1993 (compare to Table 1 time stamps). The images are of weather-radar reflectivity, displayed with the National Weather Service (NWS) six-level scale with one modification: two different colors are used for each NWS level to distinguish weather that, given the prevailing motion, is within thirty minutes of the airport (red icon). (a) At 6:29:40, AAL102 has begun its preparations for landing, making initial contact with the Fort Worth ARTCC. (b) At 06:37:05, AAL102 is still some eighty miles from the airport. (c) At about 06:50:00 the captain asks the DFW controller if the weather is moving; he also expresses doubt about making their way through (he needs to get by at least one storm region in his immediate view). The controller reports no apparent motion and reassures AAL102 that the DC-8 ahead of them is experiencing a “smooth ride”; given this information, the captain decides to continue his approach. (d) At the touchdown time of 06:59:43, storm cells sweep the runway, and—for whatever reason—the pilot cannot control the aircraft as it weathervanes in heightened winds from the southwest. Under ITWS, there are two display products that could have prevented this accident: the Storm Motion Product, which consists of vectors illustrating direction with an accompanying speed report in knots (part d shows a 40-knot NE motion) and the Storm Extrapolated Position Product, which provides direct visual cues in the form of extrapolated leading-edge contours (part c illustrates the composition of zero-reference, ten-minute, and twenty-minute expected-position contours, which together warn of a potential runway impact in ten minutes). If operational, these products would have provided an easily relayed answer regarding the storm’s movement.

minutes before the landing (and during the time of the conversation in Table 1), the cells are just intensifying. Finally, Figure 1(d) shows a mature group of storm cells sweeping the airport runways with forty-knot winds at the time Flight 102 lands. Often the spatial translation of storms will suffice for weather prediction, but translation alone will certainly not suffice when storms evolve. During the approximately thirty minutes a plane is under terminal airspace control, the morphology of a storm (i.e., its growth and decay) can change considerably.

Explicit prediction maps are beyond current modeling capabilities. In our work for the Federal Aviation Administration (FAA), we have focused on the graphical display of frequently updated precipitation maps augmented by specific weather products that warn of potential hazards or otherwise provide notification of important weather attributes such as storm motion. Figure 1 illustrates two concepts for the portrayal of storm motion. Figure 1(d) shows the Storm Motion product and Figure 1(c) shows the Storm Extrapolated Position product; both of these are scheduled components of ITWS. These particular overlays were not available at the time of Flight 102 but were sampled products from a prototype algorithm running in playback mode. Both of these simple presentations provide a controller with easy-to-relay information about storm motion. When the pilot of Flight 102 asks the controller, "Is this stuff moving," displays like those shown in Figures 1(c) and 1(d) would remind the controller that a forty-knot northeasterly motion is in the area, and that a storm will probably impact the airport within ten minutes.

### Overview of Storm-Motion Tracking

The development of our algorithm for storm-motion tracking actually has two objectives. First, as presented above, we want to provide reliable motion information in a form accessible to controllers. Second, and equally important, we want to support ITWS algorithms that need to compensate for the relative movement of data between samplings. Hence our interest in storm motion covers territory that is somewhat general. Many issues specific to our two objectives cannot be covered here; these issues are discussed elsewhere [9, 10]. We have chosen to accomplish

both of the stated objectives by using a motion-estimation strategy based on *local-area correlation*, which is the matching of regionalized image patterns from one frame to the next. The purpose of this article, then, is to present a rationale for our decision to use local-area correlation for storm-motion tracking, and to describe our adaptation of the local-area correlation method.

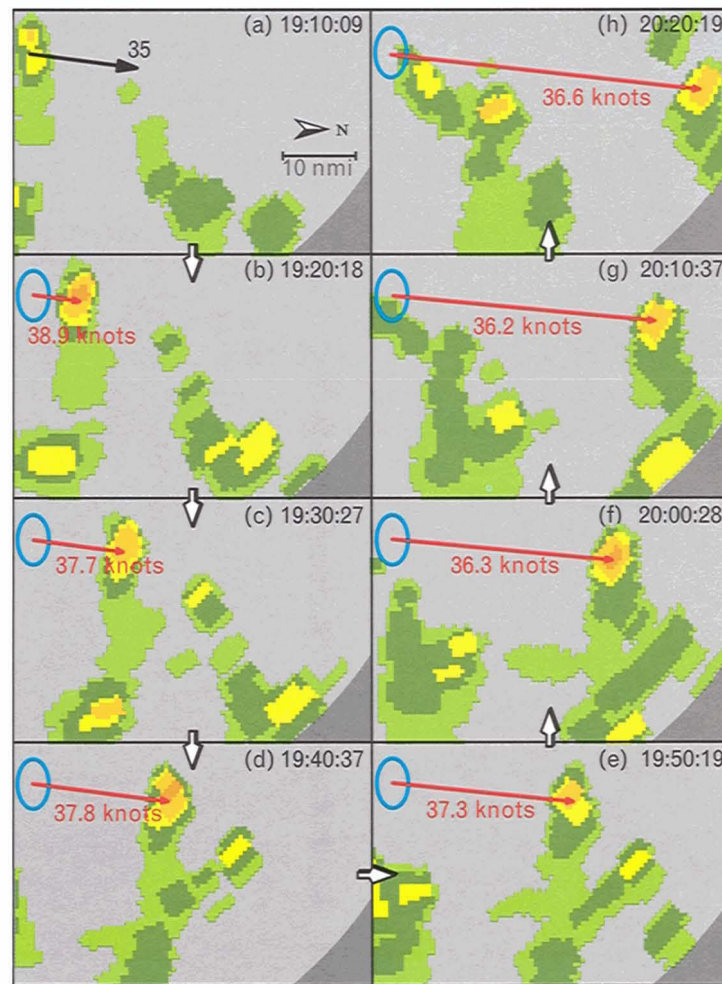
Estimating storm motion by tracking regional image patterns is a well-known procedure, with little change in approach since the 1970s. Attempts to augment and improve the method have been directed more toward the context of machine vision and deterministic (rigid body) estimation. To our knowledge, no study has focused on developing this tracking method by including evolving and varying data such as we observe in storm motion. With weather-radar data, the quality of correlation matching can be quite variable. Given that our objectives in developing a storm-tracking algorithm call for reliable, robust, and autonomous performance, there is a real need for quality-control mechanisms that can incorporate constraints and deal with outlying data points, or *outliers*, that affect the quality of motion prediction.

This article is divided into three parts. The first part provides some historical background and considers the philosophical pros and cons of correlation matching versus centroid tracking (which is the most likely alternative tracking method available for use by air traffic control). The second part presents a brief synopsis of new correlation ideas that have been proposed in the context of generic image analysis. The third part proposes an alternative approach that places the correlation technique in the context of an image-reconstruction problem. In particular, we present a linear-estimation framework that is designed to incorporate phenomenological constraints as well as outlier handling.

### Historical Background

Automated storm trackers have aligned themselves (more or less) along one or the other of two philosophical paths, which we refer to here as the *centroid method* and the *correlation method*. In the Laboratory's work for the FAA we have had the opportunity to examine both of these methods carefully [11, 12, 9].





**FIGURE 2.** When thunderstorm cells are isolated and long lived, they are excellent candidates for object-oriented centroid tracking. This sequence from Dallas–Fort Worth on 13 September 1993 shows a relatively long-lived thunderstorm cell during a seventy-minute period. (a) The output of our automated algorithm is rounded to the nearest five-knot increment. (b)–(h) Each panel recomputes the speed and direction of the thunderstorm cell relative to the displacement from part a.

Superficially, these philosophies can be contrasted as object identification with tracking versus flow-field estimation. These two methods also represent opposing positions in estimation policy; one method depends on obtaining a small number of accurate motion measurements while, in theory, the other method depends on obtaining many motion measurements with a low overall measurement error as the only goal.

#### *Centroid Tracking*

Centroid, or *cell*, tracking methods are a high-level approach to predicting storm motion because the ra-

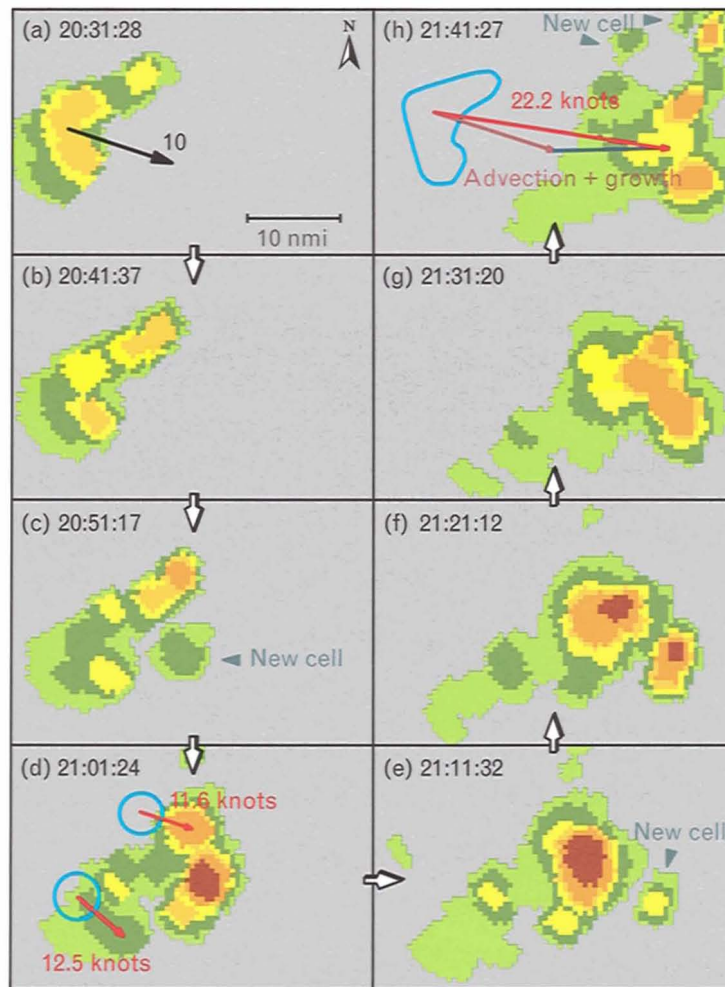
dar data are first reduced by meteorological analysis to a small number of large-scale features, which are then used as tracers of mesoscale movement [13–15]. Thunderstorms provide the best and most important example of cell tracking because they are typically viewed as aggregates of component cells that, in this case, correspond to regions of strong (localized) vertical motion [16, 17]. Because cells are a useful organizing principle, and because they can be identified and associated between successive radar scans as regions of heightened radar reflectivity, it is natural to track thunderstorms by tracking component cells. Significant hazards to aircraft are associated with thunder-

storm cells as well, which alone motivates their tracking. In “well-behaved” storms, centroid-based cell tracking and correlation-based storm tracking are often interchangeable and can be easily substituted for each other. Figure 2 shows a single-cell storm whose motion is indicative of all storm motion in its local region.

Obvious problems occur in centroid methods, however, with the subordination of motion estimation to the identification of amorphous and evolving objects such as thunderstorms. Cells, for example, being dynamic, commonly merge and split. Hence a well-known concern with object-oriented cell-centroid tracking is that the incorrect identification or

changing character of an object propagates into tracking errors, often with catastrophic results. Centroid methods for motion tracking are often confused by highly evolving conditions such as squall-line thunderstorms, when cells are poorly isolated.

Refinements in the definition of a cell and the introduction of adaptive techniques can mitigate the above problem, but the prevention of gross identification errors is far from perfect and not the only concern in the development of a centroid-tracking algorithm. Because the data are first reduced to a small number of tracers (i.e., features that exist from one frame to the next), the quality of the derived motion information relies heavily on only these few measure-



**FIGURE 3.** Organized growth and decay of storm cells can give rise to an apparent storm motion, as shown in this sequence from Orlando on 13 July 1993. (a) An instantaneous velocity estimate is output by our storm-tracking algorithm. (b)–(h) The organized appearance of new cells more than doubles the apparent motion of this initial storm-cell complex over a seventy-minute period.




ments. Hence when catastrophic errors occur, recovery is difficult because each error carries significant weight within the measurement population.

Thunderstorm cells are neither closed nor rigid; thus with typical radar-data analysis, arbitrary reflectivity thresholding redefines a cell at each analysis instant. Because the centroid method usually doesn't track a cell *per se* but an approximating feature such as a mass-weighted centroid, the concept of *centroid shift* enters as a constant source of error that limits tracking accuracy. In our experience, centroid methods are inadequate for the accurate alignment of data fields undergoing motion.

#### *Growth and Decay*

Cells exhibit life cycles with characteristic stages [18], and a typical cell does not last longer than one hour (usually less). At any given time a thunderstorm consists of a succession of cells in different stages of evolution. Many factors affect new cell generation, but it is not uncommon to see new cells appear with life cycles on the order of fifteen to twenty minutes (a time scale obviously relevant to air traffic control decision making). New cells often form next to old cells, but not necessarily along the direction of motion of the old cell. Organized growth and decay can therefore result in an apparent storm motion that veers from cellular advection (the horizontal motion due to air currents). Figure 3 illustrates just such a case; the organized growth of new cells doubles the apparent motion of the storm as a whole. In situations of severe growth and decay, what should the analysis method report—the motion of the cell, or the apparent motion of the storm, or both? This question is not an easy one because the answer is predicated on the time frame of the user's need. In the case of Figure 3, the advective component of the cell dominates for time frames of twenty minutes or less (as in predicting part *c* from part *a*) but predicting part *b* from part *a* requires the modeling of a growth-induced jump discontinuity.

As a technical challenge, could we control the sensitivity of the processing to either the motion of the cell or the motion of the storm? That is, could we reliably dissect the two forms of motion apart? In this respect, our own experimentation with tracking the



**FIGURE 4.** Organized growth and decay literally move a storm by leaps and bounds. This motion is in contrast to the more or less continuous (although by no means constant) motion effected by wind advection. The motion fields illustrated here are those of (a) a constant advective component, (b) an instantaneous growth perturbation, and (c) their sum. The obvious decomposition seen here becomes more difficult when the measurement field is corrupted by noise.

envelope of a storm (ignoring the cellular content) has been only partially successful, because it represents a compromise position between cell motion versus storm motion [9]. Nevertheless, we have observed empirically that there is at least a qualitative difference between the two motions. The translational motion of storms by steering winds is more or less continuous in nature. Growth and decay, on the other hand, are manifest as discrete perturbations of brief duration. The conceptual rule is illustrated in the sketch of Figure 4; the top row shows three successive time frames from a constant flow field, the middle row illustrates a discrete perturbation in the middle frame, and the bottom row shows their superposition. Unfortunately, the interval between growth perturba-



## MOTION ANALYSIS BY CORRELATION

LOCAL-AREA CORRELATION provides a measure of the *apparent* motion in an image sequence. It usually involves an analysis applied between two time frames, as shown in Figure A, but local-area correlation across multiple time frames is certainly feasible.

Figure A shows a windowed analysis, with its accompanying bounds on resolution. A square image mask is centered over an arbitrary pixel in one image to define a local subimage, or *correlation box*. A neighborhood of the second image is searched to find an appropriate match to the picture content of the subimage. In practice, the search is centered on zero displacement (or biased toward a prior expectation) and bounded by an excursion limit.

There are many ways to measure the goodness of match, but the more common methods are based on evaluation of an image cross-correlation. The cross-correlation computation can be normalized to desensitize the method with respect to mean-bias or scale

fluctuations. Mean-squared error also can be used, and is gaining in popularity. Some comparisons of these methods have been documented elsewhere [1].

In Figure A, the analysis is mean and variance normalized to yield a set of correlation coefficients bound between +1 and -1. If an excursion limit of  $M$  pixels is used, where  $M$  is odd and the search for best match looks no farther than  $(M-1)/2$  pixels in any direction, then the correlation analysis generates an  $M \times M$  surface of correlation values. The mask can be centered, in turn, over each pixel in the first image, providing a pixel-level analysis of the original image pair. Since each pixel gives rise to an  $M \times M$  correlation surface, a complete analysis can produce a meta-image of  $M \times M$  correlation surfaces ( $M^2$ -length vectors) indexed over the weather-image index set.

In practice, of course, the analysis is usually limited to a sampling of locations based on image content, or a sublattice of the im-

age index set. Furthermore, few analysis methods bother with the correlation surface itself and seek only the displacement that yields the maximum correlation (MAXCOR), which is then accepted as *the* displacement measurement. Previous attempts to recover a displacement-vector field have taken these MAXCOR displacement measurements as the starting point for reconstruction and treated them as uniform-quality measurements. We have been exploring alternative methods to improve the quality of motion information. Just as an intensity (scalar) image can be viewed as incomplete or degraded or both, so can a correlation-surface meta-image. We are currently interested in research that centers on reconstruction of the correlation-surface meta-image.

### References

1. P.J. Burt, C. Yen, and X. Xu, "Local Correlation Measures for Motion Analysis: A Comparative Study," *IEEE Proc. Pattern Recognition and Image Processing, Las Vegas, NV, 14-17 June 1982*, p. 269.

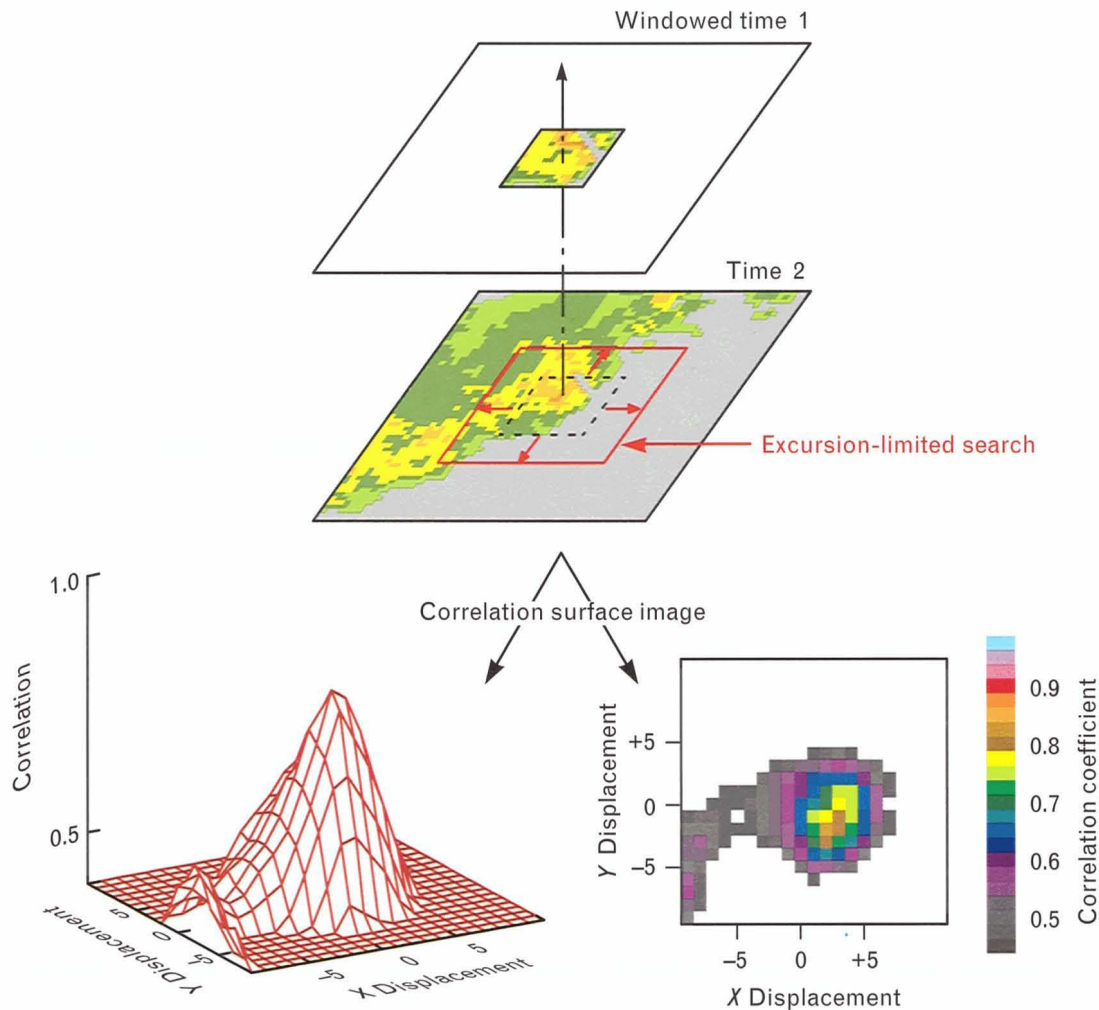
tions is not fixed *a priori*, nor is growth guaranteed to be continual. Some aspects of growth and decay may be predictable, given a history of observation (a current point of focus), but better meteorological analysis and modeling are needed here.

### Correlation Tracking

Area correlation has been applied to weather-radar

data in the past, although generally not in automated settings but in the measurement of global (whole radar field) displacements [19], regional (whole echo) movement [20], and even local (internal) motion [21]. As stated earlier, the technique of basic correlation matching is well known. Many of the image-based definitions, ideas, and issues in correlation matching are summarized in more detail elsewhere





**FIGURE A.** Motion analysis by local-area image correlation. A square image mask is centered over an arbitrary pixel in one image to define a local subimage. An excursion-limited neighborhood of a second image at a different time is searched to find an appropriate match to the first image. A surface of correlation values, shown here in both three-dimensional and false-color views, is generated for each pixel in the original image.

[22]. Our particular point of view is outlined in the sidebar entitled “Motion Analysis by Correlation.” This technique fits into the category of tracking methods that correlate regularly spaced local image segments (square regions) to obtain a grid of displacement (or, equivalently, motion) vectors.

With local-area correlation, no particular effort is made to interpret the data for its meteorological in-

formation, and this turns out to be a strong selling point because it imparts the desired immunity to cell splitting/merging. The ability to study storm motion over a range of spatial and temporal scales is also a major plus. Because the correlation method can be applied at the pixel level, it can function as a local gridded analysis (it can provide a “dense” set of measurements). This type of analysis fits in well with our

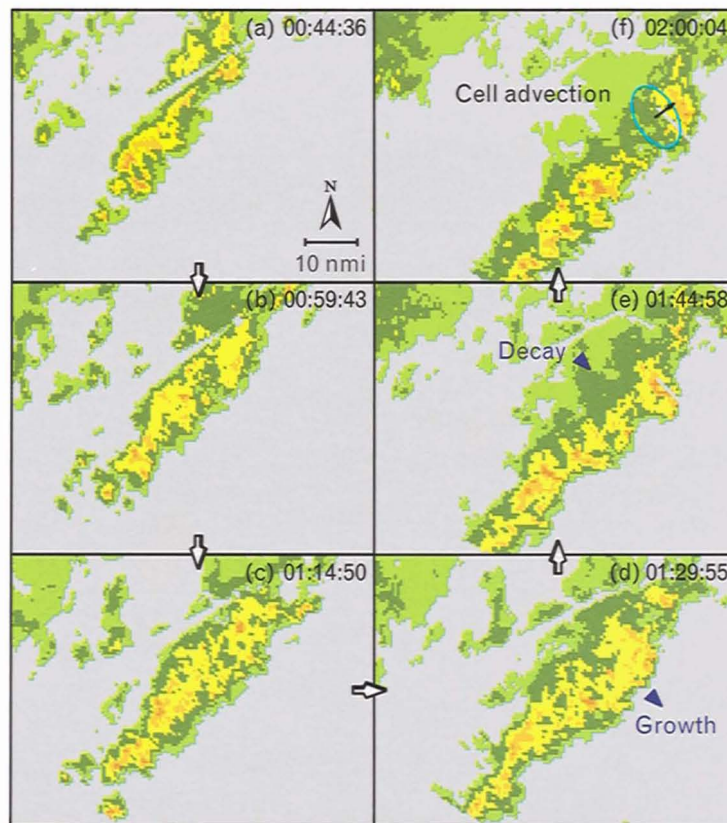
objective to support pixel-based meteorological analysis and parallels the approach used by our group at Lincoln Laboratory to develop successful gust-front [23] and microburst prediction algorithms (see the article “Automated Microburst Wind-Shear Prediction” in this issue). In contrast to the centroid approach, as the ideal argument goes, local-area correlation can provide a number of overlapping and, therefore, corroborating measurements, which (to a degree) relaxes the dependence on the quality of any individual measurement.

Our interest in the local-area correlation method stems from consideration of the virtues of the method stated above and also from the more general notion that correlation tracking can be applied across a wide range of data types and morphologies: thunderstorms, stratiform rain, and snow are all (potentially) covered. Snowstorms, and their tracking, are a major

concern at many U.S. airports, but neither they nor stratiform rain exhibits a cellular structure, and therefore cannot be tracked by centroid methods. Snowstorm tracking is a planned component of the ITWS system, and early results in applying correlation tracking to snow data have been encouraging [24].


### *A Key Challenge*

Tracking an organized linear convective system, or line storm, such as the one shown in Figure 5, is a key challenge because of its extended evolving nature (advective motion and aggressive growth are highly intermixed). Figure 6 illustrates many of the tracking problems in particular. In this figure, raw displacement vectors from a naive analysis (the maximum correlation, or MAXCOR, displacement measurements receive no additional processing) are shown aligned with one of two generating frames occurring



**FIGURE 5.** Organized linear convective systems, or *line storms*, such as this squall-line thunderstorm in Kansas City on 26 June 1989 are a problem for centroid trackers because cellular identities are highly ambiguous. This sequence also illustrates the combined actions of cell advection and organized growth and decay.





**FIGURE 6.** Motion analysis of a convective squall-line thunderstorm recorded near Kansas City on 26 June 1989. Six-level precipitation maps were analyzed by unconstrained local-area correlation. The motion-vector field is  $12 \times 12$ , with a corresponding vector separation of 14 km. Observations are clearly absent where there is little or no stormy weather. Steering winds drive established cells to the northeast while new cells form to the southeast, causing the storm's apparent motion to veer. Older stratiform storm regions to the northwest no longer exhibit a cellular structure. Confusion due to line features is evident as well.


during the time span covered in Figure 5. In both figures, there is stratiform rain to the northwest and convective thunderstorms to the south. Organized growth and decay is also present, and while steering winds drive established cells to the northeast, new cells are added to the southeast, causing the line storm's apparent motion to veer. The *aperture effect*, which is described below, is an obvious problem at the radar-data boundary, but here it is also responsible for large errant vectors (false matches) aligned with the storm's southwest-to-northeast structure. There are suspect vectors aligned northwest to southeast, too—but are these due to an aperture effect or do they indicate a growth-and-decay perturbation (as shown in Figure 3)? Processing with linear and non-linear filters could obscure the fact that here the algorithm is sensing a growth-induced perturbation.

### *The Aperture Effect*

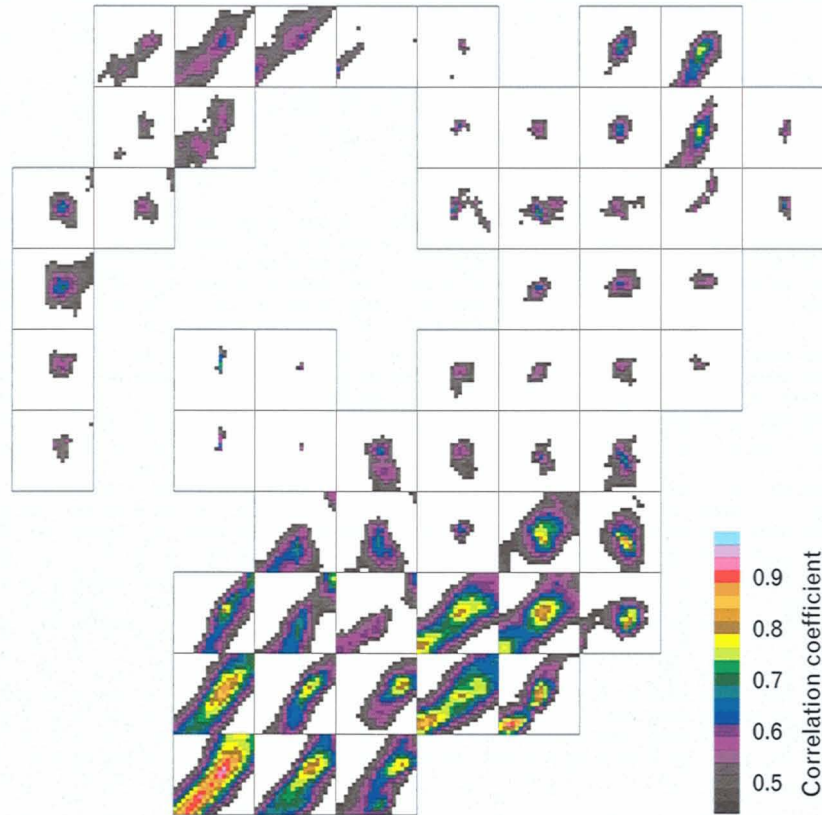
Storm growth and decay aside, local-area correlation still measures an apparent motion because any local measurement can be biased toward a particular component of motion. This happens because the matching of local (i.e., windowed) features is fundamentally limited in that linear features yield motion information only in the direction orthogonal to their edge. Figure 7 shows the degenerate case of a line that extends beyond the view window. This limitation is explicit in gradient-based techniques, such as Optic Flow, because they always track lower-dimensional features [25, 26]. With correlation matching, because of its use of local image *regions*, the problem is more implicit, existing to varying degrees depending on the local image content. The case illustrated in Figure 7, then, is an extreme example of the general problem that motion sensitivity is rarely homogeneous in direction.

### **Augmenting Local-Area Correlation**

Off-line analysis gives us the opportunity to do a thorough local-area correlation analysis centered on each image pixel. This analysis provides a natural “safety in numbers” because it is more likely that the correct displacements will outnumber the outlier displacements, leading to an easy detection of the latter. Because we have real-time processing constraints (our algorithms typically run on workstation platforms,



**FIGURE 7.** The aperture effect. A linear edge (black line) moves with velocity  $\mathbf{V}$  to its new position (blue line). Viewed through the red analysis window, the line appears to move only in the direction orthogonal to the edge.



**FIGURE 8.** Correlation surface (CS) meta-image. These unconstrained (raw) data correspond to the vector displacements shown in Figure 6. Each vector shown in Figure 6 corresponds to its similarly positioned correlation surface here. The vector is the scaled displacement corresponding to the surface maximum. A (0,0) displacement corresponds to the center of each CS image. The color scale is the same as the scale used in Figure A in the sidebar entitled "Motion Analysis by Correlation."

and two or more algorithms often share a single processor), we avoid thorough sampling and settle instead for a subsampling of the image lattice. Our measurement set is dense only in the sense that we have more measurements than would be obtained by centroid analysis. By looking at Figure 6 we can clearly see that filtering the displacement field is problematic; the actual measurements are never fixed in location (they move with the weather) and local ratios of outlier displacements to correct displacements can be large. For our application, the basic correlation mechanism is inadequate. In trial studies, we considered three improvement methods previously (and more than once) suggested in the context of generic motion analysis. These three improvement methods are summarized below.

#### *Outlier Detection and Confidence*

Figure 8 is the correlation surface (CS) meta-image corresponding to the raw vectors of Figure 6. A knowledgeable observer who looks at such a map can infer much about the true displacement field, the respective quality of measurements, and often where correction is needed. A useful local-area correlation algorithm should be capable of making the same kind of inferences. The ranking of displacement measurements (surfaces) for confidence, or at least the flagging of outlier candidates, seems a logical next step, and we considered one idea similar to suggestions found elsewhere [27, 28].

In our trial study we first tested the lesser goal of merely flagging outliers. The untested premise is that



## HIGHER-MOMENT ANALYSIS

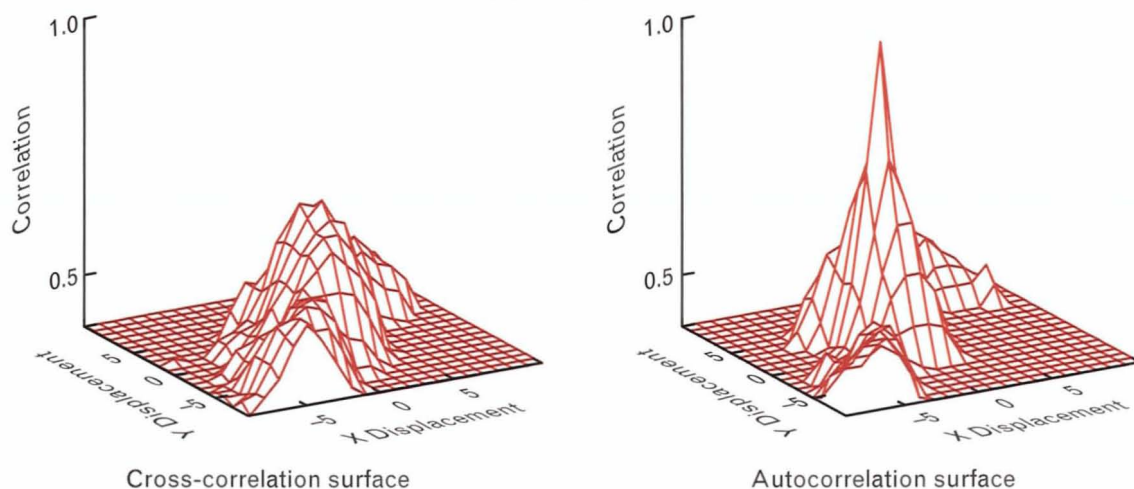
THE SHAPE OF THE cross-correlation surface provides information regarding the quality of the correlation analysis. Some surface properties are derived from the local image structure, as can be seen by comparing autocorrelation and cross-correlation surfaces. Some surface properties result from the evolutionary changes in image structure and can be interpreted as a distortion, or noise, component.

The cross-correlation surface shown in Figure A in the sidebar entitled "Motion Analysis by Correlation" is a good example of an

analysis that yields high-confidence information, namely, a well-isolated peak that slopes sharply and uniformly with direction. Unfortunately, such measurements are more the exception than the norm. The cross-correlation surface shown in Figure A in this sidebar is closer to reality. In this case, the peak is not well defined and the surface shows clear signs of the aperture effect (in other words, motion sensitivity is not uniform in direction).

The autocorrelation surface illustrates the expected morphology for the cross-correlation sur-

face. In this normalized form, the autocorrelation surface is guaranteed to exhibit a maximum peak at zero displacement. The two surfaces side by side evoke the image of a matched-filter strategy for localization (we want to find the amount by which the cross-correlation surface has been displaced from zero). The cross-correlation image can be viewed as a signal plus noise, while the autocorrelation image represents the translated signal component. Hence a correlation analysis of these two images would provide a more robust estimate of the displacement.



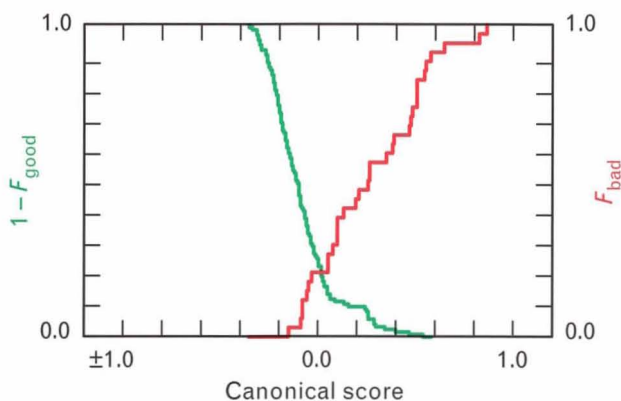
**FIGURE A.** A cross-correlation surface and its corresponding "reference" autocorrelation surface.

there is essentially a class of measurements that are accurate enough (they require little or no correction), and the rest of the measurements require attention and adjustment. We could either censor detected outliers or perhaps appeal to higher-moment correlation,

such as described in the sidebar entitled "Higher-Moment Analysis," potentially recovering these analyses. A higher-moment analysis would be more robust but also computationally expensive, so it would be necessary to limit its use.

Could the qualitative assessment of the CS data be captured by a few easily computed CS statistics? For example, when viewed from above, each correlation surface typically exhibits a major and minor axis (due to the nonhomogeneous sensitivity of the aperture effect). We computed six primary statistics, including surface spread about each axis, surface eccentricity, surface curvature at the MAXCOR value (along major and minor directions), and, of course, the number of modes (peaks). In all, nine characterizing statistics were compiled (the six listed above, plus statistics for CS signal-to-noise, weather stratiformness [i.e., diversity/entropy of intensity values], and a propensity for occlusion [at present, proximity to the data horizon]). Test subjects, working with precipitation-map movies, subjectively scored displacement measurements for consistency (but not accuracy—a vector was either “good” or “bad”).

Of the nine statistics we considered, no one statistic could adequately separate the two subjective populations. By resorting to a canonical analysis (such as given in chapter 12 of Reference 29) and linearly combining all nine measures, we were able to demonstrate some discrimination capability. This finding is summarized in Figure 9, which shows that the ca-



**FIGURE 9.** Canonical discriminant analysis of a two-class subjective labeling of raw displacement vectors. Raw displacement vectors for selected data collected during 1989 Kansas City demonstrations were human (consensus) classified as “good” or “bad.” In the opinion of the scorer, a bad vector is equivalent to censure. Nine CS statistics were combined linearly to form a normalized canonical score in an attempt to match the human discrimination. Empirical distribution functions for the two populations are plotted versus canonical score.

nonical score can be used to censor some measurements. A vertical line through canonical score “0.0” is a good first cut at separating the two classes of measurements. As would be expected, not all consistent displacements were derived from ideal correlation surfaces, and errant vectors occurred for a variety of reasons.

### *Field Constraints*

The assumption of two modes only, of course, is weak. Gross errors aside, the estimate of storm motion in Figure 6 lacks a smoothness expected of this flow field. Because there is too much uncertainty in the outcome of many correlation analyses, constraints that *apply to* and that can be *enforced upon* the velocity field are needed. For example, in wind-field recovery, restrictions are often placed on rotational and divergent components [30, 31]. Unfortunately, that type of decomposition does not work in view of the growth phenomenon illustrated in Figure 4.

Some degree of smoothness is reasonable, and elsewhere it was suggested that a penalty function be used to find a smooth field that best approximates MAXCOR (or similar criterion) measurements [32]. However, just as with similarly suggested projection methods (raw measurements, including noise, are orthogonally projected onto a suitably smooth vector space [33]), this use of a smoothness penalty function is not the same as a correlation analysis subject to a smoothness constraint. In addition, the current state of such constraint methods is that they lead to iterative, not explicit, solutions, which is limiting for our FAA application. In any event, a smoothness penalty function should be weighed against correlation-evidence *strength*, which is a difficult global optimization task.

### *Hierarchical Search*

Hierarchical themes are popular in image processing, and central to many recent treatments of motion analysis. Here, we refer to the advantages, both organizational and computational, of a processing-pyramid approach [34, 35]. Generally speaking, pyramids are regular structures upon which processing is organized into layers of increasing or decreasing resolution, as in a coarse-to-fine strategy.



With motion detection, the most common such hierarchy is one based on image resolution. That is, we construct a pyramid from a sequence of enlarged-pixel images obtained by means of repeated image coarsening [35, 22]. The idea is to use lower-resolution images to match large (low frequency) features and focus (constrain) the search space at the neighboring analysis level. Ultimately, at the finest level, the analysis duplicates that of the non-hierarchical approach except that the correlation computations are computed over what amounts to a fraction of the original search space. In this way, global and regional information is filtered through the hierarchical structure to focus the analysis and thereby reduce overall computation cost [22]. This sharing of information among levels lends an aspect of continuity as well.

In our experience, the above hierarchical search does reduce false matches from multimodal surfaces and can be augmented by temporal focusing as well. However, because this solution seeks only to localize computation of the unconstrained CS surface (it constrains only the search space), the quality of the displacement estimate is still subject to the local properties (noise) of the unconstrained CS surface (as shown in the cross-correlation surface in Figure A in the sidebar entitled "Higher-Moment Analysis"). Search-space censoring is also vulnerable to catastrophic (propagating) errors if the forcing information becomes corrupted.

### **Displacement-Field Estimation**

What hasn't been suggested, to our knowledge, is the notion of smoothing the CS meta-image. In many respects, the estimation problem is better posed in the CS domain: all displacement errors are easily associated with CS quality, and gross errors typically result because a nonlinear threshold operation has been applied to a noisy surface. Furthermore, neighboring correlation surfaces often complement one another in motion sensitivity. This all leads quite naturally to the notion of CS meta-image modeling, which has become a focal point for our study.

Our intent in this final section is to motivate the above idea further while providing an overview of the strategy. To begin, our current FAA application motivates four (self-imposed) design requirements, which

are as follows. First, because it is limited to a fixed and incomplete correlation sampling, the estimation method must nevertheless support analysis centered at each image pixel (it must solve the interpolation problem). Second, the estimation model should account for spatial structure (continuity and smoothness) by incorporating prior understanding and/or observational evidence. Third, there should be a mechanism to remove or de-emphasize outliers. Fourth, the method must remain simple; i.e., it must stay within our computational bounds.

### *Hierarchical Averaging*


To motivate our model, we present a smoothing analysis argued from a second hierarchical pyramid. In contrast to the image-resolution pyramid described above, we let each level of a pyramid represent an analysis with increasingly larger image masks (or, in other words, decreasing levels of analysis resolution). At one extreme, the apex can be a correlation analysis that pattern-matches the whole image. Let the displacement estimates be derived from the average of global, regional, and local surfaces. Hence correlation surfaces computed at one level act as a biasing "data prior" for the level below.

### *A Computational Note*

Correlation surfaces do not need to be computed for each pyramid level. When the error measure employed for CS construction is additive (unnormalized correlation and mean-squared error fall into this category; when image scale and offset are not issues, which is more or less the case for our meteorological application, these criteria work well), and when the resolution masks are defined accordingly (upper-level masks are derived from the union of lower-level masks), upper-level correlation surfaces can be obtained by adding base-level correlation surfaces. In that case, hierarchical averaging reduces to a weighted sum of primitive (i.e., base level) CS measurements. Figure 10(a) illustrates the four-level partitioning relationships in a quadtree arrangement.

### *Smoothing Results*

Figure 10(b) shows the CS pyramid using the precipitation data of Figure 6. Motion estimates at the base



**FIGURE 10.** (a) Hierarchical analysis structured on decreasing/increasing analysis resolution. At the pyramid base, the image is partitioned into primitive correlation boxes (here, 14 km x 14 km image blocks). At level two, the partition is collapsed into four primitives per correlation box; this process is repeated at level three. At the apex whole-image correlation is used. (b) Quadtree pyramid of CS surfaces. A motion estimate at the base level is obtained from an average of the base surface and its relations up the pyramid.

are derived from the average of surfaces in the lineage of each base-level location. In actuality, a weighted sum of base-level surfaces is computed. In general, the weights are not constant because of the random holes corresponding to missing data, such as the blank regions at each level in Figure 10(b). The potential of such an averaging scheme is illustrated in Figure 11, which shows the corresponding displacement vectors. Most of the gross errors seen previously in Figure 6 are corrected; (human validated) growth-and-decay perturbations remain, however, because of their CS measurement strength.

#### *Linear CS Estimation*

The pyramidal averaging procedure amounts to CS smoothing and appears to provide greatly improved displacement measurements. However, it has two deficiencies: it does not model CS quality and the pyramid weights are arbitrary. Since a “good” correlation analysis has certain characteristic properties that we *can* detect, we should utilize this information when obtaining the pyramid weights. With regard to the pyramid weights, stratiform rain and snow are quite

different from convective environments, which suggests we should perform smoothing as conditions dictate. This suggestion has been confirmed as an empirical observation from earlier attempts to smooth displacement-vector fields. Hence we have been led to consider a simple extension: a conceptualization that puts CS computation in the context of a linear-estimation problem.

Linear estimation is a good next step because it builds on the smoothing results, it requires a modeling of no more than first-order and second-order properties, and it maintains computationally simple solutions. For our estimation problem, then, the observations are defined to be the fixed-location CS primitives (data-optimized positioning is also possible but not in our current plans). In contrast to the higher-moment analysis technique presented in the sidebar “Higher Moment Analysis,” each unobserved variate is *not* likened to an autocorrelation surface; rather, it represents an idealized version of what a correlation surface should look like. Each observed variate contains an additive measurement noise representing the combined degrading effects of image



content (as represented in an autocorrelation analysis) and evolution.


For any image pixel  $i$ , we hypothesize a corresponding ideal surface  $z_i$  that describes the motion (as a likelihood, say). Let  $Z' = \{z'_1, \dots, z'_N\}$  represent the computed CS measurements from  $N$  locations. The form of the model is then

$$z'_j = z_j + \eta_{ij} \quad (j = 1, \dots, N),$$

where  $\eta = \{\eta_1, \dots, \eta_N\}$  represents an image noise process, which (for first approximations) is assumed to be zero mean, uncorrelated, and white (see Equations 8–10 in the appendix, entitled “Linear Estimation of Random Fields”). The generalized objective is, for any pixel  $i$ , take as displacement measurement the MAXCOR displacement of the estimated surface

$$\hat{z}_i = \sum_{j=1}^N a_{ij} z'_j, \quad (1)$$

where  $a_i = \{a_{i1}, \dots, a_{iN}\}$  is the weighting of data sample  $Z'$  for pixel  $i$ . Pixel  $i$  does not have to correspond to one of the measurement locations  $j$ . When the index  $i$  is not one of the measurement locations,



**FIGURE 11.** A motion analysis using the hierarchical averaging approach. The weather reflectivity data are the same as in Figure 6. This figure clearly illustrates a more accurate estimation of storm motion.

Equation 1 represents an interpolation formula.

We can appeal to the minimum mean-squared error (MMSE) criterion to determine an optimal  $a_i$ . The classic and well-known formulation of MMSE often begins with the assumption of a known mean and covariance for  $z_j$ , which is all that's needed, and from which suboptimal formulations are built when neither is actually known. A zero-mean assumption is inappropriate, and modeling and/or estimating the mean departs from our estimation objective. Given its hypothetical status, we would rather not make any statement regarding the mean of  $z_i$  (other than that we don't know it).

Luckily, there is an alternative to covariance modeling that does not require knowledge of the mean and is arguably equivalent when the mean is unknown. This is the use of the variogram function to capture second-order properties. The variogram seems particularly suited to spatial-estimation problems, and its merits are further argued in the appendix. Let  $\text{vec}(z_i)$  represent a vector created from surface  $z_i$  by a prescribed stacking of its pixel values. If

$$E[z_i] = E[z_j]$$

for all pixels  $i$  and  $j$ , the CS variogram can be defined as

$$\gamma_{ij} = \frac{1}{2} E \left[ \text{vec}(z_i - z_j)^T \text{vec}(z_i - z_j) \right].$$

The formulation for the MMSE weight  $\hat{a}_i$ , given a known variogram, essentially parallels the development for covariance modeling (see the appendix, “Linear Estimation of Random Fields”). Of course, the variogram is not known either, although a number of similar models are routinely used in the geosciences. Even with a stationarity assumption, i.e.,  $\gamma_{ij} = \gamma_{|i-j|}$ , finding an estimate for the variogram in real time would be difficult. What is feasible is adaptively selecting one from a fixed number of variogram models conditioned on the meteorological setting (convection, stratiform, snow). In real time, we would adaptively select the most appropriate variogram to use and thereby specify the degree of local smoothness. We plan to assess this technique quantitatively in the near future.

In the absence of measurement noise  $\eta$ , the weighting  $a_i$  is determined by the spatial structure captured in the variogram. If we assume that  $\gamma_{ii} = 0$ , the result is an exact interpolator. To deal with the discontinuity of outliers we must have a mechanism for de-emphasizing a particular measurement, essentially to adjust weighting  $a_i$ . For example, in linear regression a residual analysis often can be used to adjust weights [36]. In our case, we have the ancillary information of our canonical score function, which indicates a degree of corruption of the data. The modification of weighting  $a_i$  according to this confidence can follow from Equation 7 in the appendix and a model of the noise variance  $\sigma_i^2$  that is indexed by canonical score.

### Concluding Remarks

Accurate depictions of both current and future storm locations are necessary in the exchange of weather information between controller and pilot. With a rapidly updated display, systems like ITWS can easily achieve accurate depictions of current storms, but the accurate portrayal of future storms remains a difficult challenge well beyond the scope of mere tracking. Nevertheless, to achieve a practical goal, we can categorize storm motion and storm growth and decay to satisfy two of the more critical weather-information needs of air traffic controllers.

Historically, the automation of storm tracking has not been a simple task. The tracing of features, as from a meteorological analysis, is limited because of the uncertain (dynamic) nature of weather and the ambiguity in sampled radar images between storm growth and decay and storm motion. We have found local-area correlation to be the most robust approach precisely because it does not rely on explicit (invariant) features. In our work for the FAA, we have shown that reliable automated tracking can be achieved [9, 10].

To tackle storm growth and decay, we must measure and quantify it apart from any intertwined advective motion. The characteristic projection of growth and decay onto correlation measurements is promising in this respect, but more study is needed regarding the filtering of motion components and differencing of data fields. A “well-behaved” motion

field is crucial to such decomposition. Our experience with traditional (centroid) tracking methods leaves no doubt that their error performance is inadequate for such growth-and-decay computations. Local-area correlation is also subject to error, but in this method there is considerable room for improvement that comes with an understanding of the correlation “sensor” (measurement inhomogeneity, for example) and modeling of the motion-estimation process.

In dealing with correlation degradation, we have often used the heuristic relations “convection equals noise” and “stratiform equals weak signal,” and we view our motion problem in the context of image reconstruction from incomplete and noisy observations. As a result, we herein propose an extension to local-area correlation following a somewhat classical approach to the signal-in-noise problem. We have attempted to motivate our new approach through reconciliation with other image-processing heuristics. Our extension is significant in that it deals with the important issues of measurement variability (including outlier handling) and model-based constraint. In deference to covariance modeling, we believe the variogram might be better suited to our particular spatial problem. Because our approach follows a linear-estimation framework, it offers practical solutions suitable to a limited-resource environment. Future efforts directed at the issue of storm growth and decay will include the further development and validation of this estimation strategy.

### Acknowledgments

Of the many people to whom we owe thanks, we note in particular the efforts of Linda Garant and David Askey in the development of our real-time Air Traffic Control algorithm and in contributing to our scientific investigations.



## REFERENCES

1. "Aircraft Accident Report: Runway Departure Following Landing American Airlines Flight 102 McDonnell Douglas DC-10-30, N139AA Dallas/Fort Worth International Airport, Texas April 14, 1993," National Transportation Safety Board, Washington, DC, 1994.
2. L. Stevenson, "Thunderstorm Impact on Denver Air Traffic Control Operations and the Role of NEXRAD," Final Report DOT-TSC-RSPA-84-1, U.S. Department of Transportation, Research and Special Programs Administration, Transportation Systems Center, 1984.
3. J.E. Evans and D. Turnbull, "Development of an Automated Windshear Detection System Using Doppler Weather Radar," *Proc. IEEE* 77, 1661 (1989).
4. D. Turnbull, J. McCarthy, J. Evans, and D. Zrnic, "The FAA Terminal Doppler Weather Radar (TDWR) Program," *Third Int. Conf. on the Aviation Weather System, Anaheim, CA, 30 Jan.-3 Feb. 1989*, p. 414.
5. T.A. Noyes, S.W. Troxel, M.E. Weber, O.J. Newell, and J.A. Cullen, "The 1990 Airport Surveillance Radar Wind Shear Processor (ASR-WSP) Operational Test at Orlando International Airport," *Technical Report ATC-178*, MIT Lincoln Laboratory (17 July 1991).
6. J.E. Evans, "Integrated Terminal Weather System (ITWS)," *Fourth Int. Conf. on Aviation Weather Systems, Paris, 24-28 June 1991*, p. 118.
7. Martin Marietta Air Traffic Systems, "Cost-Benefit Analysis of Integrated Terminal Weather System," under contract to Dept. of Transportation, Federal Aviation Administration. ATC-92-1201 under contract DTFA01-84-C-00017, Oct. 1992.
8. Air Traffic Weather Requirements Report, prepared by the Air Traffic Requirements Team for the FAA under the auspices of the Air Traffic Plans and Requirements Service, Feb. 1993.
9. E.S. Chornoboy, "Storm Tracking for TDWR: A Correlation Algorithm Design and Evaluation," *ATC Report ATC-182*, MIT Lincoln Laboratory (14 July 1992), DOT/FAA/NR-91/8, DTIC #AD-A254701.
10. E.S. Chornoboy and A. Matlin, "Extrapolating Storm Location Using the Integrated Terminal Weather System (ITWS) Storm Motion Algorithm," *ATC Report ATC-208*, MIT Lincoln Laboratory (27 March 1994), DOT/FAA/RD-94/2, DTIC #AD-A278611.
11. J.C. Brasunas and M.W. Merritt, "Short-Term Prediction of High Reflectivity Contours for Aviation Safety," *Ninth Conf. on Aerospace and Aeronautical Meteorology, Omaha, NB, 6-9 June 1983*, p. 118.
12. J.C. Brasunas, "A Comparison of Storm Tracking and Extrapolation Algorithms," *ATC Project Report ATC-124*, MIT Lincoln Laboratory (31 July 1984), DOT-FAA-PM 84-1, DTIC #AD-A146638.
13. R.K. Crane, "Automatic Cell Detection and Tracking," *IEEE Trans. Geosci. Electron.* GE-17, 250 (1979).
14. M. Dixon and G. Wiener, "TITAN: Thunderstorm Identification, Tracking, Analysis, and Nowcasting—A Radar-Based Methodology," *J. Atmos. Ocean. Technol.* 10, 785 (1993).
15. A. Witt and J.T. Johnson, "An Enhanced Storm Cell Identification and Tracking Algorithm," *26th Int. Conf. on Radar Meteorology, Norman, OK, 24-28 May 1993*, p. 141.
16. R.A. Houze, Jr., and P.V. Hobbs, "Organization and Structure of Precipitating Cloud Systems," *Adv. Geophysics* 24, 225 (1982).
17. K.E. Browning, "Morphology and Classification of Middle-Latitude Thunderstorms," *Thunderstorm Morphology and Dynamics, Thunderstorms: A Social, Scientific, and Technological Documentary*, ed. E. Kessler, chap. 7 (University of Oklahoma Press, Norman, OK, 1985), pp. 133-152.
18. H.R. Byers and R.R. Braham, Jr., *The Thunderstorm* (U.S. Dept. of Commerce, Washington, DC, 1949).
19. G.L. Austin and A. Bellon, "Very-Short-Range Forecasting of Precipitation by the Objective Extrapolation of Radar and Satellite Data," *Nowcasting*, ed. K.A. Browning (Academic Press, London, 1982), pp. 177-190.
20. R.H. Blackmer, Jr., R.O. Duda, and R. Reboh, "Application of Pattern Recognition Techniques to Digitized Weather Radar Data," Final Report SRI Project 1287, Stanford Research Institute, 1973.
21. R.E. Rinehart, "Internal Storm Motions from a Single Non-Doppler Weather Radar," Technical Note NCAR/TN-146+STR, NCAR, 1979.
22. F.C. Glazer, "Hierarchical Motion Detection," Ph.D. thesis, University of Massachusetts, Amherst, 1987.
23. R.L. Delanoy and S.W. Troxel, "Machine Intelligent Gust Front Detection," *Linc. Lab. J.*, 6, 187 (1993).
24. P.P. Neilley and L.P. Carson, "Radar Image Tracking and Its Use in a Short-Term Snowfall Prediction System," *26th Int. Conf. on Radar Meteorology, Norman, OK, 24-28 May 1993*, p. 148.
25. B.K.P. Horn and B.G. Schunck, "Determining Optical Flow," *Artif. Intell.* 17, 185 (1981).
26. M.W. Merritt, "Use of Computer Vision Techniques in Estimating Echo Motion," *20th Conf. Radar Meteorology, Boston, 30 Nov.-3 Dec. 1981*, p. 417.
27. P.J. Burt, C. Yen, and X. Xu, "Multi-Resolution Flow-Through Motion Analysis," *IEEE CVPR Conf., Washington, DC, 19-23 June 1983*, p. 246.
28. P. Anandan, "Computing Dense Displacement Fields with Confidence Measures in Scenes Containing Occlusion," COINS Technical Report 84-32, University of Massachusetts, Amherst, 1984.
29. T.W. Anderson, *An Introduction to Multivariate Statistical Analysis* (Wiley, New York, 1984).
30. D.G. Long and J.M. Mendel, "Model-Based Estimation of Wind Fields over the Ocean from Wind Scatterometer Measurements, Part I: Development of the Wind Field Model," *IEEE Trans. Geosci. Remote Sens.* 28, 349 (1990).
31. P. Lynch, "Deducing the Wind from Vorticity and Divergence," *Mon. Weather Rev.* 116, 86 (1988).
32. P. Anandan and R. Weiss, "Introducing a Smoothness Constraint in a Matching Approach for the Computation of Displacement Fields," COINS Technical Report 85-38, University of Massachusetts, Amherst, 1985).
33. P.Y. Simard and G.E. Mailloux, "A Projection Operator for the Restoration of Divergence-Free Vector Fields," *IEEE Trans. Pattern Anal. Machine Intell.* 10, 248 (1988).
34. J.-M. Jolion and A. Rosenfeld, *A Pyramid Framework for Early Vision* (Kluwer Academic Publishers, Dordrecht, The Netherlands, 1994).
35. P.J. Burt, "The Pyramid as a Structure for Efficient Computation," in *Multiresolution Image Processing and Analysis*, ed. A. Rosenfeld (Springer-Verlag, Berlin, 1984), pp. 6-35.
36. P.J. Huber, *Robust Statistics* (Wiley, New York, 1981).
37. P.J. Burt, C. Yen, and X. Xu, "Local Correlation Measures for Motion Analysis: A Comparative Study," *IEEE Proc. Pattern Recognition and Image Processing, Las Vegas, NV, 14-17 June 1982*, p. 269.

## APPENDIX: LINEAR ESTIMATION OF RANDOM FIELDS

### Linear Mean-Square Estimation

Let  $z(x)$  ( $x \in D \subseteq R^2$ ) be a random field and let  $Z = \{z_1, \dots, z_N\}$  be a finite realization from  $N$  scattered measurement locations in  $D$ . The minimum mean-squared error (MMSE) estimate for  $z$  at  $x_0$ , where  $x_0 \in D$  is arbitrary, is easily formulated from standard theory. For example, assuming the mean,

$$m = E[z(x)],$$

and covariance,

$$R_{ij} = E\{[z(x_i) - m][z(x_j) - m]\},$$

are known, the (unbiased) linear MMSE is

$$\hat{z}_0 = \hat{z}(x_0) = m + \sum_{i=1}^N \alpha_i (z_i - m), \quad (1)$$

where the weighting coefficients  $\{\alpha_1, \dots, \alpha_N\}$  solve the  $N$  equations

$$R_{0i} = \sum_{j=1}^N \alpha_j R_{ij} \quad (i = 1, \dots, N), \quad (2)$$

(see Reference 1, for example). Of course, in practice the mean and covariance must both be given an assumed form, or they must both be estimated; clearly, the latter is difficult with random-field samples that are scattered and few.

### Kriging

In the 1950s, D.G. Krige, a South African mining engineer, established the foundations for a practical 2-D interpolation method now known as *kriging*. G. Matheron later established much of the theory for this estimation method [2, 3], which is based on modeling a phenomenon's so-called *variogram*. Assuming a constant mean as above, the spatial variogram for  $z$  can be written as

$$\gamma_{ij} = \gamma(x_i, x_j) = \frac{1}{2} E\{[z(x_i) - z(x_j)]^2\}.$$

As a characterization of second-order properties, the variogram places the focus on the process *increments*. When  $R_{ij} < \infty$ , the relationship between variogram and covariance is

$$\gamma_{ij} = \frac{1}{2} R_{ii} + \frac{1}{2} R_{jj} - R_{ij},$$

and the continuity of one implies that of the other. Relationships between process and covariance are mimicked by those of process and variogram. In particular, mean-square continuity of  $z(x)$  implies continuity of the variogram and

$$\gamma_{ij} \downarrow 0 \quad (x_i \rightarrow x_j).$$

Although a common device in the geosciences, the variogram is not as well known to systems engineers, but it is finding application as a statistic for image/scene analysis [4, 5].

If the mean is unknown, a linear estimate for  $z_0$  may have the form

$$\hat{z}_0 = \beta_0 + \sum_{i=1}^N \beta_i z_i, \quad (3)$$

but if  $\hat{z}_0$  is to be unbiased, the relations

$$\beta_0 = 0 \quad \text{and} \quad \sum_{i=1}^N \beta_i = 1 \quad (4)$$

must hold. Minimizing Equation 3 subject to the constraints of Equation 4 leads to the system

$$\gamma_{0i} = \sum_{j=1}^N \beta_j \gamma_{ij} + \mu \quad (i = 1, \dots, N), \quad (5)$$

$$\sum_{j=1}^N \beta_j = 1, \quad (6)$$



where  $\mu$  is the Lagrange multiplier for the constraint in Equation 6.

When both mean and covariance are known, Equation 3 (with Equations 5 and 6) is suboptimal in comparison to Equation 1 (with Equation 2) [6], but there is no apparent advantage should the mean be unknown. In fact, since the variogram estimator does not require knowledge of the mean, the variogram should be preferred over Equation 1 when the mean is unknown.

Given the need to estimate variogram or covariance, it is usually necessary to include an assumption of stationarity to make practical headway. An *intrinsic* random field is one for which both mean and variogram are stationary; that is, in addition to a constant mean,

$$\gamma(x_i, x_j) = \gamma(d_{ij}),$$

where  $d_{ij} = |x_i - x_j|$ . The class of intrinsic random fields includes the covariance stationary fields and also includes those nonstationary random fields with stationary increments. Hence, the intrinsic-field assumption is a weaker requirement (safer modeling choice). Phenomena with large dispersive capacities are good candidates for variogram modeling.

### Extensions

The above estimators are exact interpolators. The situation in which we observe a noise-corrupted version of  $Z$ , i.e.,

$$z'_i = z_i + \eta_i \quad (i = 1, \dots, N),$$

is straightforward [7] and leads to a modified Equation 5. This modified equation can be written as

$$\gamma_{0i} = \sum_{j=1}^N \beta_j \gamma_{ij} + \sigma_i^2 \beta_i + \mu \quad (i = 1, \dots, N), \quad (7)$$

where  $\sigma_i^2$  is the noise variance, and where it has been assumed that

$$E[\eta_i] = 0, \quad (8)$$

$$E[\eta_i \eta_j] = \sigma_i^2 \delta_{ij}, \quad (9)$$

and

$$E[\eta_i z_j] = 0 \quad (10)$$

(for all  $i$  and  $j$ ). Finally, all of the above can be extended to higher dimensions ( $M \times M$  CS vectors, for example) by using Euclidean distance to interpret all errors.

## REFERENCES

1. T. Kailath, "Lectures on Wiener and Kalman Filtering," No. 140 in *International Centre for Mechanical Sciences: Courses and Lectures* (Springer-Verlag, Berlin, 1981).
2. G. Matheron, "The Intrinsic Random Functions and Their Applications," *Adv. Appl. Prob.* 5, 439 (1973).
3. G. Matheron, "The Theory of Regionalized Variables and Its Applications," *Cahier du Centre de Morphologie Mathématique*, No. 5, 1971.
4. D.L.B. Jupp, A.H. Strahler, and C.E. Woodcock, "Autocorrelation and Regularization in Digital images I. Basic Theory," *IEEE Trans. Geosci. Remote Sens.* 26, 463 (1988).
5. J. Serra, *Image Analysis and Mathematical Morphology*, Vol. I (Academic Press, San Diego, CA, 1982).
6. M. Gevers, "On the Use of Variograms in Levinson Predictors," *Lecture Notes in Control and Information Sciences*, No. 63: *Analysis and Optimization of Systems*, eds. A.V. Balakrishnan and M. Thoma (Springer-Verlag, Berlin, 1984), pp. 268–278.
7. A. Ohsumi, "Linear Estimation of Random Fields with Second-Order Increments and Its Applications," *Automatica* 24, 203 (1988).



**EDWARD S. CHORNOBOY** is a staff member in the Radar Systems group. He received a B.E. degree in chemical engineering from Stevens Institute of Technology in 1977, an M.Sc. degree in electrical engineering and a Ph.D. degree in biomedical engineering from the Johns Hopkins University in 1985 and 1987, respectively. From 1987 to 1988 he was with the Electronic Systems and Signals Research Laboratory at Washington University. From 1988 to 1995 he was a staff member with the Air Traffic Surveillance group.



**ANNE M. MATLIN** is an assistant staff member in the Weather Sensing group. Her research is focused on providing weather products for aviation users. She received a B.A. degree in clinical psychology from the University of California, Berkeley, and she has also studied computer science and mathematics at Boston University.



**JOHN P. MORGAN** is an assistant staff member in the Weather Sensing group. His research is in software engineering and the development of algorithms for aviation weather products. He received a B.S. degree in mathematics from Pennsylvania State University, and an M.S. in mathematics from the University of Illinois at Urbana-Champaign.