4.7 GENERATION OF TURBULENCE AND WIND SHEAR ALERTS: ANATOMY OF A WARNING SYSTEM

C. S. Morse *, S. G. Carson, D. Albo, S. Mueller, S. Gerding, and R. K. Goodrich
Research Applications Program
National Center for Atmospheric Research

1. INTRODUCTION

The Juneau terrain-induced turbulence and wind shear project (Barron and Yates, 2004) is evolving from a wind information system that provides users with wind data into a warning system that additionally provides users with turbulence and wind shear alerts. These alerts are issued when the estimated turbulence or wind shear levels exceed thresholds associated with aircraft performance (Wilson, 2004). The key element in this process is estimating the hazard (turbulence or wind shear) level based on the available sensor data. A linear regression technique was chosen for this purpose based on previous work (Neilley, 1996). The statistical verification and validation of the resulting regressions are discussed by Fowler *et al.* (2004).

This paper gives a brief description of the overall warning system and describes in more detail the process by which the hazard regressions were developed.

2. SYSTEM OVERVIEW

Figure 1 shows a block diagram of the overall Juneau warning system. The system sensors consist of seven anemometer sites and three wind profiler sites in the locations shown in Fig 2.

The Anemometer Subsystem ingests 1-s wind speed and direction data from each of the ten individual anemometers in the system and performs quality control (Weekley *et al*., 2004) on that data. One minute statistics of the 1-s data, including mean and standard deviation, are computed for each of the anemometers. In a separate step the statistics from the redundant mountaintop anemometers at the same site are combined/selected as representative of the site.

The Profiler Subsystem ingests Doppler spectral data from each radar and performs quality control at the moments level using NIMA (Morse *et al*., 2002). From the quality-controlled moments, winds are calculated (Goodrich *et al*., 2002) and additional quality control is performed at the winds level. "Rapid-update" (every ~30 seconds) winds are generated by calculating new winds after every radar beam is collected.

- - -
* *Corresponding author address:* C.S. Morse, NCAR, PO Box 3000, Boulder, CO 80307
email: morse@ucar.edu

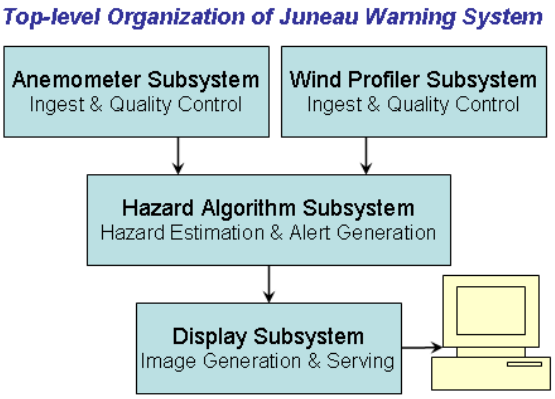

Fig. 1: The top level architecture of the Juneau warning system consists of four major subsystems: one for each of the major sensor components to perform data ingest and quality control, one to translate sensor data into hazards and alerts, and one to display the information to users.

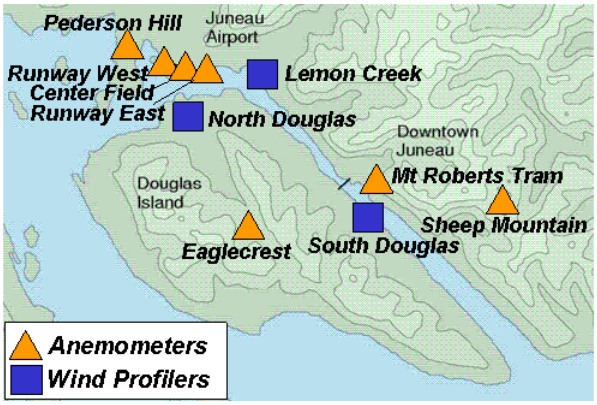

Fig. 2: The locations of the anemometers and wind profilers are shown relative to local terrain and landmarks in the Juneau area. The Gastineau Channel separates Douglas Island from the city of Juneau on the mainland.

The Hazard Algorithm Subsystem derives hazard estimates and alerts from the sensor data produced by the sensor subsystems. It consists of three major components as shown in Fig 3. The Regressor Generation component converts and formats the quality-controlled sensor data into the specific regressors that will be used to estimate the hazard levels. A fuzzy logic algorithm is applied to the wind speed and direction from the mountaintop

anemometers and the profilers to classify the current wind regime as one of Taku, Southeast, Mixed, Calm, or Unknown (Cohn *et al*., 2004).

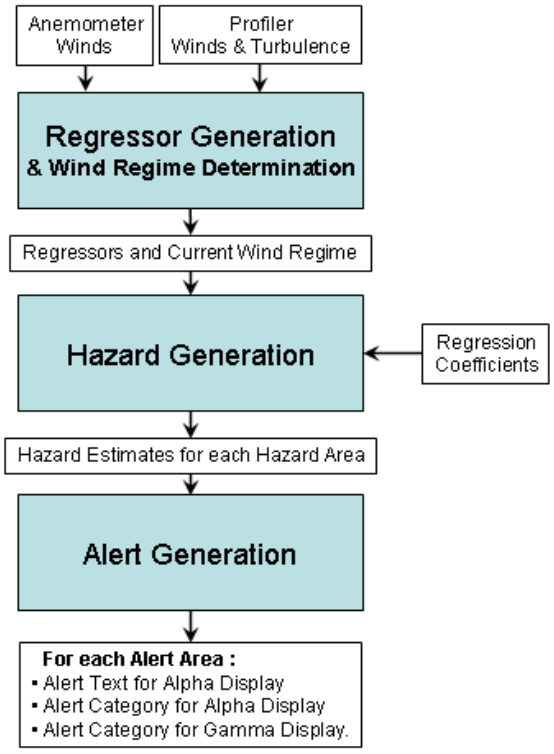

Fig. 3: The three major components of the Hazard Algorithm Subsystem. The Regression Coefficients are pre-calculated and supplied to the algorithm via configuration files.

The Hazard Generation component uses the regressors and wind regime, in combination with pre-calculated regression coefficients, to estimate the hazard levels in each of the hazard areas. The determination of these regression coefficients is described in Section 3. The wind regime is used to select the appropriate set of regressions (and their associated coefficients) to use for each hazard area. From this set of regressions, a subset is chosen to apply based on the skill of the regression and the availability of high confidence regressors. Simply stated, each regression in the set is considered, starting with the most skillful based on r^2, the square of the sample correlation coefficient. If each of its input regressors is available and has sufficiently high confidence, the regression is included in the subset, otherwise it is excluded. When the target number of regressions, currently 10, has been included in the subset, or the regression set has been exhausted, the subset is complete. Each of the regressions in the subset is evaluated to estimate the hazard level and these estimates are averaged to give the final hazard estimate. This process is repeated for each hazard area and hazard type. These hazard estimates are used by the Alert Generation component to determine what alerts should be displayed to users.

The behavior of the Alert Generation component is largely dependent upon the requirements of the FAA. By design, all the alerting complexity is embodied in this component. Because alerting strategy is an area where requirements often are subject to change, this component was designed to be flexible and has the capability of generating and prioritizing multiple types of alerts, e.g., wind shear and turbulence, as well as implementing different alerting philosophies, e.g., worst encounter, first encounter. The current prototype alerting system being fielded produces turbulence alerts only at three specific alert levels; each alert area is categorized based on the highest turbulence estimate from any of the hazard areas that overlap that alert area. Each area is categorized as having no alert, or having an alert at one of three levels. If there are no hazard estimates available from the associated hazard areas, the alert area is categorized as "out of service" (OTS). Additionally, if only some of the associated hazard areas do not have hazard estimates available, an "impaired" flag is set for the alert area.

The Display Subsystem receives the quality-controlled sensor data and the alerts and reformats this information for display to the users (Mueller *et al*., 2004). By design, components in the Display Subsystem are "dumb" and do not include any significant algorithmic processing. The prototype uses a variety of web-based technology. Text-based web pages that display the current anemometer and profiler winds data have been a part of the wind information system for several years. The prototype system also has two Java-based display clients. The Alpha display displays only the textual alert messages for each alert area and is intended for use in the air traffic control tower. The Gamma Display, shown in Fig. 4, provides a greater variety of information including a geographical situation display and time history of the sensor data, and is intended for use by the Flight Service Center, airline operations, and other general aviation users.

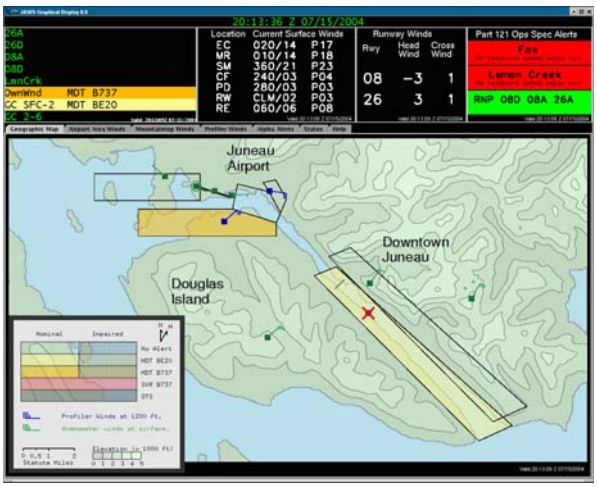

Fig. 4: Example of the Gamma Display showing the Alpha alerts in the upper left corner and the geographic situation display in the main window. Two alert areas show alerts.

3. DETERMINING THE HAZARD REGRESSIONS

The key component of the Hazard Generation algorithm is the set of regression coefficients used to estimate the hazard levels. The method by which these coefficients are generated is discussed in this section.

3.1 Where and how can we skillfully predict hazards?

The first step was to identify geographic regions where aircraft hazards are observed (Cohn *et al*., 2004 and Braid, 2004). Figure 5 shows research aircraft tracks flown in the Juneau area. The indicated hazard boxes were chosen so that the turbulence "hot spots" were located near the center of the boxes in order to generate good statistics.

Because turbulence occurs in different regions under different wind regimes, improved skill is expected if a separate analysis is done for each of the different wind regimes. This expectation was corroborated by the results of the feasibility study (2000). This study also indicated that better skill was achieved when the Gastineau channel hazard areas were further refined by altitude. Each of the four channel hazard areas shown in Fig. 5 represents two hazard areas, one below 2000 ft and another at 2000 ft and higher.

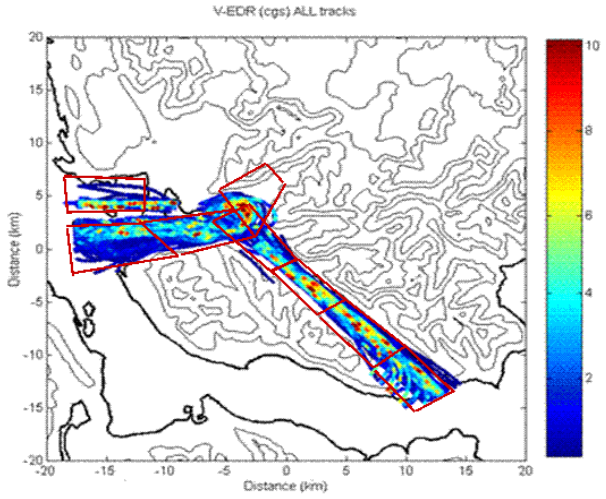

Fig. 5: Eddy dissipation rate as measured by research aircraft along flight tracks in the Juneau area are shown with the hazard prediction boxes overlaid.

A linear regression approach was chosen as the method to predict the hazards based on the successful use of this technique by Neilley (1996) for estimating terrain-induced turbulence hazards from anemometer data in Hong Kong. For this application, regressions of the form shown in Eqn. 3.1 were used.

$$p = c_0 + \sum_{i=1}^{4} c_i r_i \qquad (3.1)$$

where p is the predicted hazard,
r_i are the regressor values, and
c_i are the regression coefficients.

3.2 What is truth?

The source of the truth data for the regressions was research aircraft data collected during three field programs in the winters of 1998, 1999/2000, and 2002/2003 (Cohn *et al*. 2004). In the 1998 field program the University of North Dakota Citation was employed as the research aircraft and in the latter two field programs the University of Wyoming King Air was employed. During the 2002 field program, data was also collected from an Alaska Airlines B737 aircraft.

It was critical to perform in-depth review and quality-control on this aircraft data (Gilbert *et al*., 2004) in order to insure that only valid truth data were used in generating the regression coefficients. The aircraft data analysis resulted in eddy dissipation rate (EDR) values calculated at time resolution of one second, each value representing an average over a distance of one kilometer, about 12 seconds. EDR values for the time period when the aircraft was within the geographic confines of a specific hazard box were collected and statistics generated. Because of the intermittent nature of turbulence it was desired to use a target statistic that captured the worst hazard observed in the box but one that was still averaged enough to produce stable statistics. To this end, the one-second data were replaced with the 3-point median value of the one-second values at each point and the target statistic for the regressions was chosen as the largest of those median values within the box. In the following discussions, a "data point" is one excursion of the aircraft through a hazard box and the associated target hazard value.

3.3 Choosing the candidate regressors

The set of candidate regressors was chosen based on the measurements that would be likely to have a positive correlation with the atmospheric turbulence as measured by the research aircraft. A set of simple regressors could be obtained from the one-minute statistics generated for each anemometer site, i.e., the mean wind speed, the wind speed standard deviation, and the wind direction standard deviation. Because the dimensional analysis suggests that the square of the wind speed might be more likely to be linearly correlated with the EDR than the non-squared quantity, the square of the mean wind speed is also included as a candidate regressor along with the wind speed and wind direction variances.

Because of the much larger amount of raw winds data available from the profilers, additional averaging was

performed to generate regressors from the profiler data. The range of the profilers was divided into six altitude regimes each with a depth of 300 m. These regimes were characterized by the lowest included altitude, i.e., 300, 600, 900, 1200, 1500, and 1800 m, respectively. Over each of these altitude regimes the wind speed and direction data from each of the included profiler range gates were included in the statistics. Candidate regressors included the same quantities as used from the anemometer statistics with the addition of the eddy dissipation rate measured from the spectral widths and the standard deviation and variance of those measurements over the included range gates. Separate statistics were calculated for EDR measured as $\varepsilon^{1/3}$ and measured as $\varepsilon^{2/3}$, where ε is the eddy dissipation rate (m^2s^{-3}). For each altitude regime, the average vertical shear of the horizontal wind was also calculated as a candidate regressor.

In addition to these "simple" regressors, certain derived quantities were considered likely to be better correlated with the truth data. For example, the wind direction standard deviation scaled by the wind speed, might be better correlated than the sigma value alone since at a low wind speed, a large sigma value is expected and is not particularly interesting, but in conjunction with a large wind speed might suggest the presence of turbulence. Both wind speed and wind direction sigmas scaled by the mean wind speed were included as candidate regressors from each of the anemometers and profiler altitude regimes.

Another category of derived regressors is the wind speed component along some specific direction vector, as it is possible that the setup of a turbulent environment might be sensitive to winds coming from some specific direction. Component calculations, Eqn. 3.2, were carried out for values of d at ten degree increments, but only positive wind speed component values were used. Negative values of w_d were not used as regressor values. The squares of the positive w_d values were also considered as candidate regressors.

$$w_d = wspd * \cos(wdir - d) \qquad (3.2)$$

A final category of derived regressors were the "point-to-point shear" calculations. Here an attempt is made to directly calculate an estimate of the headwind shear experienced between two points, p_0 and p_1, along a specific flight path, as illustrated in Fig 6.

Let $\vec{u}_0$ and $\vec{v}_0$ be the unit vectors along and perpendicular to the flight path at p0, respectively. Similarly, let $\vec{u}_1$ and $\vec{v}_1$ be these unit vectors at p_1.

The head wind shear s_h is calculated as the difference in the dot product of the unit vector along the flight path and the wind vector at the two points, normalized by the path length, as shown in Eqn. 3.3.

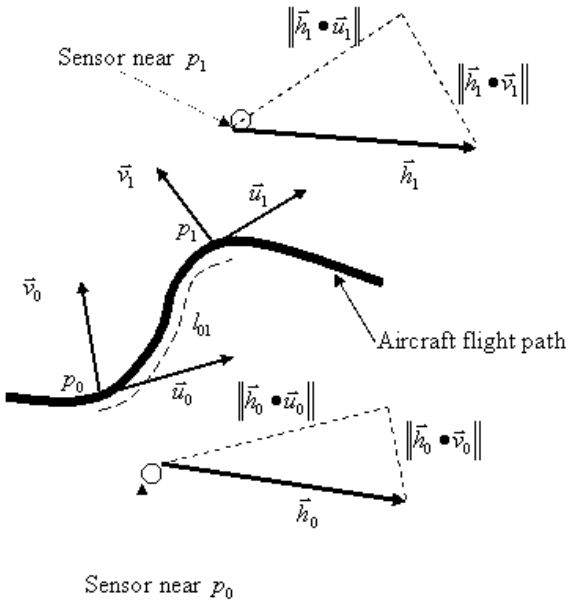

Fig. 6: Schematic diagram of the aircraft track between points p_0 and p_1 and the wind vectors, h, at the nearby sensors.

$$s_h = \frac{\vec{h}_1 \bullet \vec{u}_1 - \vec{h}_0 \bullet \vec{u}_0}{l_{01}} \qquad (3.3)$$

Similarly, the calculation of the cross wind shear s_x takes dot products between the wind vectors and the cross direction unit vectors, Eqn 3.4.

$$s_x = \frac{\vec{h}_1 \bullet \vec{v}_1 - \vec{h}_0 \bullet \vec{v}_0}{l_{01}} \qquad (3.4)$$

Headwind and crosswind shear regressors for six flight path segments were included among the candidate regressors. A total of 2138 candidate regressors were considered for use in predicting the hazards.

3.4 Finding skillful regressions

In preparation for the regression analysis, a tabular database was created, each row representing one excursion of the research aircraft through a hazard box. The columns of the database consisted of the various hazard values calculated from the aircraft data for that excursion, status information such as the identity of the hazard box, and the time-associated regressor values calculated from the sensor data. The latest sensor data received prior to the time the aircraft entered the hazard area were used to calculate these regressor values.

With such a large number of candidate regressors, a combinatoric explosion occurs if a brute force method is applied to testing all possible combinations of four

regressors. Several techniques were applied to mitigate this problem. The process described below was performed for each hazard area and hazard type for which regressions were required.

The first step was to test the individual skill of each regressor by calculating the single variable linear regression of each candidate regressor against the truth values using a Singular Value Decomposition technique. For a given hazard area and hazard type, the regressors with the largest r^2 values were selected for the next step with the restrictions that there be at least 20 data points for calculating the regression, that each selected regressor could have an r^2 value of no less than 0.1, and that no more than 500 regressors total were retained.

The next task was to generate sets of four-variable regressions to evaluate. The theoretical upper limit for the r^2 value of the four-variable regression is the sum of the r^2 values from the single variable regressions. This theoretical maximum can be achieved only if the variables are not correlated with each other. So as a prerequisite for selecting the four-variable regressions, the "covariance ratio" between each pair of regressors was calculated using the hazard estimates from their respective single variable regressions. The covariance of these hazard estimates was calculated resulting in a 2x2 covariance matrix, M. The "covariance ratio", c_{12}, between regressors R_1 and R_2 is calculated as shown in Eqn. 3.5.

$$c_{12} = \sqrt{\frac{M_{12}^2}{M_{11}M_{22}}} \tag{3.5}$$

To insure representation by all sensor sites and to limit the combinatoric explosion, a fixed number of regressors from each of the ten sensor sites were selected for use in the four-variable combinations. The site regressor with the largest r^2 is selected as S_0. Subsequent site regressors S_m are selected as the remaining regressor R_j with the highest score, Eqn 3.6, based on its value and its correlation with the previously selected regressors.

$$score_m(R_j) = r_j^{\,2} / \sum_{i=0}^{m-1} c_{ij} \tag{3.6}$$

For our analysis, five regressors from each sensor site were selected limiting the total number of possible combinations of four regressors to 230,300. The number of combinations was further reduced by eliminating combinations whose r^2 sum did not exceed the final r^2 minimum target of 0.6. Also eliminated were any which included pairs of regressors whose "covariance ratio" exceeded 0.7. Even with these *a-priori* exclusions, the resulting test sets typically consisted of many thousands of four-variable regressions.

Each four-variable regression in the test set was evaluated by examining its r^2 value. Regressions for which there were 20 or fewer data points were eliminated from consideration. Regressions with too high a condition number from the SVD analysis were also eliminated. For clarity in the condition number analysis, the regressor values were first normalized over their own ranges, i.e., regressor and target values, x, were replaced by $(x-\mu)/\sigma$, where μ is the mean value and σ is the square root of the variance of all the values for that regressor. This normalization also results in regression coefficients that are proportional to their contribution to the regression, uncontaminated by the effects of the varied scale and units of the unnormalized regressors.

Originally a fixed number of the regressions with the highest r^2 values were selected for use in the system. However, as the real time system would be averaging the estimates from the ten best-performing regressions, this approach might result in including a large number of regressions that might not ever be used.

An alternate method for selecting the final set of regressions to use was to consider possible sensor outage scenarios. First the entire set of regressions was ranked by a "figure of merit" (FOM) that was comprised principally of the regression r^2 value. As shown in Eqn. 3.7, the FOM was decreased for regressions calculated with fewer than 80 points.

$$FOM = r^2 f^{0.25} \tag{3.7}$$

where f is the smaller of unity and the ratio between n, the number of data points, and the product of $n_{\min}$, the minimum acceptable number of data points for a single variable regression, i.e., 20, and n_{var}, the number of variables in the regression, i.e., 4.

With this ranked set of regressions, each possible sensor outage scenario was considered beginning with no outage, a single sensor outage, dual sensor outage, etc. The ten physical sensor sites were considered for this exercise, resulting in 1021 possible scenarios. For each scenario, the best regressions in the list that did not use any regressors from the excluded sensor(s), up to a maximum of ten, were added to the final list. For a given hazard area and hazard type this resulted in final regression sets on the order of 1000 regressions.

Each regression set, consisting of the regressor names, associated coefficients and the r^2 value for each four-variable regression selected, was output to a configuration file for use by the real time warning system. One configuration file was generated for each hazard area/hazard type combination.

4. SUMMARY AND FUTURE WORK

We have outlined the architecture and design of a system that utilizes linear regression to estimate turbulence and wind shear in real time to produce alerts for pilots in the Juneau area. The linear regressions utilize quality-controlled data from wind profilers and from anemometers located at mountaintop and near the airport to estimate turbulence and wind shear hazards in each of several geographic areas. Separate regressions are used in different wind regime conditions. We have outlined the approach taken to generate the regression coefficients used in the system and described the approaches used to mitigate the combinatorial explosion that results from considering a large number of potential regressors.

It remains to evaluate the performance of the hazards regressions on a larger sample of data than that included from the field programs. This analysis can be performed by applying the hazard algorithms to the three years of sensor data archived thus far over the course of the project, as well as observing the performance of the prototype system in the field. Of particular interest is whether the regressions appear to be over- or under-warning as compared with pilot reports. As required, strategies to mitigate the under- or over-warning will need to be developed.

ACKNOWLEDGEMENTS

This research is in response to requirements and funding by the Federal Aviation Administration (FAA). The views expressed are those of the authors and do not necessarily represent the official policy or position of the FAA.

REFERENCES

Barron, R. and V. Yates, 2004: Overview of the Juneau terrain-induced turbulence and wind shear project. *Preprints, 11th Conference on Aviation, Range, and Aerospace Meteorology*, Hyannis, MA

Braid, J. T., 2004: Turbulence PIREPs in Juneau – an analysis. *Preprints, 11th Conference on Aviation, Range, and Aerospace Meteorology*, Hyannis, MA

Cohn, S. A., J. T. Braid, C. Dierking, M. K. Politovich and C. G. Wade, 2004: Weather patterns of Juneau Alaska and their relationship to aircraft hazards. *Preprints, 11th Conference on Aviation, Range, and Aerospace Meteorology*, Hyannis, MA

Cohn, S. A., R. Barron, A. Yates, S. Mueller, A. R. Rodi, P. P. Neilley, A. Praskovsky and L. B. Cornman, 2004: Field programs to investigate hazards to aviation in Juneau, Alaska. *Preprints, 11th Conference on Aviation, Range, and Aerospace Meteorology*, Hyannis, MA

Fowler, T. L., J. T. Braid and M. J. Pocernich, 2004: A performance analysis of the Juneau wind hazard alert system. *Preprints, 11th Conference on Aviation, Range, and Aerospace Meteorology*, Hyannis, MA

Gilbert, D., L.B. Cornman, A. R. Rodi, R. J. Frehlich and R. K. Goodrich, 2004: Calculating EDR from aircraft wind data during flight in and out of Juneau, AK: techniques and challenges associated with non-straight and level flight patterns. *Preprints, 11th Conference on Aviation, Range, and Aerospace Meteorology*, Hyannis, MA

Goodrich, R. K., C. S. Morse, L. B. Cornman and S. A. Cohn, 2002: A horizontal wind and wind confidence algorithm for Doppler wind profilers. *J. Atmos. Oceanic Technol.,* **19**, 257-273.

Juneau Alert Generation Feasibility Study, September 29, 2000.

Morse, C. S., R. K. Goodrich and L. B. Cornman, 2002: The NIMA method for improved moment estimation from Doppler spectra. *J. Atmos. Oceanic Technol.*, **19**, 274-295.

Mueller, S., C. S. Morse, D. Garvey, R. Barron, D. Albo and P. Prestopnik, 2004: Juneau airport wind hazard alert system display products. *Preprints, 11th Conference on Aviation, Range, and Aerospace Meteorology*, Hyannis, MA

Neilley, P. P., 1996: Using anemometers to detect and forecast turbulence at Hong Kong's New Airport. *Proceedings, Workshop on Wind Shear and Wind Shear Alert Systems*, Oklahoma City, OK, Amer. Meteor. Soc., 133-136.

Weekley, R. A., R. K. Goodrich, A. Praskovsky and L. B. Cornman, 2004: An anemometer data quality control method designed for a turbulence and wind shear prediction algorithm. *Preprints, 11th Conference on Aviation, Range, and Aerospace Meteorology*, Hyannis, MA

Wilson, F. W., 2004: Aviation impacts of terrain-induced wind shear. *Preprints, 11th Conference on Aviation, Range, and Aerospace Meteorology*, Hyannis, MA