J1.4 Quality Assessment of the National Ceiling and Visibility Analysis Product

Tressa L. Fowler*, Matthew J. Pocernich, Jamie T. Braid,
Anne Holmes and Richard E. Bateman
National Center for Atmospheric Research, Boulder, Colorado

1 INTRODUCTION

The National Ceiling and Visibility analysis product was evaluated over the winter of 2004/2005 and is compared with measurements from routine aviation weather reporting sites. The results are summarized via standard verification statistics and presented through a variety of plots. For purposes of comparison, the statistics for the operational ceiling and visibility aviation advisories are also included. The goal of this study is to evaluate the performance of the NCV product at analysis locations between METAR locations. A cross-validation approach makes this evaluation possible.

Data sources and methods are discussed in sections 2 and 3, respectively. Results from an analysis of flight category are given in section 4 while diagnostic measures of the skill of the ceiling and visibility components are given in sections 5 and 6 respectively. Section 7 contains the results of a sensitivity analysis to the verification method. Finally, conclusions are presented in section 8.

2 DATA

For this study, NCV hourly analyses and surface ceiling and visibility observations (METARs) over the CONUS during the period 25 October 2004 through 18 January 2005 are examined. These datasets are described in more detail in the following subsections.

2.1 NCV analysis product

The NCV analysis product is a ceiling and visibility diagnostic that combines observational data from satellite and surface observations to produce an analysis of ceiling and visibility conditions on a grid across the Continental U.S. with output from the numerical weather prediction model. To produce ceiling and visibility diagnosis operationally, forecasters subjectively examine these datasets and use established "rules of thumb" to reach conclusions about ceiling and visibility conditions that might be hazardous to aircraft. The NCV algorithm interpolates surface observations and satellite information to produce a grid of ceiling and visibility values. The current operational version of the NCV analysis produces these ceiling and visibility values on a two-dimensional grid corresponding to the horizontal grid structure of the Rapid Update Cycle (RUC; Benjamin et al., 2001) numerical weather prediction model. The NCV analysis product uses ceiling and visibility observations to determine ceiling and visibility values at the observation sites, and then applies an interpolation scheme to estimate ceiling and visibility values between sites. The product is updated hourly. For this report, hourly NCV analyses are evaluated for the period 25 October 2004 through 18 January 2005.

The ceiling and visibility values from the NCV product are converted into flight categories using the rule set in Table 1. The lowest of the ceiling and visibility conditions determine the flight rule. For instance, if the ceiling is 800 ft and visibility is 6 mi. the 800 ft ceiling causes the flight rule to be IFR.

Table 1: Ceiling and Visibility bounds for each flight category.

	Ceiling	Visibility
LIFR	Less than 500 ft	Less than 1 mi
IFR	Between 500 and 1000 ft	Between 1 and 3 mi
MVFR	Between 1000 and 3000 ft	Between 3 and 5 mi
VFR	Greater than 3000 ft	Greater than 5 mi

An example of the NCV analysis product is shown in Figure 1. Note that the product has a speckled appearance in some locations, for example, west of the Great Lakes. Each individual grid point on the NCV analysis grid is treated separately in the verification process, and there is no need for the flight category at one grid point to correspond in any way to the points around it. In fact, some IFR-or-worse conditions exist at a single location which is surrounded by less severe conditions.

***Corresponding author address:** Tressa L. Fowler,
P.O. Box 3000, Boulder, CO 80307-3000.
tressa@ucar.edu

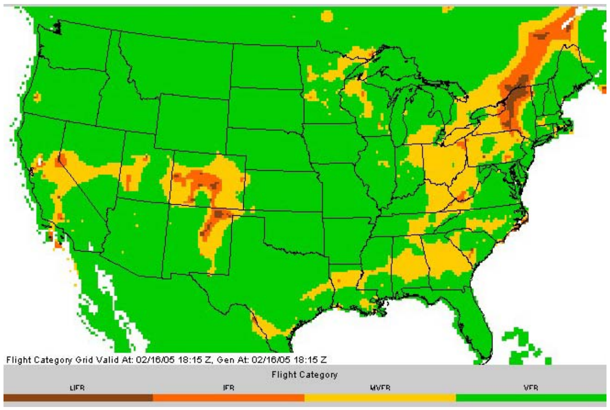

Figure 1: Example of the NCV analysis product

2.2 AIRMETs

The operational forecasts of IFR-or-worse flight rules are the AIRMETs. The AIRMETs are aviation advisories and are issued when expected ceilings are below 1,000 ft and/or expected visibility is below 3 mi. The AIRMETs are six hour forecasts issued every six hours that are often amended or canceled when ceiling and visibility conditions change during the forecast period. Furthermore, the AIRMET must cover a minimum area of 3,000 sq. mi. where the conditions are expected to cover most of the forecast area (NWS 1991). In this regard, the AIRMETs are quite different from the NCV analysis, which only attempts to analyze conditions at points on a grid. Figure 2 shows an example of some AIRMET polygons (outlined in red). The METARs have been overlayed in Fig. 2 in green (for IFR-or-worse conditions) and purple (for VFR-or-better conditions).

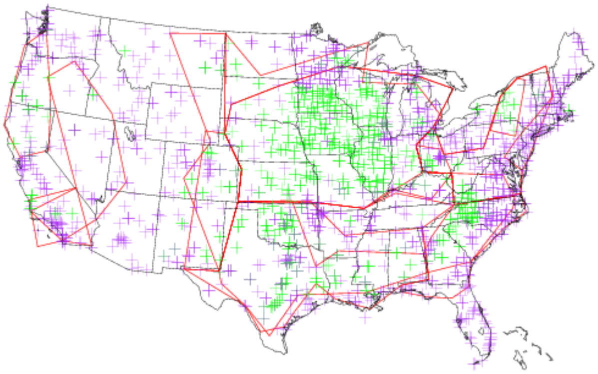

Figure 2: Example map of CONUS showing IFR AIRMETs (outlined in red; solid polygon shapes) and METARs (IFR conditions in green, VFR in purple).

2.3 METARs

Surface observations of ceiling and visibility, which are available in METARs (Aviation Routine Surface Weather Reports), are used to evaluate the NCV analysis product. The METARs measure ceiling in hundreds of feet near the surface and thousands of feet at higher levels. Similarly, visibility measures are given in fractions of a mile for lower visibility values, and in whole miles for higher visibility values. Observations are taken at least once per hour, though special or changing weather conditions can result in more frequent observations.

During the time period of interest, approximately 1600 METAR stations across the CONUS gave reports of ceiling and visibility values. Figure 2 shows the locations of these stations. Some states, such as North Carolina and Iowa, have a dense network of METAR stations. Other states, like Nevada and Montana, have a sparse network. Still other states, like California and Texas, have many stations located in the vicinity of metropolitan areas but relatively few stations in the remaining areas. The verification statistics may be impacted by the station density. For example, the San Francisco Bay area has several stations while the coast just north of this area has very few. Correctly identified IFR-or-worse conditions will be rewarded by more correct matching observations near San Francisco than along the more northerly coast. Thus, conditions in areas with a great density of stations are naturally given somewhat more weight in the statistics than conditions in areas with fewer stations.

Figure 3: Map showing METAR sites across the CONUS.

Recently, most U.S. METAR stations have been converted from human observing systems to Automated Surface Observing Systems (ASOS) (Bradley and Imbembo 1985; US DoC 1992). However, some stations still have observations from a human observer. Thus, the METARs may have some internal inconsistencies. While inconsistencies such as these are commonplace in meteorological observations, awareness of such issues is essential when interpreting results.

3 METHODS

3.1 *Matching the gridded NCV product to METAR sites*

The NCV analysis product is matched to the METAR observations. For each METAR location to be used in verification, the minimum ceiling and visibility measurements from the four surrounding grid points from the NCV analysis grid are selected. In particular, for each variable (ceiling and visibility) the minimum value from one of the four surrounding grid points on the NCV analysis grid is matched to the ceiling/visibility measurement from the METAR station observation to create a verification pair. The ceiling and visibility measurements may come from different gridpoints. Since the NCV product is on the RUC 20-km grid, the maximum distance between a METAR site and its matching grid location for the NCV analysis is less than 30 km.

3.2 *Cross-validation*

A cross-validation technique (Neter et al. 1996) was applied in the evaluation of the NCV analyses to ensure independence between the METAR stations used for verification and those used to create the product. Because the NCV analysis product uses METAR observations to determine ceiling and visibility values at the METAR sites, verification using the same METAR observations as were used to create the analysis would produce perfect verification statistics. In particular, in this case, the METAR observations would serve as both nowcasts and observations. The METARs will always match themselves exactly. *Thus, the goal of this study is to evaluate the performance at analysis locations between METAR locations. The cross-validation approach makes this evaluation possible.*

Using the cross-validation approach, 1300 METAR reports (referred to as the *training* set) out of nearly 1600 METAR sites, were randomly selected to produce each NCV analysis. The remaining 300 METAR stations (referred to as the *testing* set) were used to verify the product. In order to prevent a “bad” selection of METAR sites from affecting the statistical results, and to ensure that enough locations were chosen for verification, this procedure was repeated ten times for each analysis time to provide ten different testing/training METAR sets. Thus, ten different NCV analyses were produced at each time: one for each of the ten training sets. The verification statistics are based on the ten testing sets of METAR reports, accumulated across all of the NCV analyses included in the verification sample.[1,2]

Since ten different cross-validation versions of the NCV analyses were produced on each hour during the 25 October 2004 to 18 January 2005 time period, a total of 20,140 NCV analyses are available for verification.

3.3 *Verification statistics*

Overall verification statistics are calculated based on binary event/non-event categories. The four flight categories listed in section 2.1 are condensed into two by combining the bottom and top two categories, yielding the categories IFR or worse and MVFR or better. The verification statistics computed include the probability of detection (POD), the probability of detection for non-events (PODNo), Bias, and the False Alarm Ratio (FAR). In addition, three skill scores are included: (a) the Heidke Skill Score (HSS), (b) the Gilbert Skill Score (GSS), and (c) the True Skill Statistic (TSS). The percent of the CONUS covered by the average event area (Percent Area) is used as a measure of over-warning. Finally, the POD per unit area, known as Area Efficiency, is also included. Each statistic is calculated using the formulas listed in Table 3, based on a standard 2 by 2 contingency table as shown in Table 2. For more information on verification statistics for categorical forecasts, see Wilks (1995).

Table 2: Standard 2x2 contingency table for verification statistics. Entries in the table represent counts of each forecast/observation pair.

	METAR Flight Category	
NCV Flight Category	IFR or worse	MVFR or better
IFR or worse	YY	YN
MVFR or better	NY	NN

---

[1] A smaller number of METAR sites may be available at any given time for producing or verifying the analysis in the event of sensor or data outages.

[2] The results of the verification study may be somewhat sensitive to the proportion of “held-out” stations. Although this sensitivity is not expected to be large, it may have some impact on the results and is currently being investigated further.

Table 3: Verification statistics and their associated formulas based on counts from Table 2.

Statistic	Formula
POD	YY / (YY + NY)
PODNo	NN / (NN + YN)
Bias	(YY + YN) / (YY + NY)
FAR	YN / (YN + YY)
HSS	(YY + NN – C1) / (N – C1) { where C1 = [(YY + YN) (YY + NY) + (NY + NN) (YN + NN)] / (YY + YN + NY + NN) }
GSS	(YY– C2) / (YY - C2 + YN +NY) [where C2 = (YY + YN) (YY + NY) / (YY + YN + NY + NN)]
TSS	POD + PODNo -1
Percent Area	Average Event Area * 100 / Total CONUS Area
Area Efficiency	100 * POD / Area

The actual ceiling and visibility values are examined separately as well. In particular, the bias in the NCV ceiling and visibility values is assessed. Boxplots, histograms, and a contour plot (essentially a 3-dimensional scatter plot) are used to examine errors in and agreement between NCV and METAR values. Quantile-quantile (qq) plots are used to compare the distributions of NCV versus METAR values. Linear models are overlaid on the qq-plot to quantify the differences in distributions.

4 RESULTS OF FLIGHT CATEGORY ANALYSES

The verification results are summarized by flight category using the 2x2 verification statistics. Verification statistics for the NCV analysis product as compared to the AIRMETs are provided in Table 4. The NCV analysis achieves a lower POD (0.57 vs. 0.83) and higher PODn (0.97 vs. 0.81) as compared to the AIRMETS. The NCV product has a much lower false alarm ratio than the AIRMETs (0.19 vs. 0.43). On average, the NCV analysis product covers roughly 75% of the area covered by the AIRMETs. Both products are quite biased, but in opposite directions. The NCV product has a bias of 0.7, and thus identifies IFR-or-worse conditions less often than they occur. The AIRMETs' bias is 1.45, so they identify these conditions more often than they occur. The minimum size and time restrictions placed on the AIRMETs almost require an over-warning bias. Thus, the bias statistic for the AIRMETs should be viewed as a characteristic rather than a performance measure. The NCV product has slightly larger values of both the HSS and the GSS than the AIRMETs, but a smaller TSS value. The area efficiencies of the two products are roughly comparable, 35 for the NCV vs. 39 for the AIRMETs.

Table 4: Verification statistics for the NCV analysis product.

	NCV	AIRMETs
POD	0.57	0.83
POD NO	0.97	0.81
FAR	0.19	0.43
Bias	0.70	1.45
HSS	0.60	0.55
GSS	0.43	0.38
TSS	0.54	0.64
Percent Area	17	22
Area Efficiency	35	39

When examining the statistics in Table 4, it is important to remember that these statistics represent the algorithm performance between METAR stations. The POD at the METAR locations included in the analysis is close to 1 and the FAR is close to 0.

Verification statistics were also computed separately for day and night. However, the results differed very little from each other and from the overall results presented in Table 4. Thus, those statistics have been excluded from this report since they contain no new information.

5 CEILING RESULTS FROM CROSS-VALIDATION ANALYSES

For this analysis, METAR observations and NCV analyses of ceiling (again, at locations between the stations included in the product) are compared. In the majority of cases, about 6.6 million, the ceiling heights observed from the METAR reports and analyzed by the NCV are "unlimited". Since these correctly identified non-event cases are the least interesting and are difficult to analyze since "unlimited" is not a numeric value, they are excluded from this analysis. Only cases that have measurable NCV or METAR ceilings are examined, resulting in over 4 million cases. However, this large sample size is often difficult to examine, so a random selection of 10,000 cases was used for some of the analyses and displays.

Ceiling observations are censored at 20K ft. Thus, any observation or forecast for ceilings above 20K ft is set to 20K ft. Censoring prevents large but meaningless differences, say between 25K ft and 35K ft, from overwhelming the analyses. Furthermore, when instruments are used to measure ceiling, the ceiling height is often capped, which is not the case with a human observed ceiling. By censoring the data, the instrument and human observations are more likely to be consistent.

A histogram of errors in the ceiling field (METAR – NCV) is shown in Figure 4. The great majority of the errors are small, typically less than one thousand feet. However, the negative skew in the histogram indicates that the NCV analysis is more likely to indicate the ceiling is too high than too low. In other words, the NCV product is biased toward higher ceilings than are observed. The average error is -1417 feet while the median error is -348 feet, also indicating that the NCV product typically tends to produce somewhat higher ceilings than are observed.

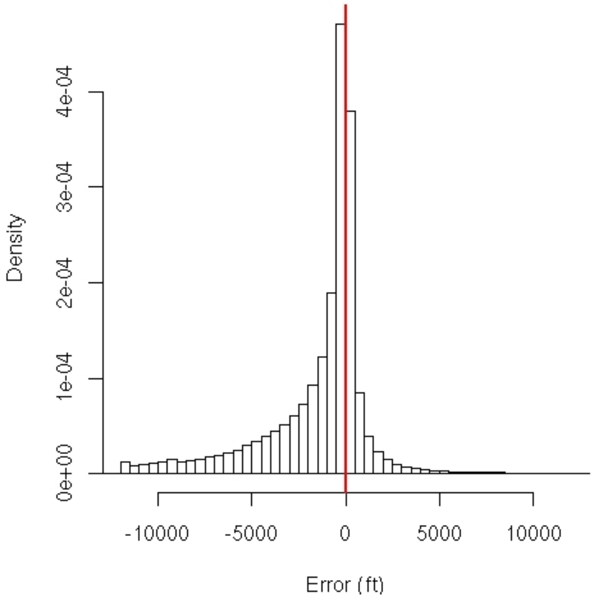

Figure 4: Histogram showing error in ceiling height (METAR – NCV).

A quantile-quantile (qq) plot shown in Figure 5 compares the distributions of NCV and METAR ceiling fields on a log-log scale. This type of plot shows the relationships between the overall distributional characteristics of each variable (e.g., the range, variance) rather than characteristics of their individual differences. The vertical stacks of values on the lower left of the plot, near the origin, are due to the discreteness of the METAR ceiling values, which is especially noticeable near the surface on the log scale. Multiple NCV ceiling points match each discrete METAR ceiling value at those levels.

If the distributions of these two fields were identical, all points would fall along the one-to-one line. Instead, the points are shifted almost linearly above the line. This result indicates that (in the original scale) the distribution of the NCV ceiling field is approximately the same as the distribution of the METAR ceiling field except that it is shifted higher (by about 0.65, the intercept of the linear model shown in the figure) in the log-log scale. Thus, in the original scale, the NCV ceiling distribution is approximately the same as the METAR ceiling distribution times 1.9 ($e^{0.65}$).

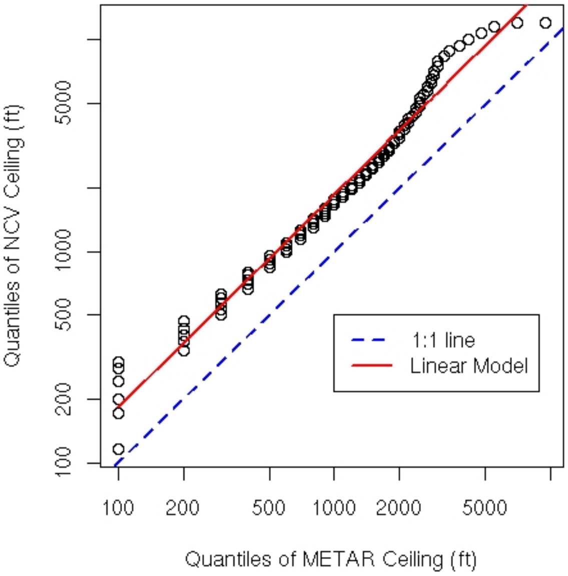

Figure 5: Quantile-quantile plot showing distribution of NCV vs. METAR ceiling values on log-log scale.

At each end of the distribution, the points are somewhat non-linear. At the bottom end, this is probably due to the discreteness of the METAR measurements rather than any real difference in the distributions of ceiling values. However, at the top end, the departure from linearity implies that the difference in NCV and METAR ceiling values are even larger than would be expected based on the estimated shift of the rest of the distribution (i.e., the NCV product is even more biased in the higher ranges).

Figure 6 shows boxplots[3] of the NCV ceiling versus the observed METAR ceiling for cases in which at least one of the METAR or NCV ceilings were less than unlimited, stratified by the NCV analysis ceiling value. Uneven ranges of NCV ceiling values were used in this plot for two reasons. First, operationally, it is more important to distinguish between lower ceiling values than higher, so the lower ranges are smaller and the higher ranges are larger. Second, the great majority of the ceiling measurements are concentrated at lower levels, so the boxes representing the lowest three altitude ranges represent approximately the same number of cases in spite of their differing altitude ranges. Some outliers extend up to 20K ft, where the measurements are censored.

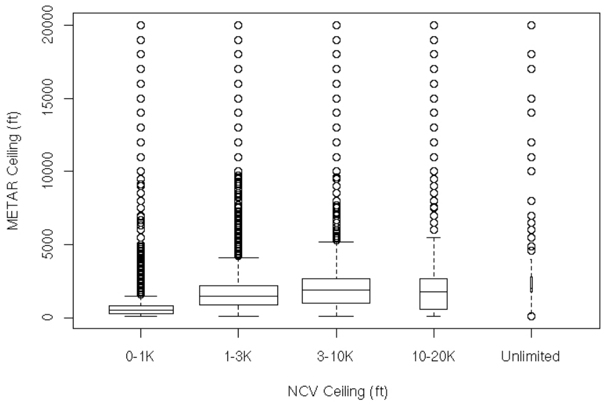

Figure 6: Boxplots of METAR ceiling below 20K ft by range of NCV ceiling (ft).

With over 4 million observations used to create this graphic, the outliers represent a very small number of cases; thus, it is informative to focus on the bulk of the observations (i.e. the boxes) below 6K ft. A similar graphic, with the outliers removed, is shown in Figure 7. Ideally, the boxes should be centered along the diagonal line from the lower left corner to the upper right corner. Although the boxes do not follow the diagonal, the first three do increase from left to right. The box for the NCV ceiling measurements of 10-20K ft generally corresponds to lower METAR ceiling values. The last box, for unlimited NCV ceiling values, is based on too few cases to be meaningful. However, typically when the NCV ceiling value was unlimited, the associated METAR value was also unlimited.

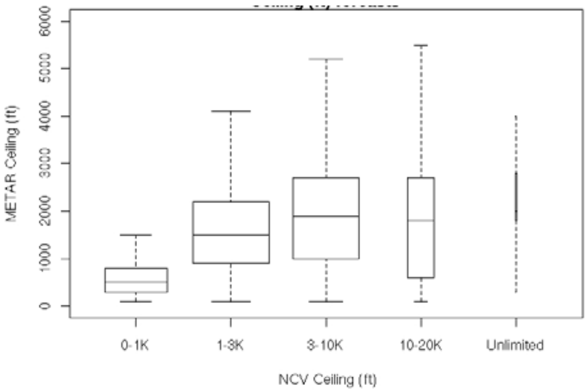

Figure 7: Boxplots of METAR ceiling below 6K ft by range of NCV ceiling (outliers eliminated).

A contour plot of the density of a random sample of 100,000 ceiling observations below 5,000 ft is given in Figure 8. Areas with a great number of points are shaded in warmer colors. Cooler colors indicate areas with fewer points. This plot is an alternative to a scatter plot. In a scatter plot, the areas with warm colors would be an indecipherable mass of points; the blue areas would have some points and the purple areas would be nearly empty. Ideally, warm colors should fall along the one-to-one line (in red) with cooler colors filling the remaining areas of the plot, which would indicate a good correspondence between the NCV analysis ceiling values and the observed ceilings provided by the METARs. Indeed, this is nearly the case as shown by the warm colors located along the diagonal. To the upper left of the one-to one line, there is an area with a small group of points in dark blue. These points are those for which the NCV product gave slightly higher ceiling values than were observed by the METAR.

---

[3] Box plots show the distribution of values. The line at the center of each box is the median, while the top and bottom of the box represent the 75th and 25th percentiles, respectively. Thus, the box shows the range of the center half of the data. The whiskers extend to the maximum and minimum values that are not outliers, each showing the range of the top and bottom quarters of the data. The dots above or below the whiskers are outliers. The width of each box is scaled to the number of cases represented by that category. Thus, narrower boxes represent fewer cases than wider boxes.

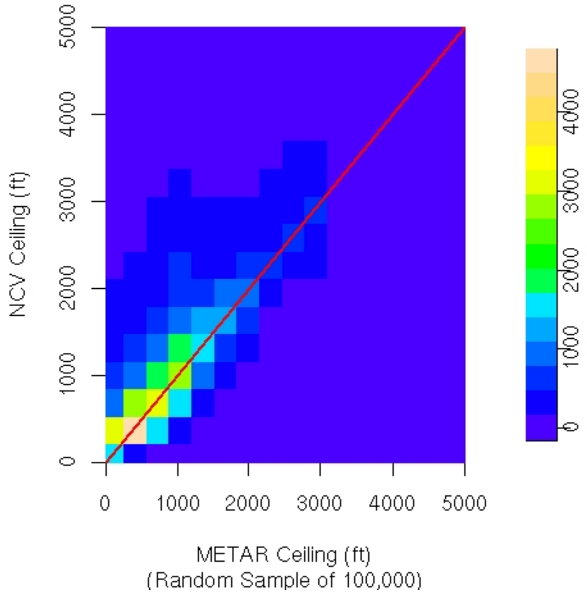

Figure 8:Contour plot showing density of METAR and NCV ceiling pairs.

6 VISIBILITY RESULTS FROM CROSS-VALIDATION ANALYSES

This section presents a comparison of METAR and NCV visibility values, again for locations representing the interpolation points between METAR stations. Once again, more than 6 million cases where both the METAR and NCV reported "unlimited" visibility were excluded from this analysis. Visibility measures are censored at 10 miles, as visibility greater than ten miles is essentially considered unlimited.

A quantile-quantile (qq) plot, provided in Figure 9, compares the distributions of the NCV analyzed visibility and observed visibility from the METAR. If the distributions of these two measures are the same, then the points on this plot will fall along the one-to-one line. Although the discreteness of the METAR measurements makes this nearly impossible, the two distributions are very similar as most of the points fall near to the one-to-one line. The slope of the linear model fit to these points is about 0.8, not quite the slope of one that would indicate perfect agreement. The NCV visibility field somewhat overestimates visibilities on the lower end and underestimates them on the higher end, resulting in a slightly narrower distribution of values than is observed. This is fairly common behavior when measurements are created using linear methods, such as those used to derive the NCV visibility field.

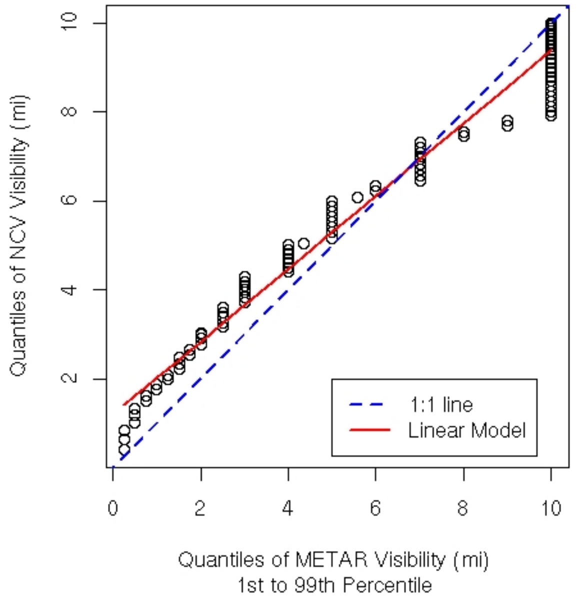

Figure 9: Quantile-quantile plot showing relationship between distributions of METAR and NCV visibility fields.

The histogram in Fig. 10 shows the relative frequencies of the visibility errors (METAR – NCV) for cases in which at least one of the two visibility values (METAR or NCV analysis) is less than unlimited. The great majority of the errors in the visibility field are small, less than 1 mile. Further, the errors appear approximately symmetrical, indicating that the NCV is relatively unbiased (i.e., it is not likely to consistently over- or under-estimate the visibility, overall).

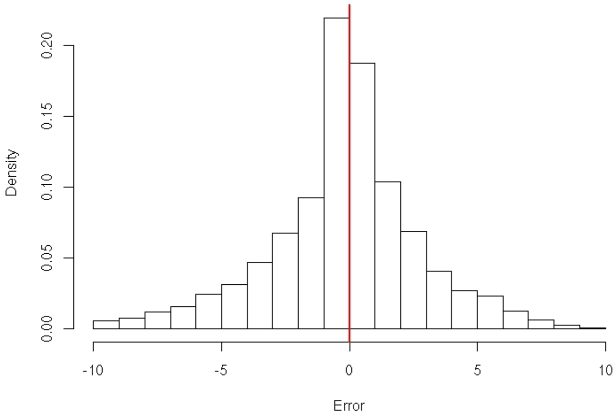

Figure 10: Histogram of errors in the NCV visibility field (METAR – NCV).

Figure 11 shows boxplots of NCV and METAR visibility. The center lines of the boxes (i.e. the medians) tend to fall along the diagonal line from the lower left to the upper right. This result indicates that the NCV visibility analysis roughly corresponds to the measured METAR visibility. The spread, or variability, as measured by the height of the boxes increases as the visibility increases. This common behavior indicates that the uncertainty in the NCV visibility increases as the NCV visibility value increases (i.e., smaller values are more certain than larger values).

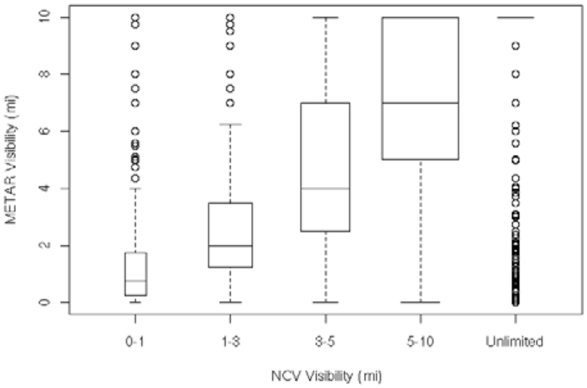

Figure 11: Boxplots showing METAR visibility values for categories of NCV visibility.

Figure 12 shows a similar plot to Figure 11, with the axes switched. Again the centers of the boxes tend to increase from left to right as they should, indicating that typically, the observed and analyzed data agree well. However, the spread (i.e., variability) of the NCV visibility does not increase as the observed visibility value increases, it stays about the same. Thus, the confidence interval around the NCV visibility value is about the same regardless of whether a user observes high or low visibility conditions.

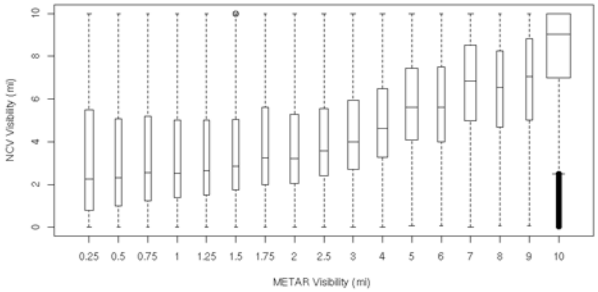

Figure 12: Boxplots of NCV visibility (mi.) by METAR visibility category (mi.)

7 SENSITIVITY OF STATISTICAL RESULTS TO CROSS-VALIDATION

The analyses presented in all previous sections used a cross-validation methodology with a roughly 300/1300 split of the stations into testing and training sets. If too many stations are included in the testing set, then the product may suffer. In order to determine if the selected 300/1300 split was “too much” a sensitivity analysis was undertaken. In this analysis, statistics for a single run of the 300/1300 split were produced. (Analyses in other sections of this paper repeat this analysis 10 times and accumulate the testing stations.) Additionally, two other testing/training set sizes were selected and the verification statistics for each produced. For the first, the verification is completed twice using only half of the 300 stations for testing. This resulted in two sets of data with a 150/1450 split. For the second, half of each of the previous two sets are kept for testing and each used to verify the NCV analysis produced with the remaining stations, yielding 4 sets of data with a 75/1525 split.

The results from the analysis of 150 and 75 holdout files were grouped to give 3 data sets each verified on 300 stations. The difference is that the NCV analysis product was produced using either 1300, 1450, or 1525 stations. The sets are also described in Table 5.

Table 5: Number of METAR stations used in cross-validation sensitivity analysis.

Training Set Stations	Testing Set Stations	Accumulated Testing Set Stations	Resulting Data Set
1300	300	300	Set 1
1450	150	300	Set 2
1450	150		
1525	75	300	Set 3
1525	75		
1525	75		
1525	75		

Verification statistics for all three sets of these 300 stations are presented in Table 6. The NCV product fares as well when produced using information from 1300 stations as it does when it is produced using information from 1525 stations. The statistics change almost imperceptibly from larger to smaller training set sizes. The maximum change is in the POD statistic, which changed by 0.008, from 0.765 to 0.773.

Table 6: Verification statistics for cross validation sets of size 75, 150 and 300.

	POD	POD No	FAR
Set 1	0.765	0.910	0.291
Set 2	0.770	0.908	0.294
Set 3	0.773	0.907	0.294

The verification statistics do not appear to be sensitive to the size of the cross-validation testing set, at least for sets comprising 20% of the data or less.

8 DISCUSSION AND CONCLUSIONS

This study used a cross-validation approach to evaluation the performance of the NCV analysis algorithm at interpolated locations between METAR stations. Overall results indicate that the algorithm is skillful at these locations, with somewhat varying performance depending on the component being evaluated.

The analyzed NCV visibility field closely matched the observed METAR visibility at all levels, as indicated by the small errors in the results and matching distributions.

Overall, the analyzed NCV ceilings matched well with the METAR ceilings, especially when ceilings were unlimited or below 10K ft. However, the NCV analyzed ceiling field is biased, producing higher ceiling values than are typically observed. For instance, when the NCV ceiling field is between 10 and 20K ft, the matching METAR ceiling was often below 5K ft.

The flight category verification statistics are somewhat mixed. As shown by the bias, the NCV analyses under-identify IFR events. The false alarm ratio is very low and the product does a good job of detecting events and a great job of detecting non-events.

With these caveats, the NCV analysis product shows positive skill in identifying IFR conditions and ceiling and visibility values, and thus it shows promise for future use as an operational tool.

ACKNOWLEDGEMENTS

NCAR is sponsored by the National Science Foundation. 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 of the FAA.

REFERENCES

Benjamin, S. J. and G. A. Grell, S. S. Weygandt, T. L. Smith, T. G. Smirnova, B. E. Schwartz, D. Kim, D. Devenyi, K. J. Brundage, J. M. Brown, and G. S. Manikin, 2001: The 20-km version of the RUC. 14th Conference on Numerical Weather Prediction, Ft. Lauderdale, FL. Amer. Meteor. Soc., Boston

Bradley, J.J. and Imbembo, S.M., 1985: "Automated Visibility Measurements for Airports," American Institute of Aeronautics and Astronautics: Preprints, Reno, NV.

Neter, J., M. H. Kunter, C. J. Nachtsheim, and W. Wasserman, 1996: Applied Linear Statistical Models. McGraw Hill / Irwin, Chicago, Ill.

NWS, 1991: National Weather Service Operations Manual, D-22. National Weather Service. (Available at Website http://wsom.nws.noaa.gov).

Petty, K., A. Bruce Carmichael, Gerry M. Wiener, Melissa A. Petty, and Martha N. Limber, 2000: A fuzzy logic system for the analysis and prediction of cloud ceiling and visibility. Ninth Conference on Aviation, Range, and Aerospace Meteorology. Orlando, FL. September. Amer. Meteor. Soc., Boston, 331-333.

Takacs, A., L. Holland, M. Chapman, B. Brown, T. Fowler, and A. Holmes, 2004: Graphical turbulence guidance 2 (GTG2): Results of the 2004 post analyses. Quality Assessment Report. Report to the FAA, October 2004.

U.S. Dept. of Commerce, NOAA, 1992: ASOS Users Guide, Government Printing Office

Wilks, D., 1995: Statistical Methods in the Atmospheric Sciences. Academic Press, San Diego.