Automated Extraction of Weather Variables from Camera Imagery [1]

Robert G. Hallowell & Michael P. Matthews
MIT Lincoln Laboratory
244 Wood Street
Lexington, MA 02420
Voice: 781-981-3645
bobh@ll.mit.edu

Paul A. Pisano
U.S. Federal Highway Administration - Road Weather Management Program
400 7th Street, S.W.
HOTO-1 Room 3408
Washington, D.C. 20590 USA
Voice: 202-366-1301
paul.pisano@fhwa.dot.gov

ABSTRACT

Thousands of traffic and safety monitoring cameras are deployed or are being deployed all across the country and throughout the world. These cameras serve a wide range of uses from monitoring building access to adjusting timing cycles of traffic lights at clogged intersections. Currently, these images are typically viewed on a wall of monitors in a traffic operations or security center where observers manually monitor potentially hazardous or congested conditions and notify the appropriate authorities. However, the proliferation of camera imagery taxes the ability of the manual observer to track and respond to all incidents. In addition, the images contain a wealth of information, including visibility, precipitation type, road conditions, camera outages, etc., that often goes unreported because these variables are not always critical or go undetected. Camera deployments continue to expand and the corresponding rapid increases in both the volume and complexity of camera imagery demand that automated algorithms be developed to condense the discernable information into a form that can be easily used operationally by users.

MIT Lincoln Laboratory (MIT/LL) under funding from the Federal Highway Administration (FHWA) is investigating new techniques to extract weather and road condition parameters from standard traffic camera imagery. To date, work has focused on developing an algorithm to measure atmospheric visibility and prove the algorithm concept. The initial algorithm examines the natural edges within the image (the horizon, tree lines, roadways, permanent buildings, etc) and performs a comparison of each image with a historical composite image. This comparison enables the system to determine the visibility in the direction of the sensor by detecting which edges are visible and which are not. A primary goal of the automated camera imagery feature extraction system is to ingest digital imagery with limited specific site information such as location, height, angle, and visual extent, thereby making the system easier for users to implement. There are, of course, many challenges in providing a reliable automated estimate of the visibility under all conditions (camera blockage/movement, dirt/raindrops on lens, etc) and the system attempts to compensate for these situations. This paper details the work-to-date on the visibility algorithm and defines a path for further development of the overall system.
Key words: weather—camera—visibility—extraction

INTRODUCTION

The first video surveillance cameras used for surveying traffic congestion were deployed in England in the mid-1970s and in the United States in the 1980s by the New York City Police Department. These early analog camera images were useful for general viewing but often the images were too washed out to even determine general traffic conditions. The advent of digital camera technology, high bandwidth wired and wireless communications, dramatic reductions in camera costs, have combined with increased funding of Intelligent Transportation Systems (ITS) infrastructure and a heightened need for security surveillance of wide areas due to post-9/11 terrorist concerns have fueled an explosion in available camera assets. Today, there are thousands of traffic and safety monitoring cameras deployed or being deployed across the country. Video cameras are being employed for a wide range of uses from monitoring building access to adjusting timing cycles of traffic lights at clogged intersections. Many of these traffic related cameras are available on the web with near real-time access, and are located at key traffic bottleneck or hazard points across the country. There are currently over 4000 State/City Department of Transportation owned and publicly available traffic cameras across the US (Figure 1), and there are many state, county and municipal centers planning more installations (USDOT, 2004).

The DOT owned cameras are located along major arterial roadways, key traffic intersections, remote intersections in mountain passes, and along side Road Weather Information System (RWIS) sites. The cameras are used to monitor road conditions for snow and blowing snow, monitor and adjust, where possible, traffic flows and verify RWIS observations and external road condition reporting. Currently, much of this monitoring is done manually. Images are typically viewed on a wall of monitors in a traffic operations or security center. Image processing applications for traffic management are focused mainly upon arterial management, to detect the presence of stopped vehicles at signalized intersections. Often the imagery is also displayable from a web-based map allowing commuters and commercial vehicle operators to view current traffic and weather conditions. However, it is generally left to observers to monitor potentially hazardous or congested conditions and to notify the appropriate agency.
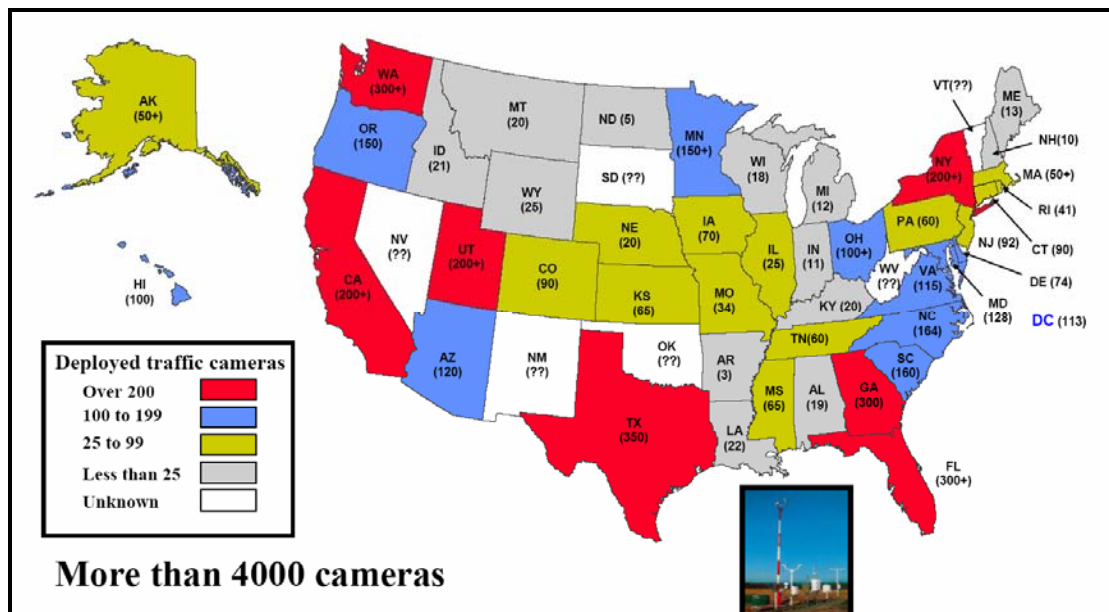

Figure 1. Survey of state DOT traffic cameras (combination of web survey and USDOT, 2004)

PROBLEM STATEMENT

The proliferation of camera imagery taxes the ability of the manual observer to track and respond to all incidents. In addition, the images contain a wealth of information, including visibility, precipitation type, road surface conditions, etc., that often goes unreported because these variables are not always critical to operations or go undetected by the observer. Camera deployments continue to expand and the corresponding rapid increases in both the volume and complexity of camera imagery demand that automated algorithms be developed to condense the discernable information into a form that can be easily used operationally by users. Recently, a number of companies have stepped forward to examine ways in which automated image processing might assist security and safety officials. Several companies offer automated license plate detection and reading for automated toll-way and red light enforcement, still others analyze images for security breaches and a few companies have begun to use video for traffic incident and flow monitoring. However, the area of weather and road condition analysis from video imagery is relatively new. The Japanese meteorological institute has done some work using image analysis to determine road conditions (Yamada, 2001). Similar road condition studies using neural networks and IR cameras have been performed by the Swedish National Road Agency (SNRA, 2002). The University of Minnesota has performed visibility tests utilizing fixed distance targets (Kwon, 2004). However, most of these programs require either new hardware (IR cameras and fixed sign placements for example) or extensive site surveys to determine object distances and land types.

MIT Lincoln Laboratory (MIT/LL) under previous funding from Department of Defense (DOD) and under new funding from the Federal Highway Administration (FHWA) is investigating new techniques to extract weather and road condition parameters from standard traffic camera imagery without additional hardware, signage or site surveys. This paper examines the work accomplished to date, including results from an early prototype algorithm, and discusses future efforts under the FHWA *Clarus* initiative, discussed later, to develop the algorithm further.

HISTORICAL WORK

The initial work on camera imagery for weather sensing was performed for the Army as part of a program to gather real-time weather data in the battlefield to support ground operations. The Weather Web (WxWeb) program was part of the overall Smart Sensor Web program, which also had a component called Image Web. Image Web envisioned deploying digital cameras to use in a military setting to monitor movement of enemy forces. MIT/LL was tasked with evaluating the use of camera images for weather surveillance. As a first step, MIT/LL deployed two fixed digital cameras, at different times, to two field site locations. The cameras were co-located with meteorological sensors for measuring, temperature, dew point, pressure, wind speed and direction, and visibility. The primary focus of WxWeb was mountainous terrain and fog conditions, thus initial algorithm development focused on visibility restrictions due to fog. From this initial work, a conceptual algorithm was created for automatically processing incoming images to derive a variety of weather variables (a patent is pending for the technique described below; (MIT, 2002)). The images and data collected during the WxWeb program provide the basis for the analysis presented here. Future funding from the FHWA under the *Clarus* research initiative will be used to generalize and extend the current algorithm.

ALGORITHM GOALS

The goal of this research is to develop an automated weather variable extraction system that utilizes existing visible camera technology. In addition, the user should only be required to enter rudimentary location (latitude, longitude, elevation) and viewing information (minimum/maximum viewing distance).

The initial focus is on daylight imagery, as nighttime imagery requires some ambient or fixed point source lighting and more extensive analyses that are outside the scope of this initial research.

The initial goal is to automatically estimate the overall visibility based on the characteristics of the image. Meteorological visibility is defined in several ways in the Glossary of Meteorology (AMS, 2005). The general definition is the furthest distance that a manual observer can see a defined reference object with the naked eye in a given direction. The manual observation reported is the *prevailing* visibility which refers to the median visibility gathered from around the horizon. Automated visibility sensors used on ASOS (Automated Surface Observing Systems) are designed to measure the prevailing visibility by assuming that the conditions between the sensor's receiver and transmitter represent the nominal conditions around the horizon. Since the actual visibility may not be homogeneous over the entire domain, it is quite possible that the visibility estimate of the laser sensor could differ from a manual observation. Similarly, visibility measured by a fixed camera and viewing angle may not be the same as either the manual or laser sensor. The camera will capture prevailing visibility in the direction the camera is pointed, we call this *directional* visibility. Directional visibility may differ significantly from the local laser sensor when the contaminant causing the visibility reduction is not present at the point of measurement. These situations might occur in the case of approaching rain/snow showers or cloud decks, fog banks growing or decaying, and other more localized atmospheric phenomena.

In addition to visibility, the algorithm is designed to incorporate other weather variable algorithms in the future (fog or precipitation detection and trend, road condition detection and trend, etc). Of equal importance will be the need to recognize when the camera image is impaired either due to hardware failure or external objects (dirt/precipitation on the lens or unexpected blockages from vehicles).

DATA GATHERING

Camera Locations

Two site locations were fitted with a mid-resolution (320x240) digital camera. The first site was located at MIT Lincoln Laboratory's Annex-2 facilities atop a 25-meter tower on Katahdin hill at Hanscom AFB in Lexington, MA. During testing, the camera collected images on a five-minute interval from February 2000 through March 2000. The primary camera view for Annex-2 was northwest over the airfield. Figure 2(a) is an image collected on a clear day with visibility in excess of 50 kilometers. There are several distinctive features in the camera image. First, the distant horizon is the small mountains located in southwest New Hampshire, all at distances greater than 50 kilometers. Second, Hanscom airfield is located in the center of the image, and is at a distance of 3.2 kilometers to the "far" side, and 1.5 kilometers on the "near" side. The smokestacks visible on the left side of the image are at a distance of 400 meters. Finally, the ASR radar located in the foreground is 60 meters from the camera. The image has significant clutter in the foreground from several trailers and other vehicles. The camera housing is also visible in the upper-left corner of the image.

The second camera was located atop an instrument shelter at the Harriman & West airport in North Adams, Massachusetts and operated capturing images at 1 minute intervals from March 2000 to March 2001. Although the camera was remotely controllable, the primary camera view is to the west of the airport over the Williamstown valley located along the Taconic Mountain Range. Figure 2(b) is an image collected from the Harriman camera on a clear day. There are several distinctive features in the camera image. First, the most distant horizon is the Taconic Range, located 6.7 kilometers to the west of the instrument shelter. Second, the airport's fuel storage facilities are located in the foreground at a distance of 42 meters. Third, the airport's hangar facilities are located in the center and the north (right) side of image. Two aircraft hangars are also visible, one behind the fuel storage facility (400 meters), and one to

the right and behind the facility (200 meters). On the east (left) side of the image are the runways, taxiways, and parked aircraft. Also visible on the left side of the image is a "near" ridgeline located approximately 2.1 kilometers from the camera. Finally, the camera's housing bracket is seen at the top of the image.
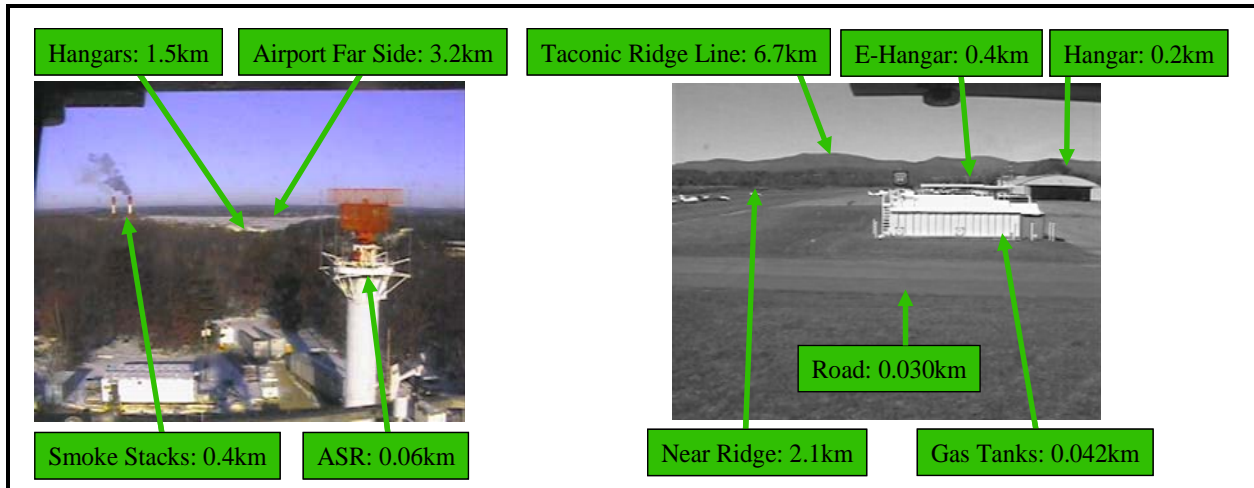

Figure 2. (a) Hanscom AFB camera view (left) and (b) Harriman and West camera view (right)

Meteorological Data

Weather data was gathered from sensors co-located with each camera. Standard measurements of temperature, dew point, pressure, wind speed and direction and visibility were gathered continuously over the test period. Of key interest for verification in this study were visibility measurements gathered using a Vaisala FD12P laser to estimate visibility by analyzing the scatter of the laser beam. The FD12P makes for an excellent automated method of generating standard meteorological visibility estimates. However, as mentioned above, the FD12P produces an estimate of the visibility using a small spatial sample (Figure 3) as the distance from the laser transmitter to the receiver is only three feet.

Figure 3. Vaisala FD12-P laser visibility sensor

EXAMPLES OF TEST SITE IMAGES

Figure 4 depicts several examples of images from Harriman with fog at the airport. Figure 4(a) is from August 31, 2000 at 11:00Z, the visibility at this time was reported at 130 meters. Figure 4(b) is from August 21, 2000 at 12:00Z, visibility is estimated at 450 meters. Figure 4(c) is from August 7, 2000 at 12:00Z, visibility is estimated at 3,500 meters. Figure 4(d) is from August 16, 2000 at 12Z, visibility is estimated at 9,200 meters. Several important things can be observed in this series of images. First, the clarity of the foreground objects (most notable the fuel storage facility) improves as the visibility increases. In Figure 4(a) the fuel facility is visible; however, it appears fuzzy. In Figure 4(b) the fuel facility is visible and the edges are more distinct than in 4(a). In Figure 4(c) the image of the fuel facility is appears sharp, but the more distant Taconic ridge line is missing. Conversely, in Figure 4(d) the

foreground buildings and the distant ridge line can be seen clearly. The background image also changes from the lowest to highest visibility images. In the first image, the background (grassy areas and pavement) has very little texture, but by Figure 4(d) the background has more texture, and is of a lower grayscale value. We can also observe in the last image the presence of rain drops on the camera enclosure. This data quality issue presents a problem for weather algorithms and is addressed in the algorithm development section.
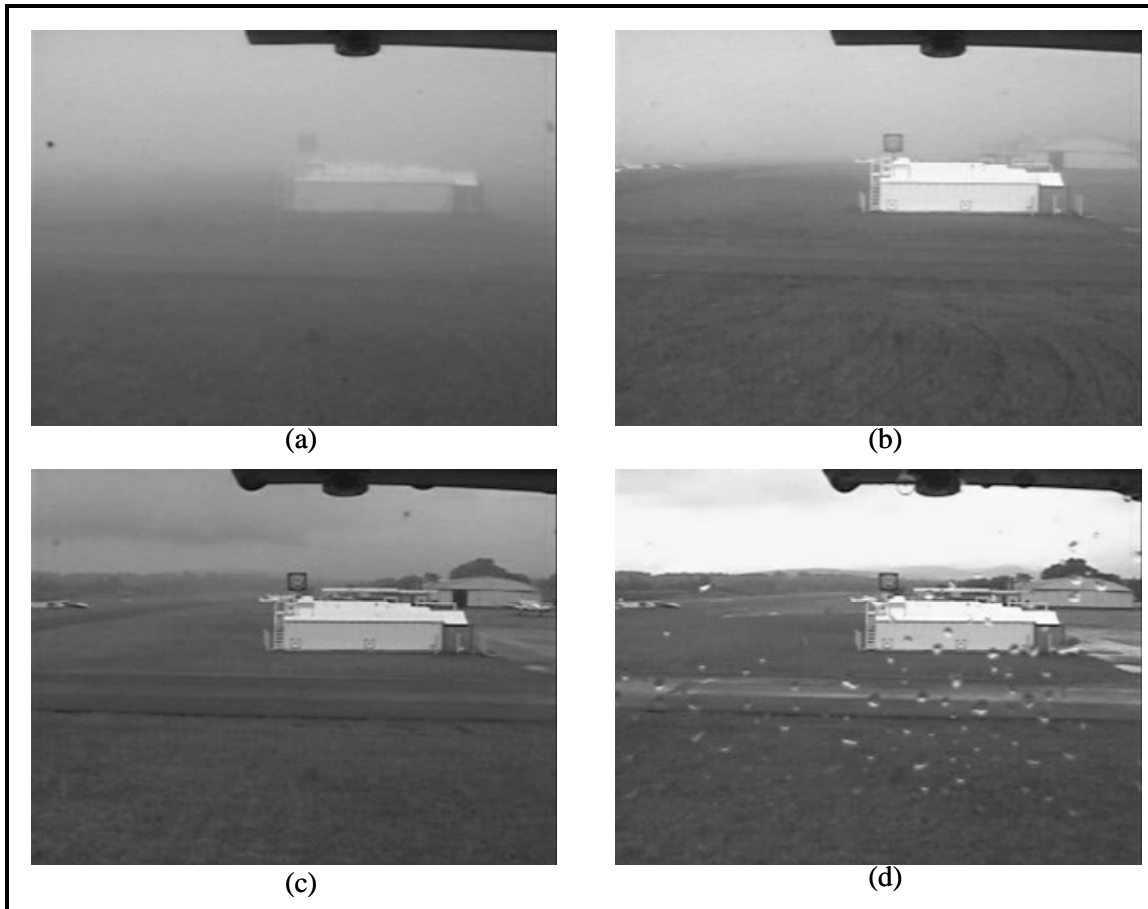

Figure 4. (a) August 31, 2000 at 11Z, visibility of 130 meters (b) August 21, 2000 at 12Z, visibility of 450 meters (c) August 7, 2000 at 12Z, visibility of 3,500 meters (d) August 16, 2000 at 12Z, visibility of 9,200 meters

Figure 5 depicts several examples of images from Annex-2 with fog present. In the first image, Figure 5(a), the visibility is less than 100 meters. This image is from February 11, 2000 at 20:00Z. The ASR radar located 60 meters from the camera is faintly visible, almost blending in with the fog. The second image, Figure 5(b), is from February 18, 2000 at 20:00Z. The visibility at this time was just over 1,000 meters. Most notable in this image is that the smokestacks are now visible and a ridgeline is now visible that was not detectable in the clear day image. Figure 5(c) is from February 27, 2000 at 17:00Z. The visibility at that time was just over 5 kilometers. Now the airfield is visible and the far side of the airfield is faintly visible. The final image, Figure 5(d), is taken on February 20, 2000 at 14:00Z. The visibility is greater than 10 kilometers at this time. Snow cover is plainly visible in Figure 5(b) and 5(d), while in Figure 5(b) snow is actually falling.
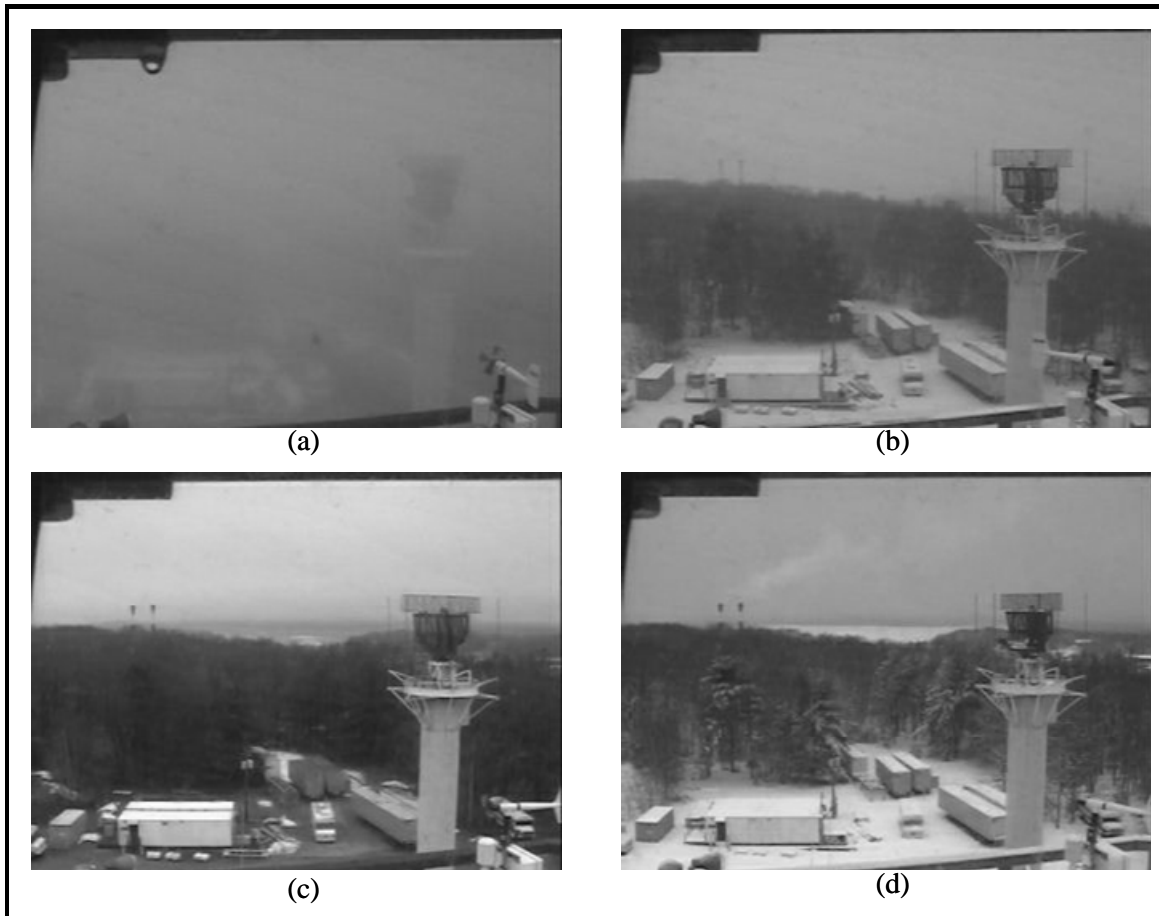
Figure 5. (a) February 11, 2000 at 20Z, visibility less than 100 meters (b) February 18, 2000 at 20Z, visibility of 1,000 meters (c) February 27, 2000 at 17Z, visibility of 5 kilometers (d) February 20, 2000 at 14Z, visibility of 10 kilometers

ALGORITHM DEVELOPMENT

Defining Features

One of the primary goals of the envisioned algorithm is that it be easily deployed in a variety of environments with little manual site setup. As such, it is better for the algorithm to rely on overall features in a subject image rather than explicit knowledge of the distance to various objects. Therefore, the core of the algorithm is based on analyzing the entire image itself and edge features within the entire image. The images in Figure 6 show the clear image from Figure 2(b) two and the heavy fog image from Figure 4(a), but with an edge extraction technique applied. Edges both near and far can be seen quite clearly in the high visibility case on the left. In the image on the right, however, the furthest edges that can be seen are those related to the gas tanks and associated building located some 42 meters from the camera. The hangar sitting at 200 meters is completely obscured. Indeed the laser measured visibility for this image is 130 meters. Similarly, Figure 7 shows the same edge losses in the Hansom AFB images from Figures 2(a) and 5(a). The nearest edges in these images are those from the radar tower some 60 meters from the camera. No other edges can be seen in the low visibility image on the right and the measured laser visibility in this case is approximately 100 meters.

Figure 6. Edge extracted images for the Harriman airport from a clear (>20km visibility) day on the left and a low visibility (130 meter) day.
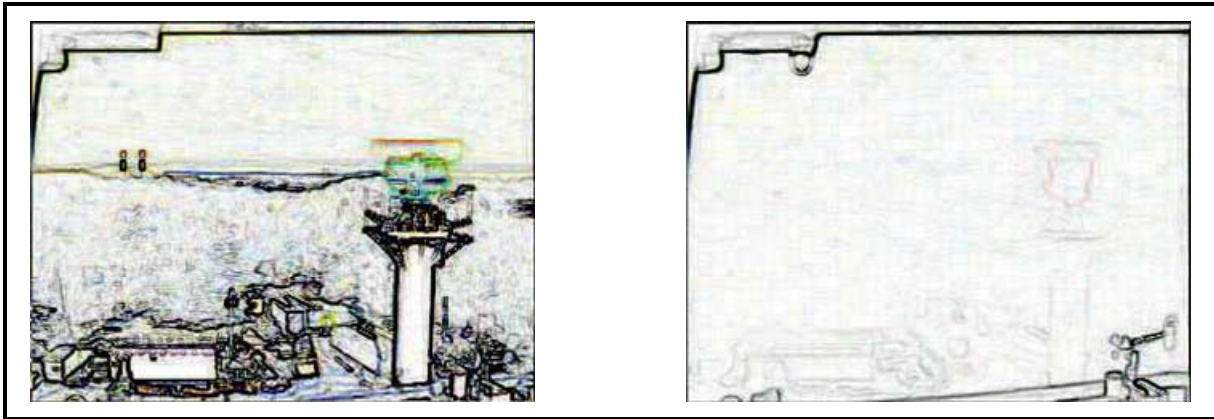

Figure 7. Edge extracted images for the Hanscom AFB camera from a clear (>40km visibility) day on the left and a low visibility (< 100 meter) day.

Based on reviewing dozens of low visibility events it was clear that finding a way to correlate edge loss with visibility was a concept worth exploring. A clear image contains a full set of "expected" edges, these are the strong edges associated with buildings, trees, the horizon, roads, etc. As visibility decreases, fewer and fewer expected images are visible, and the loss of these edges occurs from the furthest edge to the closest as visibility approaches zero. Determining expected edges is accomplished by maintaining a composite image compiled from an historical average of daylight imagery within the system. In addition, constant but weak edges are also removed from the composite image leaving only high signal edges that should be found in any clear image. In each image, of course, there are "unexpected" edges; these are edges associated with traffic, dust/water on the camera lens, snow piles, and other varying phenomena. Figure 8 illustrates the concept of separating expected and unexpected edges within the system. Composite edges are shown in the upper left, a building shaped edge near the bottom with an average weighting of 0.8 (on a hypothetical scale of 0.0-1.0 edge strength, with 1.0 being a strong edge) and a horizon edge with an average weighting of 0.5. Weaker edges (below some threshold, in this example 0.5) are removed from the composite image. The current edges in the lower left represent the edges from an incoming image, in addition to the "expected" edges seen in the composite image there are "unexpected" edges from transient objects (in this case rain drops on the camera shield). Expected edges are extracted from the current edge field, by finding matching edges within the composite edge field. The relative strength or weakness of expected edges as compared to the composite field is directly proportional to the reduction in visibility. Unexpected edges are strong edges (>0.5) that are not associated with a

corresponding composite edge. This illustration is conceptual, the system examines each pixel within an image to determine its' edge strength and while those strong pixels will make lines similar to the ones shown, the signal strength may vary significantly.
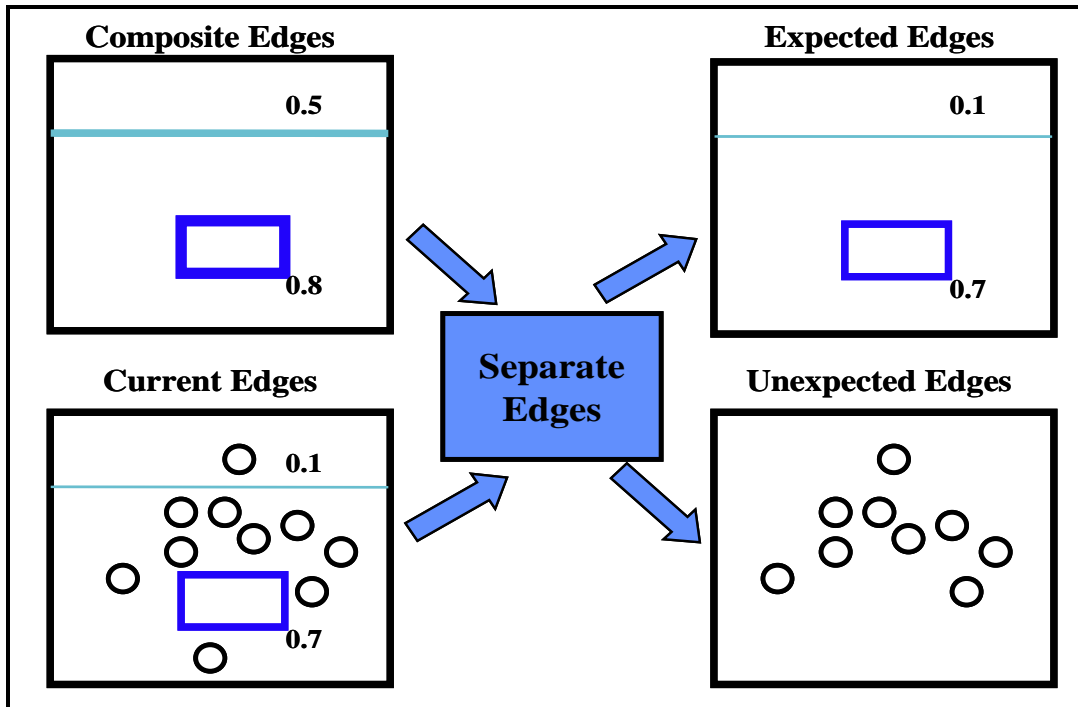

Figure 8. Illustration of edge analysis to separate strong, long term "expected" edges from strong, but transient "unexpected" edges.

Removing unexpected edges is crucial to calculating an accurate estimate of the true visibility. The example shown in Figure 9 shows how the system can effectively eliminate these unexpected edges. Both the top (Feb 28[th]) and bottom (Feb 21[st]) images on the left are days with visibilities > 10 km. However, the top image has a large number of raindrops on the camera shield. Edge detection (middle row) and expected edge extraction (right) relative to the composite image edges are performed on both input images. As you can see the image on the right are quite similar, the strong expected edges have been preserved, and both images yield algorithm visibility estimates > 10 km.

Analyzing Images

There are a multitude of edge detection algorithms and there are a variety of ways to quantify the strength of the edges found. The processing used for this analysis is the Sobel edge detection algorithm (Parker, 1997). This algorithm looks for discontinuities in the image and generates a scaling (0-255) that represents the intensity of the edge. The Sobel algorithm yields images like the ones shown in Figures 6 and 7. Another approach to image analysis is the Fast-Fourier Transform (FFT) also referred to as the power of an image. The FFT is an efficient algorithm to determine the frequency spectra of digital signals. The FFT will produce an image with the same dimensions as the original image. However, the FFT image would not be visually informative to the analyst. Once the FFT is computed, the image magnitude is computed by performing a summation over all relevant pixels in the image. This summation provides a single measure of the relative frequency amplitudes in the input and composite image. Typical low visibility images, for example those caused by fog, tend to wash out high frequency edges and therefore yield lower overall magnitudes than images on high visibility days.
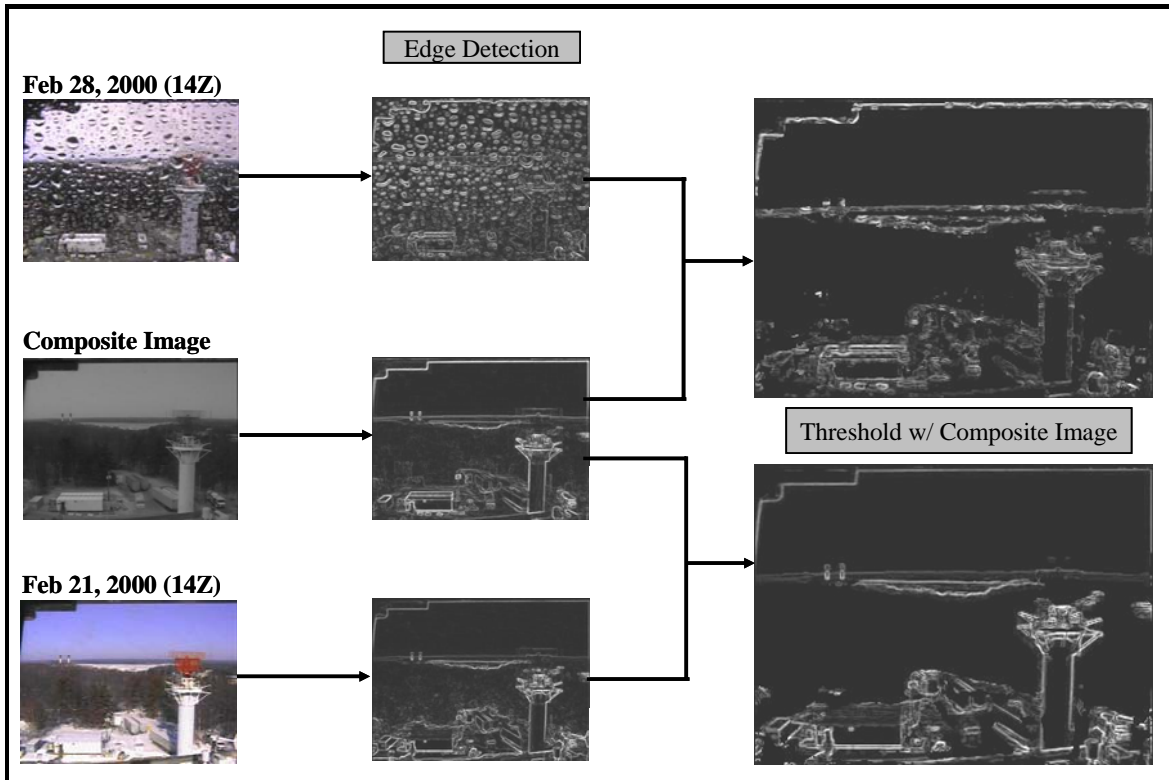
Figure 9. Example applications of data quality algorithms to remove transient edges.

For each camera image used in this analysis, the Interactive Development Library (IDL) was used to generate the Sobel edge detection. IDL was then used to compute the image grayscale mean, standard deviation, and power from both the original image as well as the expected edge detection image. For each image, five measures of the relative strength of the input image were calculated. The first measure, normalized image magnitude was calculated by calculating the magnitude of the input raw image and dividing it by the magnitude of the composite image. The next four measures were based on measures of edge strength as opposed to raw image pixel values. A normalization of each pixel's edge strength was performed first, by dividing the input image by the composite image edge strength. Three of the measures, the edge mean, edge median, and edge magnitude were based only on expected edges. The final determinant, total edge magnitude, was based on all strong edges (expected and unexpected) within the image. Figure 10 shows each of these measures for all daylight hours from the Harriman airport camera as compared to the true visibility as measured by the FD12P visibility sensor. There is good agreement between the normalized values for the mean of the edges ($r=0.70$), the median of the edges ($r=0.65$), and the magnitude of the entire image ($r=0.62$), the expected edges ($r=0.63$), and all edges ($r=0.77$). There are a few outliers and closer scrutiny of each of these cases reveals that, as noted earlier, they are due to valid differences from visibility sensor in the FD12P.

Initial Results

Rather than choose only one predictor for visibility, the algorithm uses an average visibility value as predicted by each of the predictors. This process, sometimes called fuzzy logic, often yields better results than any single predictor because the consensus reduces the impact of outliers. Figure 11 shows the algorithm results by using the fitted lines shown above as predictor functions. The visibility values are broken down into categories: less than 1km, 1-5 km and greater than 10km. The overall probability of a correct visibility categorization was 73.4%; however the crucial "less than 1 km" category was correctly

predicted 90.3% of the time. The worst performing category was "greater than 10 km", but, this range is often the area where the camera might see an incoming front, where as the lidar sensor can only estimated based on the local conditions. While these results are promising, they are based on a single camera and the low visibility comparisons are based on a small set of data.
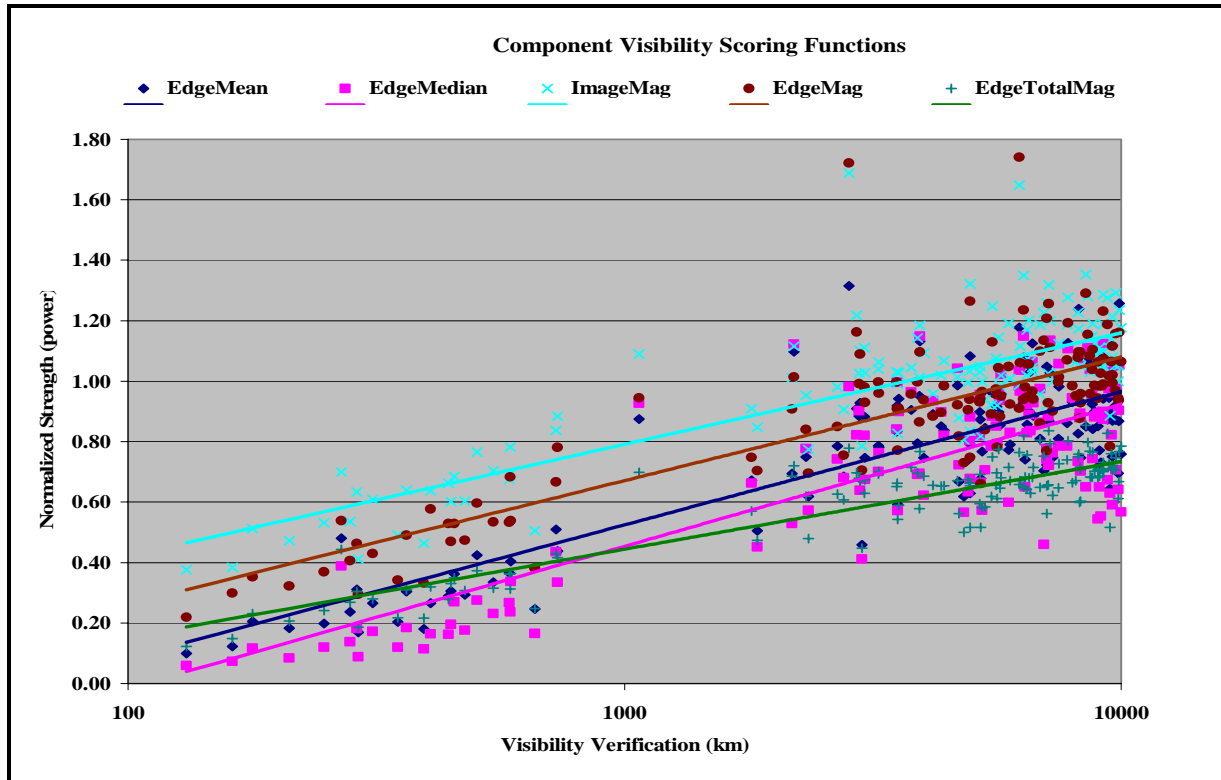

Figure 10. Comparison of edge and image normalized strength ratios to true visibility as measured by the FD12P lidar for all daylight images at the Harriman airport camera site (7/1/200 to 10/31/2000).

CLARUS RESEARCH INITIATIVE

Clarus (which is Latin for "Clear") is an initiative to develop and demonstrate an integrated surface transportation weather observing, forecasting and data management system, and to establish a partnership to create a nationwide surface transportation weather observing and forecasting system (www.clarusinitiative.org). Part of the *Clarus* charter is to investigate new technologies that have potential application to surface transportation weather. As such, FHWA will be funding a one year effort to further develop the automated camera visibility algorithm. While the algorithm detailed above shows promise, there are many technical difficulties to overcome. The *Clarus* research effort will have three main components: 1) survey the current camera usage by state DOTs, 2) Analyze and enhance the current algorithm utilizing MIT/LL controlled cameras, with particular emphasis on snowplow cab-height visibility, and 3) Extend the prototype to operate on a set of existing state DOT camera images.

The initial results were created from one camera location, using correlation curves defined from that same camera. A first step under the Clarus research will be to perform the same algorithm analysis shown above on the Hanscom AFB test images. This will provide valuable information on the types of modifications that may be needed to make the algorithm more generic. In addition, FHWA funding will be used to install a set of two camera installations at MIT/LL facilities. A survey will be performed to

determine the characteristic of the various state DOT camera installations. The MIT/LL cameras will be similar to those used by DOTs across the country. The heights of the two installations will be near-ground (5-15 meters) and above ground (25-50 meters) to mimic typical DOT installations. This tiered camera installation will allow us to analyze the correlation between standard height cameras and the visibility at "cab-height". Cab-height visibility is critical for snowplow operators so that they can be warned when road conditions are reaching zero-visibility conditions making operations too risky. Additionally, a lidar visibility sensor (the Vaisala FD12P) and solar radiometer will be co-located with the cameras along with standard weather data sensors. These test cameras will be used to test and tune the prototype camera visibility algorithm. The biggest challenge will be to develop a system that is transportable to a wide variety of imagery but with as little tuning and site surveying as possible. Once the system is improved and tuned, we will access and process standard DOT cameras that are available in a variety of cities and states.

		Video Estimated Visibility (km)			
		<1	1-5	5-10	≥ 10
Lidar (FD12P) Visibility (km)	<1	28	2	0	0
	1-5	1	24	11	6
	5-10	0	10	64	32
	≥ 10	0	46	174	663

Figure 11. Prototype visibility algorithm scoring results on daylight imagery for the Harriman airport camera test site 7/1 to 10/31/2000.

REFERENCES

AMS, 2005, *Glossary of Meteorology*, American Meteorological Society, Boston, MA. (book)

Kwon, T.M., et.al. 2004, *Atmospheric Visibility Measurements Using Video Cameras: Relative Visibility* Minnesota University, Minneapolis, MN and Minnesota Department of Transportation, St. Paul, MN. Report: CTS-04-03, Jul 2004. (CD-ROM)

MIT, 2002, *Video system for monitoring and reporting weather conditions*, US Patent application #20020181739, December 5, 2002. (patent application)

SNRA, 2002, *Final Report on Signal and Image Processing for Road Condition Classification*, AerotechTelub and Dalarma University under the Swedish National Road Agency. Report #2002-02-06. (report)

Parker, James R., 1997, Algorithms for Image Processing and Computer Vision, New York, John Wiley & Sons, Inc., 1997, pp. 23-29 (book)

USDOT, 2004, *ITS Deployment Tracking survey*, United States Department of Transportation (USDOT) Intelligent Transportation Systems (ITS) Joint Program Office (JPO), Washington, DC. http://itsdeployment2.ed.ornl.gov/its2004/Results.asp?ID=325&rpt=A&filter=1. (website)

Yamada, M., et.al., 2001, *Discrimination of the Road Condition Toward Understanding of Vehicle Driving Environments*, IEEE Transactions on Intelligent Transportation Systems, Vol. 2., No. 1, March 2001. (paper)