4A.7 AN AUTOMATED VISIBILITY DETECTION ALGORITHM UTILIZING CAMERA IMAGERY[†]

Robert G. Hallowell* and Michael P. Matthews
Massachusetts Institute of Technology Lincoln Laboratory, Lexington, MA

Paul A. Pisano
Federal Highway Administration, Washington, D.C.

1. INTRODUCTION

The Federal Highway Administration (FHWA) has had a focused program to improve the integration of weather decision support systems into surface transportation operations since 1999. *Clarus* (Latin for clear) is the FHWA's most recent surface transportation weather initiative. The *Clarus* concept is to develop and demonstrate an integrated surface transportation weather observing, forecasting and data management system (Pisano, 2006a). As part of this effort, the FHWA is also promoting research into methods for applying new and existing sensor or probe data. These efforts include utilizing new in-vehicle sensor data that will be part of the vehicle infrastructure initiative (VII) (Pisano, 2006b), and finding innovative ways to use existing camera imagery. MIT Lincoln Laboratory (MIT/LL) was tasked to evaluate the usefulness of camera imagery for sensing ambient and road weather conditions and the feasibility for creating a portable visibility estimation algorithm.

This paper gives a general background on the current utilization of camera imagery, including past and ongoing research of automated weather/condition algorithms. This is followed by a description of the MIT/LL camera test site, the analyses performed and the resultant prototype visibility estimation algorithm. In addition, the paper details application of the prototype algorithm to existing state DOT cameras in Utah. The final section discusses the future possibilities of camera-based weather and road condition algorithms.

2. BACKGROUND

Cameras have been used for decades to remotely monitor traffic and to protect life and property. The deployment and utilization of cameras has expanded dramatically in the last decade with support from the Department of Transportation for traffic and emergency management, "511" services, and monitoring of Intelligent Transportation Systems (ITS). In addition, the Department of Homeland Security has funded camera deployments in support of threat surveillance and emergency management operations. Camera sensors are particularly important for surface transportation applications because they directly observe the road/rail environment. Richard Cressey, a former counterterrorism official for the US Government, was quoted recently stating that "*there are about 30 million video surveillance cameras in the United States shooting about four billion hours of footage every week*" (Johnson, 2006). Four billion hours of imagery: that is like watching every minute of every day of your life and the lives of all your direct relatives up to and including your great-great-great-great grandparents, all in one week. And that number is growing every day.

Only a fraction of the surveillance cameras are designed to observe transportation assets. But, even if you just consider State Department of Transportation (DOT) and Traffic Management Center (TMC) camera assets there are over 10,000 cameras continuously monitoring major

[†]This work is sponsored by the Federal Highway Administration (FHWA) under Air Force Contract No. FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the author and are not necessarily endorsed by the United States Government.

[*]*Corresponding author address*: Robert Hallowell, MIT Lincoln Laboratory, 244 Wood Street, Lexington, MA 02420-9185; e-mail: bobh@ll.mit.edu

roadways in the US (USDOT, 2004, and surveys done in support of this research). When you have manual observers of transportation assets, the observations are generally focused on the primary tasks of traffic management or security. Ancillary information, such as weather or road conditions, are not routinely reported or archived. Clearly, manual observation of all video footage is impossible, and more and more information will be lost unless automated algorithms are developed to extract and report this valuable information.

2.1 Related research

Recently, a number of companies have stepped forward to examine ways in which automated image processing might assist security and safety officials. Most of the research and private sector involvement has focused on security surveillance: identifying terrorists, finding bombs, or other security concerns. In the area of traffic management a host of companies and associated research offer automated license plate detection for automated toll-way and red light enforcement, 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 still limited to research organizations.

Camera-based weather research has largely focused on road condition, driver-level visibility measurement and fog detection. The FHWA's successful project to create a Maintenance Decision Support System (MDSS) for winter maintenance operations cited the need for better road condition information to provide enhanced feedback into the system (Pisano, 2004). A wide variety of international researchers have been at the forefront of using camera imagery for meteorological analysis. The Japanese Meteorological Institute examined overall image characteristics for gauging road surface conditions during winter storms (Yamada, 2001). Similar road condition studies using neural networks and IR cameras have been performed by the Swedish National Road Administration (SNRA, 2002). The SNRA prototype is currently being evaluated throughout Sweden.

Previous MIT/LL research for the Department of Defense (DoD) (Clark, 2000) has shown that statistical edge analysis of camera imagery could be used to estimate the *meteorological visibility* (defined below) of a region. A follow on study of the DoD work illustrated that an automated algorithm for visibility using standard camera imagery was plausible (Hallowell, 2005). Recent studies by the Hokkaido Regional Development Bureau have shown that manual observers are able to consistently perceive the measured visibility (very poor, poor, fair, and clear) by a single camera image of a road scene (Nagata, et.al, 2006).

The Nagata research confirms that an automated algorithm that mimics the cues used by human observers should yield high-quality results. The University of Minnesota has performed successful visibility tests utilizing fixed distance targets and defined a visibility index called Relative Visibility to score such efforts (Kwon, 2004). Hagiwara (2002, 2004) performed visualization experiments in artificial fog in initial research and then utilized sub-images to measure image contrast. Algorithm results for the weighted intensity power spectra (WIPS) compared favorably to manually estimates of the test image's visibility. Finally, French researchers have used on-board cameras in vehicles to estimate visibility conditions (Hautiere, 2005). The work presented here is complimentary in nature to these other studies.

3. RESEARCH GOALS

The goal of this *Clarus*-sponsored research is to develop an automated weather variable extraction system that utilizes existing visible spectrum camera technology to estimate meteorological visibility for surface transportation. The goal focuses on visible imagery because the vast majority of cameras currently deployed as DOT sensors are visible-only cameras. State DOTs are primarily concerned when visibilities drop below 1 km and more critically, below 100

meters. However, the system should detect all ranges of visibility as the rate of change in visibility will be a key factor in any eventual visibility forecast system. Given the large number of camera assets, the algorithm should be developed in such a way as to maximize the ease and speed with which it can be deployed. As such, the user should only be required to enter rudimentary location (latitude, longitude, elevation) and viewing information (minimum/maximum viewing distance). Daytime imagery is the primary focus, as nighttime imagery requires more complicated analyses of ambient or fixed point source lighting. Finally, a secondary goal is to compare the visibility as seen by an elevated camera to that of a winter maintenance operator driving on a road. Zero visibility conditions are an extreme hazard for winter maintenance crews working to clear roads during winter storms.

4. MIT/LL CAMERA TEST SITE

Two cameras were installed at Hanscom Air Force Base (HAFB is co-located with MIT/LL in Lexington, MA) in support of algorithm development. COHU-3960 visible imagery cameras (Figure 1) were utilized in this project; this camera is a high-resolution digital camera providing full color imagery with full remote pan, tilt and zoom capabilities. Many State DOT cameras are installed on light poles or in other elevated positions.

Therefore, one camera was installed on top of the Lincoln Lab flight facility which has an elevation of approximately 30 meters (Figure 2, FF1). A second camera was installed at the base of the facility on a small tower at a height of 5 meters (Figure 2, FF2). This low level camera had two purposes: (1) zero visibility conditions in the truck cab are a major concern for snow plow operators and this camera was a surrogate for conditions as seen from the cab, (2) some DOT cameras are mounted on ground level stations or on bridge overpasses; processing the algorithm on the low-level camera would allow us to discover differences in how the visibility algorithm might work at varying heights.

Meteorological visibility is defined in several ways in the American Meteorological Society's (AMS) 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 Automated Surface Observing Systems (ASOS) 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 that of a manual observer. 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; this is called *directional* visibility (akin to what the Glossary defines as "sector 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, growing or decaying fog banks, and other more localized atmospheric phenomena.

Figure 1: COHU 3960 Series environmental camera.

Both cameras were positioned to look west over a series of buildings and access roads with a low ridge line in the near distance. Figure 3 (a) and (b) show representative views of the upper and lower camera views, respectively. Camera images were collected at one-minute intervals starting with the upper camera on February 10th, 2006 and the lower camera on March 16th, 2006. The images were captured at a resolution of 320x240 pixels. While below the maximum resolution of this camera, the reduced resolution is typical of State DOT camera imagery. Camera images continue to be collected to this day; however, the images used for analysis in this paper ended on May 31st, 2006.

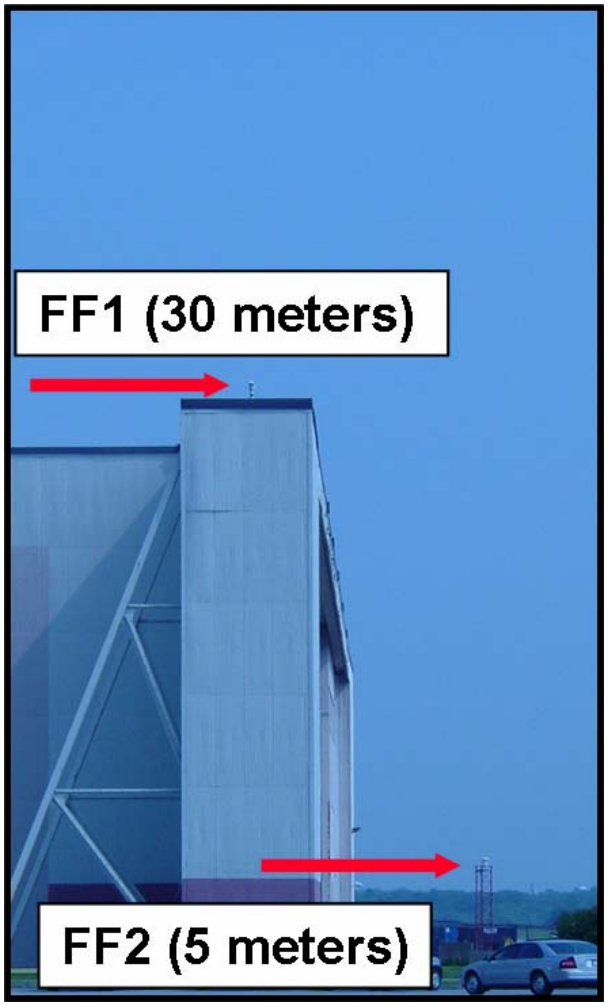

Figure 2: Field site tiered cameras mounted at the MIT/LL flight facility on Hansom Air Force Base, Lexington, MA

Figure 3: Example clear day images from the MIT/LL field site cameras; elevated, roof-top camera (a) on the top and low-level camera view on the bottom (b).

4.1 Meteorological Data

HAFB has an Air Force maintained ASOS within one kilometer of the test cameras. An example ASOS station is shown in Figure 4. Consequently, weather data were gathered from the ASOS sensor including measurements of temperature, dew point, pressure, wind speed and direction and visibility. Of key interest for verification in this study were visibility measurements gathered using the ASOS' Vaisala FD12P laser to estimate visibility by analyzing the scatter of the laser beam. These data were also gathered at one-minute intervals over the course of the study. For truthing and tuning purposes, only hourly observations (as reported to the NWS) were used. Limiting the tuning data set made it practical for a manual observer to compare the ASOS reported visibility to the image being analyzed.

Figure 4: Example ASOS station configuration, the FD12P laser visibility sensor is shown at the top of the mounting pole.

The FD12P laser makes for an excellent automated method of generating standard meteorological visibility estimates. However, the FD12P laser produces an estimate of the visibility using a small spatial sample, as the distance from the laser transmitter to the receiver is only three meters. As mentioned above, the localized estimation of visibility from the ASOS may sometimes be different from the visibility as seen in a camera image but nevertheless be consistent with prevailing conditions. Therefore, when observations and model estimates were significantly different (e.g. ASOS reported ten miles and the model predicted one), the image was manually verified by an observer. In the cases where the ASOS was found to be in error a correction was made to the truth. Less than 3% of the more than 3000 ASOS observations were corrected, but those that were represented gross errors by the ASOS. There were two main reasons for corrections: (1) the delayed response of the ten-minute time-averaged ASOS report to rapidly improving or decaying visibility conditions and, (2) distant weather viewed by the camera which was not yet impacting the ASOS. However, the correction rate likely underestimates the number of corrections that might have been made because only gross differences in truth and model estimate were investigated. In addition, the maximum visibility range reported by the ASOS is ten miles (16.1 km). However, many camera images can detect features that are much further than the ASOS reported distance. This difference certainly had an impact on the algorithm's ability to estimate visibilities of distances over ten miles. However, a nominal clear day visibility of ten miles is more than sufficient for most meteorological applications.

5. MIT/LL CAMERA ALGORITHM

The MIT/LL visibility 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. Currently, the system ingests only grey-scale imagery (color images are converted to grey-scale on ingest). The composite is simply the average of all daylight clear-day imagery over a fixed time period (up to 30 days). As shown in the algorithm flow diagram (Figure 5), a camera image is captured and then an edge extraction technique is applied. Both the raw image and the edge detections are sent to the composite image generator where the images are checked to determine if enough edges are

visible for this to be a clear daytime image. If the time period is acceptable then the raw and edge images are added to their respective composites. The system is flexible in that it allows access to the overall composite of the raw or edge detection images as well as the same images broken down by the hour of the day.

The raw edge detection is also sent to the normalized edge extraction process where a comparison is done between the current image and the composite image. This process allows the system to focus only on the edges that might be expected on an average clear day. By averaging, transient edges in the image such as cars, people, etc are removed. Finally, a set of image and edge characteristics are calculated for the current image and, in turn, compared to the tuned visibility curves for each detector. The detector estimations are then averaged together to provide the estimated visibility for the current image.

5.1 Overall image analysis vs. "targets"

When studying how people perceive visibility, researchers will often utilize targets of fixed shapes and sizes. For example, Kwon (2004) set up a series of fixed sign targets at various distances. Visibility was then determined by asking users to select the furthest target that they could see clearly. This research was prototyped by having a State DOT deploy signage at fixed distances from camera locations. An image analysis algorithm was then employed to detect those specific targets within the image. However, there are two problems with deploying such a system on a wide scale. First, deploying visibility markers at every camera location would become expensive and time-consuming. And, second, making such an algorithm robust to be able to recover from sign blockage or damage (if a sign falls down or is removed) would be challenging. It may be better to utilize the entire image in the analysis.

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. Overall image analysis is based on using feature detectors that focus on characteristics of the full image along with edge features throughout the entire image. It is expected that by including the full range of imagery that the resultant algorithm will be more stable than a fixed target system. For these reasons, overall image analysis was utilized for the *Clarus* research effort.

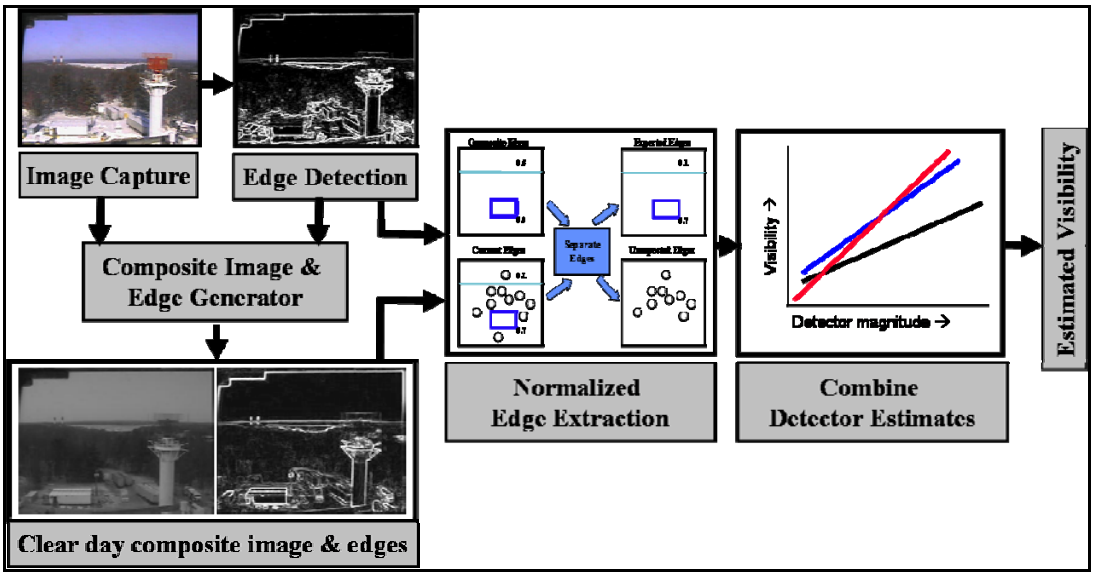

Figure 5: Algorithm flow diagram for the MIT/LL automated camera visibility algorithm.

5.2 Edge detections

Edges are the key to visibility estimation. A manual observer utilizes edges, or contrasts between one object and the background, to estimate which objects can be seen at various distances, as he looks out over the landscape. Similarly, the automated algorithm starts by surveying the scene with an edge detection algorithm. The method chosen for this analysis is a 'Sobel' algorithm (Parker, 1997), although other edge detection algorithms were also experimented with and appeared equivalent. The images in Figure 6 were taken overlooking a radar installation at HAFB. The top portion of Figure 6 depicts a clear day image and its associated edge detection image. Edges both near and far can be seen quite clearly in this high visibility case. The bottom portion of Figure 6 shows low visibility conditions with laser-measured visibilities of approximately 100 meters. Correspondingly, the furthest edges that can be seen are those related to the radar tower sitting just over 60 meters from the camera position.

5.3 Composite generation and usage

The composite generator is designed to create a running average of the most recent clear images and edges ingested into the system. The composite can be generated, stored and utilized in a variety of ways. Initially, the algorithm was designed to create the composite by averaging each edge detection image; however, it was discovered that the composite edge was cleaner and more precise if the raw image was used to create the composite. Each time the algorithm called for a composite edge comparison, the edges would be generated in real-time from the raw composite image. Initially, the reference composite was generated in real-time, but removing low visibility or bad images from the composite in real-time is challenging. The algorithm verifies that a significant number of edges are in the image (as a percentage of the expected edges), but if the camera were to move or was overwhelmed with transient edges (the former being a formidable challenge in the DOT images discussed later) then the composite would become corrupted with undesirable images and edges. In practice, the composite was generated offline with prior knowledge that the images selected were clear. In fact, it was found that composite image generation required relatively few images for it to become satisfactory for the system. An 8-hour day of clear images was comparable to a month's worth of daytime imagery when generating a composite.

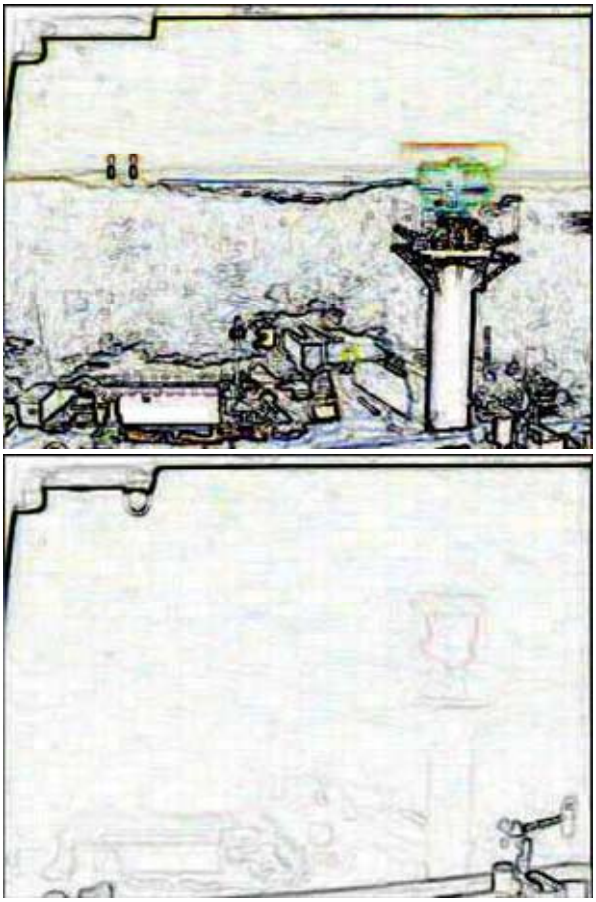

Figure 6: Edge extracted images for an example Hanscom AFB camera from a clear (>40km visibility) day on the top, and a low visibility (<100 meter) day on the bottom.

Further improvements were made by creating composites broken down by hour of the day instead of averaging over all daylight hours. Shading and light reflections can cause significant shifts in the edges found within an image and obviously changes in sun

angle impact both. The analysis presented in the results section reflects the implementation of a pre-calculated composite averaged over all daylight hours. Results improve a few percentage points when an hour-by-hour composite is utilized.

5.4 Normalized edge extraction

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 the expected edges is accomplished by maintaining a composite image as discussed above. 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 7 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.

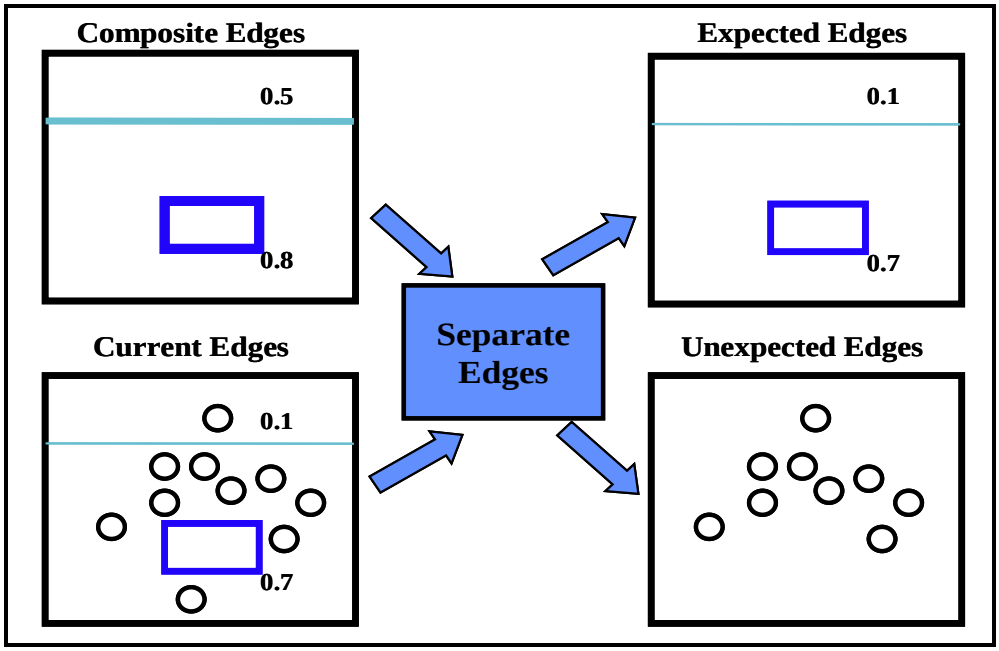

Figure 7: Illustration of edge analysis to separate strong, long-term "expected" edges from strong, but transient "unexpected" edges.

5.5 Calculating and correlating estimators

The algorithm uses fuzzy logic integration to determine a consensus estimate of the visibility in the image. Initially, a wide set of potential estimators was calculated, ranging from simple sums of the image pixel strengths to Fast-Fourier Transforms (FFT) of the normalized expected edges. Table 1 lists the 22 estimators examined along with a brief definition of the statistic they represent. The image magnitude is calculated by performing an FFT and then summing the absolute value of all 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.

Table 1. Estimator descriptions for visibility algorithm

Estimator name	Description
MeanImage	The mean grey-scale value of the entire image.
StdImage	The standard deviation of all grey-scale values in the image.
MagImage	The magnitude of all grey-scale values in the image.
MeanEdge	The mean edge strength value of all edges within the image.
StdEdge	The standard deviation of all edge strength values within the image.
MagEdge	The magnitude of all edge strength value of all edges within the image.
MaskedMeanEdge	The mean edge strength value for all edges within the current image that match expected edges that are in the composite edge image.
MaskedStdEdge	The standard deviation of all edge strength values for all edges within the current image that match expected edges that are in the composite edge image.
MaskedMagEdge	The magnitude of edge strength values for all edges within the current image that match expected edges that are in the composite edge image.
ExMeanEdge	The mean edge strength value of all unexpected edges within the current image.
ExStdEdge	The standard deviation of all unexpected edges within the current image.
ExMagEdge	The magnitude of all unexpected edges within the current image.
MissMeanEdge	The mean edge strength value for all expected edges in the composite image that are missing from the current image.
MissStdEdge	The standard deviation of edge strength value for all expected edges in the composite image that are missing from the current image.
MissMagEdge	The magnitude of the edge strength values for all expected edges in the composite image that are missing from the current image.
DiffMeanEdge	The mean of all edge strength differences between the composite and the current image for all composite edges.
DiffStdEdge	The standard deviation of all edge strength differences between the composite and the current image for all composite edges.
DiffMagEdge	The magnitude of all edge strength difference between the composite and the current image for all composite edges..
RatioMedianEdge	The median of all edge strength ratios (composite over current image) for all composite edges.
RatioMeanEdge	The mean of all edge strength ratios (composite over current image) for all composite edges.
RatioStdEdge	The standard deviation of all edge strength ratios (composite over current image) for all composite edges.
RatioMagEdge	The magnitude of all edge strength ratios (composite over current image) for all composite edges.

Each potential detector was calculated for each image and then the full set of calculated values was correlated against the hourly ASOS visibility measurement (corrected as noted above). Figures 8 and 9 show two examples of estimator correlations. Figure 8 shows the RatioMedianEdge estimator and Figure 9 shows the MissMeanEdge estimator. The scoring curves used to map the estimator value to the visibility are also shown as solid lines. There were a total of 17 estimators that were of sufficient quality to be incorporated into the visibility estimates for the roof-level camera.

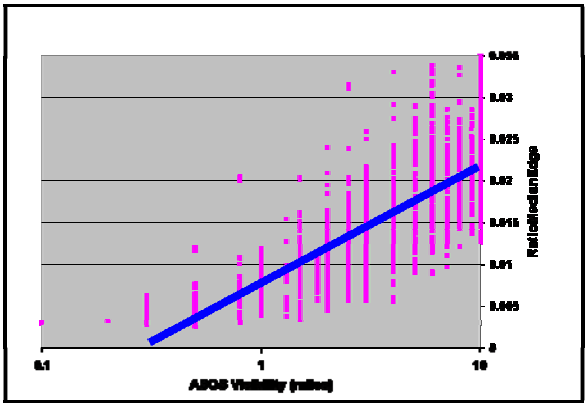

Figure 8: RatioMedianEdge vs. ASOS Visibility for roof-level camera estimator. The solid blue line represents the scoring function used to convert RatioMedianEdge values to visibility.

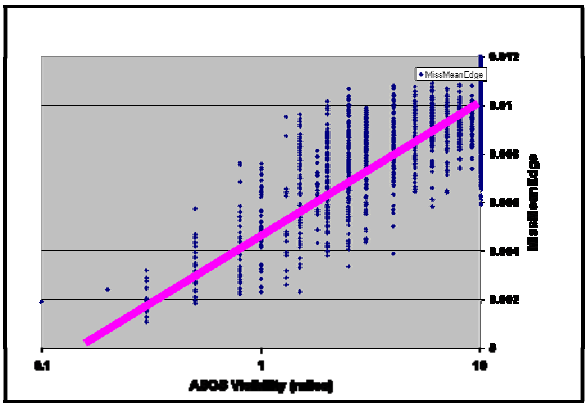

Figure 9: MissMeanEdge vs. ASOS Visibility for roof-level camera visibility estimator. The solid magenta line represents the scoring function used to convert MissMeanEdge values to visibility.

It would be ideal if the same set of estimators could be used to predict visibility from any camera or viewing angle. However, attempting to simply apply the same detectors as the roof-level camera to the road-level camera imagery resulted in extremely poor estimates of visibility. So, each camera site was allowed to have its own optimized set of detectors, weights and curves. Attempts to use an overall optimization for both cameras were also unsuccessful and resulted in significantly poorer individual results.

6. TEST SITE ALGORITHM RESULTS

Algorithm results for the roof-level camera are shown in Figure 10. The data set used was the same as that used to optimize the algorithm scoring functions. Data was scored for all camera images that were gathered within two minutes of the ASOS observation (therefore there may be multiple images for each ASOS observation). Overall, the algorithm performed extremely well on the 3114 camera images verified. The calculated Critical Success Index (CSI) over all ranges was 78% with a standard deviation between observed and predicted visibility of +/- 1.2 miles. But, images with very low visibilities had a significantly lower CSI of just 49% (primarily due to the high probability of false alarms when the algorithm estimated a <1 mile visibility) and a standard deviation of +/- 0.77 miles. Testing on an independent data set of roughly 1000 camera images for one month following the optimization of data yielded only a slightly lower overall CSI of 72%.

The road-level camera, however, had significantly poorer results on roughly the same data set (see Figure 11). The overall CSI value was just 57% with a standard deviation between observed and predicted visibilities of +/- 3.1 miles. In addition almost none of the ASOS visibilities below 1 mile were estimated correctly. The road-level camera was almost incapable of distinguishing between edges at different ranges as the low incidence angle of the

camera to the ground tended to blur all the edges together. This result strongly indicates that the height and angle of the camera are extremely important to algorithm performance.

Video Algorithm Visibility (miles)	ASOS Visibility (miles) <1	1-5	5-10	≥ 10
<1	110	80	0	0
1-5	29	495	66	45
5-10	4	105	119	193
≥ 10	2	55	113	1698

Figure 10: Visibility algorithm results for the MIT/LL roof mounted camera.

Video Algorithm Visibility (miles)	ASOS Visibility (miles) <1	1-5	5-10	≥ 10
<1	1	2	0	1
1-5	92	335	75	149
5-10	4	19	32	387
≥ 10	7	14	7	646

Figure 11: Visibility algorithm results for the MIT/LL road-level camera.

7. APPLICATION TO STATE DOT CAMERAS

Camera data were downloaded for several cameras from the Utah Commuterlink traveler information site (www.utahcommuterlink.com) starting on January 10th, 2006 and ending for this analysis on June 30th, 2006. Images were collected every 10 minutes, although the update rate of each camera varied from 10 to 20 minutes between cameras. Weather data, as with the MIT/LL cameras, were gathered from the standard hourly NWS ASOS station (the additional one-minute ASOS data were not available for this dataset). Two of the downloaded camera images are shown annotated with key features and distances in Figures 12 and 13.

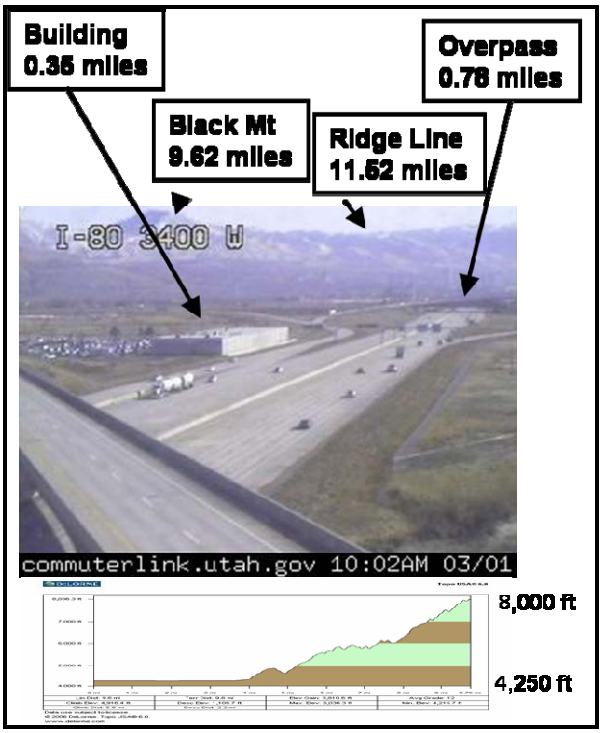

Figure 12: Salt Lake City, Utah camera number 48 located at I-80 and 3400W Street. Key edges within the image are identified; the image shown is the nominal operating image looking west.

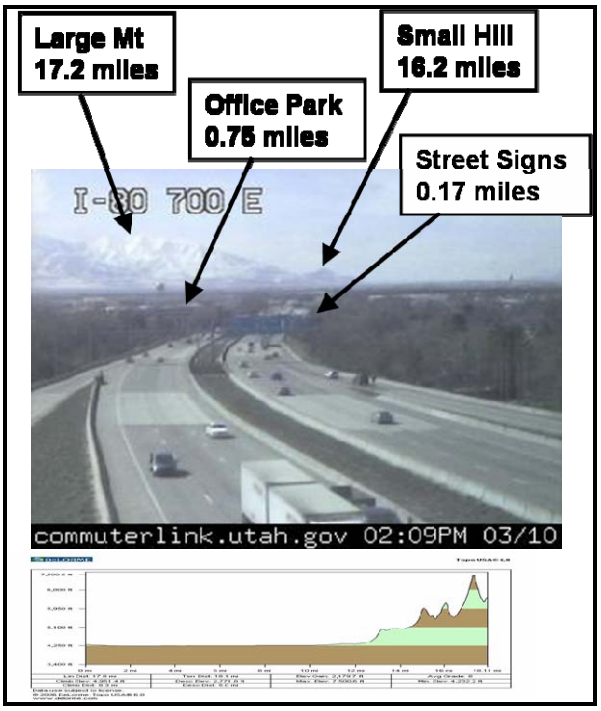

Figure 13: Salt Lake City, Utah camera number 56 located at I-80 and 700E Street. Key edges within the image are identified and annotated with the distance of each feature. The image shown is the nominal operating image looking west.

7.1 State DOT algorithm results

The extension of the algorithm methodology to additional existing DOT cameras in Utah was challenging. There were three main obstacles in analyzing the DOT cameras: (1) a lack of a homing setting, (2) random changes in camera views, sometimes for extended periods of time, and (3) the camera was incapable of reporting back its pan, tilt, and zoom settings for automated processing.

Although the camera images shown in the figures are the nominal views for each camera, there was no pre-set homing setting that brought the camera back to its base position and setting. So, if the camera was moved to view other areas and then returned to the nominal view the image would often be slightly shifted or (un)zoomed. This reset error often resulted in the composite imagery becoming blurry as the edges became misaligned. To alleviate this problem an image registration process was added to the algorithm. Each incoming edge image is correlated to the expected composite image edges, and shifted up to 15 pixels to compensate (this had the adverse affect of reducing the overall image size that could be processed). If the camera was shifted to an entirely different view, the algorithm would ignore it and wait for the next acceptable image.

Algorithm results for Camera 48 are shown in Figure 14. The data set used was the same as that used to optimize the algorithm scoring functions. Data were scored for all camera images that were gathered within ten minutes of the ASOS observation. Overall the algorithm performed with a calculated CSI over all ranges of 61% with a standard deviation between observed and predicted visibility of +/- 2.3 miles. Images with very low visibilities also had a CSI of 61% and a standard deviation of +/- 0.8 miles. Similar results can be seen for Camera 56 in Figure 15. The overall CSI score was 62% and a standard deviation of +/- 2.5 miles, although the CSI for less than 1 mile was 53% and a standard deviation of 1.1.

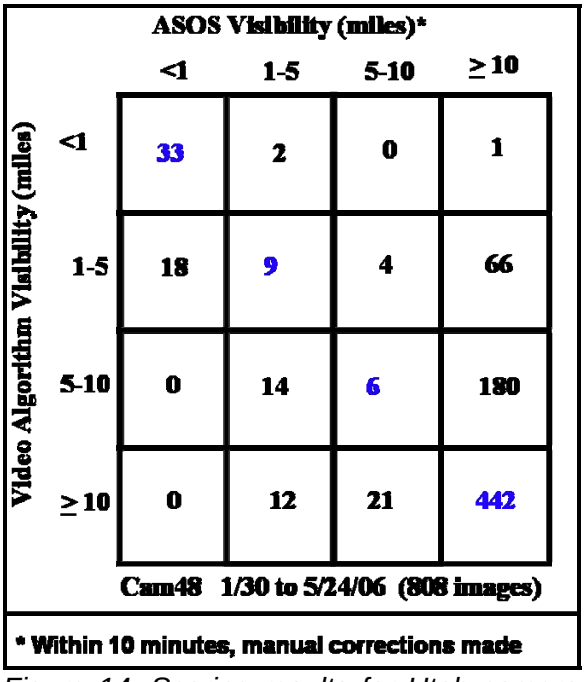

Figure 14: Scoring results for Utah camera #48.

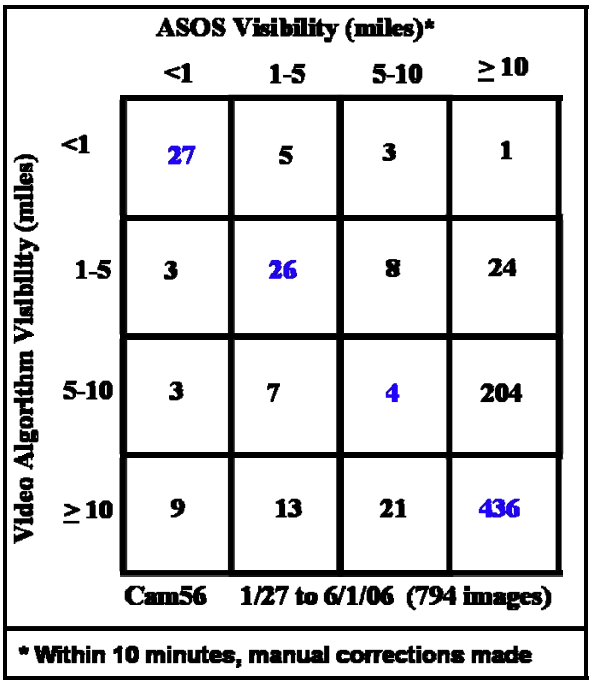

Figure 15: Scoring results for Utah camera #56.

8. CONCLUSIONS

Camera and video data availability is growing rapidly and automated tools are needed to extract derived information from images that are often the only source of valuable weather and surface condition data in the transportation corridor. An algorithm was presented that shows promise in estimating the critical variable of roadway visibility. Estimators were defined that generated visibility estimates that are optimal for providing qualitative estimates of visibility conditions (e.g. low, moderate, high) and moderately successful at predicting specific visibility values. Tuning of the system required only 4-6 weeks of variable visibilities. The technique of extracting expected and unexpected edges as a method of both finding the most important edges and controlling for contaminants on the lens performed quite well. In addition, there are several important findings from this work that should be noted for further camera research:

Camera siting:

Elevated cameras that have a diverse view of the surrounding landscape show the most promise for the edge-based visibility algorithm. The road-level camera at the test site performed poorly during low visibility events, however, Utah camera 48 was a low camera yet it performed equally well on low and high visibility events. **Therefore, the sensitivity of the algorithm to camera siting appears to be more a function of the range of edges that are visible in the image rather than the relative height of the camera itself.**

Camera meta-data:

Standard State DOT camera imagery offers many challenges from alternating image scenes to misaligned edges due to repositioning errors. Techniques were successfully applied to control for slight misalignments. Alternating scenes, when a camera switches between preset views could be handled by allowing the system to maintain multiple expected composites (and some experimentation was successfully performed as a proof-of-concept). However, manual shifting of images to capture accidents or monitor road work was extremely difficult to differentiate from poor visibility images that should be processed. As part of the *Clarus* system design, metadata for camera imagery is being proposed that will capture attributes such as pan, tilt, and zoom factors. **State DOTs should strongly consider capturing camera orientation data to make all forms of automated processing simpler and more exact.**

How good is good enough?:

A number of states operate fog or low visibility warning systems that can be used as a measure of desirable operating ranges and accuracies (FHWA, 2006). The average system begins to warn motorists that low visibility conditions are present when visibilities are between 250 to 400 meters, speed reductions are recommended between 60 and 300 meters, and roads are often closed when visibilities drop below 60-

90 meters. Formal requirements for detection/forecast accuracy of visibility conditions for surface transportation operations could not be found. In general, users express a need for a better than 90% probability of detection with less than 10% probability of false alarm (corresponding to a CSI of 82%).

The overall CSI value for the four cameras ranged from 57-78%. CSI scores for low visibility events (<1600 meters) were significantly poorer, ranging from near zero to 61%. Scores for visibilities in the ranges applicable to the fog warning systems cited above (<400 meters) were not possible because of the infrequency of the events within the test data set. **When examining the larger dataset, the algorithm performed close to user expectations and with further refinement could meet performance criteria. However, the lack of events with these extremely low visibilities was a shortcoming of the analysis effort. Future research should capitalize on fog warning data systems as a source of very low visibility events.**

Towards a portable visibility algorithm:

The algorithm parameters were able to be tuned on multiple cameras using only a small subset of verification data. Ideally one would prefer a single set of estimators that allowed users to estimate visibility from any camera given limited meta-data about the environment. However, generically applying the original test suite algorithm by adding camera characteristic variables was only marginally successful. **The existing algorithm could be set up in a dense camera environment (such as Salt Lake City, UT) allowing many cameras to be tuned simultaneously. The wide variety of camera views sensing in a similar visibility environment would provide a rich dataset to optimize the generalization of this algorithm.**

9. FUTURE CAMERA RESEARCH

The field of mining digital imagery for weather and road condition data is in its infancy. By leveraging the growing industry of security surveillance for homeland defense and traffic monitoring operations, new algorithms to detect and/or measure precipitation type/intensity, wet roads, icy/snow covered roads and bridges, and refined algorithms for visibility are on the horizon.

But these algorithms will only be marginally useful if they stand alone as simple point sources of data. These data must be integrated into an overall forecasting system such as the one visualized under *Clarus*. The power of integration allows the marginal estimate of visibility from a new video sensor to become the final piece that confirms hazardous conditions in a lower valley, hail on a roadway, or other weather/road hazards. In addition, connecting the camera sensor into the larger network allows there to be feedback to the camera itself. High wind or hail warnings could trigger the camera to transition to a protective mode. A system like MDSS might activate a camera to pan and zoom to a potentially icy intersection to confirm treatment recommendations.

Finally, automated algorithms for weather, traffic, emergency management and security could be integrated themselves. Integration at this level would provide priority ratings of various images based on content and allow users to choose only those images, or roadways, or vehicles that are of interest to their specific application. For example, Traffic Management Centers could benefit from this priority rating by automatically selecting images where traffic is stopped or an accident is detected.

10. REFERENCES

AMS, 2005: Glossary of Meteorology, American Meteorological Society, Boston, MA. 2005.

Clark, D.A., 2000: An integrated weather testbed for support of theater operations in complex terrain, Battlespace Atmospheric and Cloud Impacts on Military Operations

Conference, Fort Collins, CO, 25-27 April 2000.

FHWA, 2006: Road Weather Management Program website: Best Practices for low visibility. 2006. http://ops.fhwa.dot.gov/weather/weather_events/low_visibility.htm .

Hagiwara, 2002: Assessment of visibility on roads under snow conditions using digital images, 11th International Road Weather Conference. Sapporo, Japan. January 26, 2002.

Hagiwara, 2006: A Simple Approach for Using CCTV Road Images to Provide Poor-Visibility Information, Transportation Research Board, Washington, DC January 2006.

Hallowell, R.G., et.al, 2005: Automated Extraction of Weather Variables from Camera Imagery, Mid-Continent Transportation Research Symposium, Ames, IA, August 2005.

Hautiere, 2006: "Automatic fog detection and estimation of visibility distance through use of an onboard camera", Machine Vision and Applications, v 17, n 1, April, 2006, p 8-20

Johnson, Mark, 2006: Video, sound advances aimed at war on terror, Associated Press State and Local News Wire, August 7, 2006, Niskayuna, NY

Kwon, T.M., 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)

Parker, J.R., 1997: Algorithms for Image Processing and Computer Vision, New York, John Wiley & Sons, Inc., 1997, pp. 23-29

Pisano, P.A., 2004: The winter maintenance decision support system (MDSS): Demonstration results and future plans, American Meteorological Society, 20th International Conference on Interactive Information and Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology, Paper 18.4, January 2004.

Pisano, P.A., 2006a: "Concept of Operations for Clarus – The Nationwide Surface Transportation Weather Observing and Forecasting System, Transportation Research Board, Paper 06-1193, January, 2006

Pisano, P.A., 2006b: FHWA's Clarus Initiative: Concept of Operations and Associated Research, 22nd International Conference on Interactive Information Processing Systems for Meteorology, Oceanography, and Hydrology, American Meteorological Society, Atlanta, GA, January 2006.

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. June, 2002.

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

Yamada, M., 2001: Discrimination of the Road Condition Toward Understanding of Vehicle Driving Environments, IEEE Transactions on Intelligent Transportation Systems, Vol. 2., No. 1, March 2001.