VISIBILITY ESTIMATION FROM A MOVING VEHICLE USING THE RALPH VISION SYSTEM

Dean Pomerleau

Robotics Institute, Carnegie Mellon University

5000 Forbes Avenue, Pittsburgh, PA 15213 USA

Keywords: Visibility Estimation, Lane Tracking, Vision System, CCD Camera

ABSTRACT

Reduced visibility is a common casual factor in many traffic accidents. This paper describes a forward looking vision system which simultaneously track the lane and estimate visibility. The system estimates visibility by measuring the attenuation of contrast between consistent road features at various distances ahead of the vehicle. Results of experiments on simulated images, as well as live vehicle tests are presented.

1. Introduction

Reduced visibility caused by fog, rain, snow, darkness and glare is a frequent contributing factor to traffic accidents [Najm et al., 1995]. In fact, some of the most serious of all highway incidents, sometimes involving dozens or even hundreds of vehicles, occur when reduced visibility conditions result in a chain reaction of crashes. Paradoxically, some advanced technology, like Adaptive Cruise Control (ACC) systems have the potential to decrease, rather than increase safety in these situation by encouraging drivers to travel at a speed and headway distance that may not be safe for the ambient environmental conditions. This paper describes the first step in the solution to this problem, a system that can estimate the ambient visibility from a moving vehicle.

There are several technologies typically employed to estimate visibility, including transmissometers, which measure the transmittance of the atmosphere over a baseline distance, and nephelometers which measure the scattering coefficient of an air sample caused by suspended particles [Federal Meteorological Handbook, 1996]. Unfortunately, these systems suffer from several drawback for automotive applications. Transmissometers require a transmitter and a receiver a substantial distance (typically hundreds of meters) apart, which is very difficult to implement on a moving vehicle. Stationary transmissometers located near stretches of roadway commonly plagued with poor visibility can be effective for a local area, but may miss nearby reduce visibility conditions because of the very localized nature of some reduced visibility phenomena.

Nephelometers can be mobile, since they use a collocated transmitter and received to measure the backscatter of light off particles in the air. However they are prone to miss many of the important phenomena effecting how far a driver can truly see. These phenomena include:

- Opacity of the atmosphere due to particulates
- Ambient lighting conditions - sun, moon, overhead lights, direction of lighting
- Headlights from the driver's own vehicle and other vehicles
- Windshield transmissive properties due to dirt, water, snow or ice buildup.

The only way to automatically estimate the cumulative influence of these factors on the driver's ability to see potential obstacles ahead is to employ a sensing system which reasonably match the driver's perceptual characteristics. The system described in this paper accomplishes this match by using a CCD video camera pointing out the windshield of the vehicle, and processing the same features as the human driver to estimate visibility.

0-7803-4269-0/97/$10.00 © 1998 IEEE

2. Approach

Manual visibility estimates are typically made by attempting to detect high contrast targets at various known distances. The farthest distance at which a target can be reliably detected is considered the visibility distance. Ideally, an automated visibility estimation system should work the same way. Unfortunately, it is very difficult to consistently find high contrast targets at various known ranges from a moving vehicle. Even the features that are supposed to be consistent on a roadway, the lane markings, vary greatly in their appearance, and are in fact frequently missing or obscured. The RALPH (Rapidly Adapting Lateral Position Handler) system [Pomerleau et al., 1996] overcomes this difficulty when detecting the position and curvature of the road ahead in camera images by utilizing whatever features are visible on the roadway, including lane markings, road/shoulder boundaries, tracks left by other vehicles, and even subtle pavement discolorations like the oil stripe down the lane center when necessary.

The visibility estimation system described in this paper exploits RALPH's ability to find and track arbitrary road features. In short, the system estimates visibility by measuring the attenuation of contrast between consistent road features at various distances ahead of the vehicle.

2.1. Road Feature Detection

To measure contrast between consistent road features, first these features must be detected in images of the road ahead. The algorithm the RALPH system uses to find road features is based on the observation that when viewed from above, a road resembles a ribbon of parallel bands formed by lane markings and other road features. To exploit this characteristics, RALPH first extracts from the image a trapezoidal region of the road ahead (See Figure 1). RALPH automatically varies the position of this trapezoid based on the vehicle's velocity and the current visibility, but under good conditions the top of the trapezoid is typically viewing the road between 50m and 120m ahead of the vehicle. RALPH resamples the image from this trapezoid. The horizontal extend of the trapezoid is set so that its width on the ground plane is identical at each row of the image. The horizontal distance that each row of the trapezoid encompasses is approximately 7.0 meters, about twice the width of

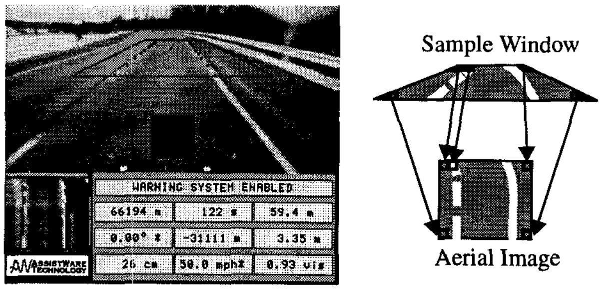

Figure 1: Forward looking image (left), and RALPH's sampling strategy (right).

a typical lane. This trapezoid is selectively sampled according to the strategy depicted in the schematic on the right of Figure 1 to create an aerial view of the road ahead. This sampling process results a low resolution (35x50 pixel) image in which important features such as lane markings, now appear parallel in the low resolution image (see schematic aerial view in the lower right of Figure 1, and the actual aerial view show in the lower left of Figure 1). Note that this image resampling is a simple geometric transformation (based on the assumption that the road is locally planar), and requires no explicit feature detection.

RALPH then uses this aerial image to locate the road ahead. To accomplish this, RALPH uses a one-dimensional representation of the road, created by taking a cross section of the aerial image perpendicular to the road, called the road template. The aerial image for the road in Figure 1 and road template created from a cross sections at the bottom of the image, are shown in Figure 2

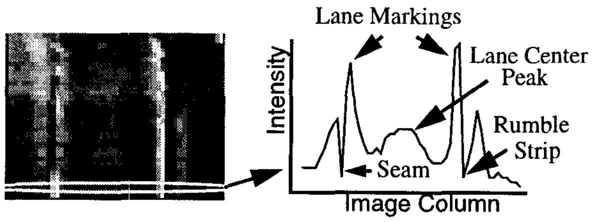

Figure 2: An aerial road image (left) and cross sections taken from the bottom of the image (right).

There are several things to note about the template cross section. First, the lane markings show up quite distinctly as the two highest peaks. Also apparent in the cross section are two sharp dips just outside the lane markings, caused by a black filled seam in the pavement on the left side of the lane, and the dark banding of a rumble strip on the right side. Finally, down the center of the lane the pavement is slightly lighter in intensity than the more heavily worn pavement closer to the lane boundaries, causing a wide shallow peak in the center of the cross section.

RALPH exploits all of these features to find the road ahead by using the entire one-dimensional cross section as a template. For each row of the aerial image, RALPH shifts the template left or right until it best matches the particular row's cross section. The amount of shift required to match a particular row is proportional to the lateral displacement of the lane center at that row of the image. For more details on the algorithm RALPH employs to generate and maintain the template, and how RALPH finds the position and curvature of the road ahead using the template, see [Pomerleau et al., 1996].

2.2. Visibility Estimation

In order to estimate visibility, the system uses the shifted road cross sections generated during the road detection process. Two such cross sections, one from the top of the aerial image, and one from the bottom, are shown in Figure 3. Notice that at

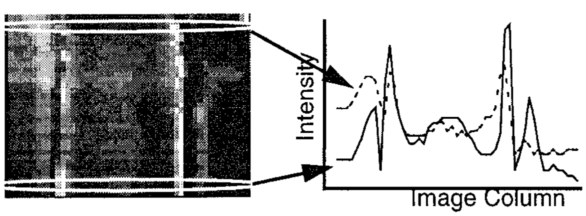

Figure 3: Road cross sections from top (dashed curve) and bottom (solid curve) of the aerial image

the top of the image, relative far ahead of the vehicle, the peaks in the cross section are not quite as high, and the dips are not quite as low as the at the bottom of the image, close ahead of the vehicle. Qualitatively, it is this attenuation of contrast between features with increasing distance from the vehicle that the visibility estimation algorithm (described below) is measuring.

To quantify the feature attenuation, the system estimates for several rows at the top and bottom of the image, the median intensity around the lane center, as well as the maximum deviation from this median intensity within the row. The system averages the maximum intensity deviation for the rows at the top, and the rows at the bottom of the image, to overcome the effects of intermittent dashed lane boundaries and other image artifacts. The difference between the average maximum intensity deviation at the bottom and the top of the aerial image is the system's estimate of contrast attenuation.

In order to estimate visibility, it is not enough to simply measure contrast attenuation, since visibility should be a function of distance. Therefore, the contrast attenuation as measured above is scaled based on the distance between the top and bottom of the RALPH's view trapezoid (which can vary as mentioned previously). The resulting value is a measure of contrast attenuation per meter.

The final step in estimating visibility is normalization. Even under clear conditions like that shown in Figure 1, the contrast in the aerial image is significantly attenuated, even over the relative short distance between the bottom and the top of the image (see Figure 3). This is cause primarily by imaging artifact relating to the pixel spacing on the CCD array, and the camera's limited depth of field. Together these artifacts result in a blurring towards the top of the aerial image under all conditions. To eliminate the effect of this blurring on the visibility estimate, the contrast attenuation per meter value is normalized, so that the rate of attenuation on a bright clear day is equivalent to a visibility of 1.0, and visibility under degraded conditions are expressed relative to this baseline.

Figure 4 depicts an example of a reduced visibility condition, night driving. In this situation, the driver's visual range is reduced due to the limited range of the vehicle's headlights. This can be seen in the reduced contrast towards the top of the view trapezoid. Cross sections from the top and bottom of the aerial image for this night image are shown on the right of Figure 4. Note how the absolute intensity of the cross section, as well as the maximum contrast in the cross section, are greatly reduced towards the top of the image when compared with results from the daytime scene shown in Figure 3. As a result of the greater feature attenuation, the visibility for this situation, as computed with the algorithm described above, has dropped to33% of the clear daytime visibility (reported as "0.33 vis" in the lower right corner of Figure 4).

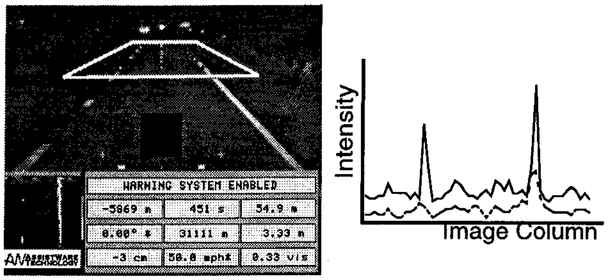

Figure 4: Night scene (left) with cross sections (right) from top (dashed) and bottom (solid) of the aerial image

3. Results

Two sets of experiments were conducted to test the visibility estimation algorithm's performance under a wide range of conditions. The first set of tests involved running the algorithm on a sequence of real road images in which various levels of simulated fog had been introduced through image manipulation. The second set of experiments involved live on-road tests of the visibility estimation algorithm.

3.1. Simulated Fog Experiments

As part of a project to test lane tracking systems under reduced visibility conditions [Pomerleau et al., 1995], Battelle Memorial Institute previously generated a set of images depicting various levels of fog from an image sequence collected on Carnegie Mellon's test vehicle, using Battelle's Electo-Optical Visualization and Simulation Tool (EOVAST) software. Given an original image, and accompanying estimates of camera characteristics, scene geometry and lighting conditions, the EOVAST software generates degraded versions of the same image as they would appear under user specified adverse weather conditions. The EOVAST software was originally developed for military targeting applications, and has been extensively validated for accuracy. For more details on EOVAST, and the results of the lane tracking tests under reduced visibility conditions see [Pomerleau et al., 1995].

In total, EOVAST was used to generate 120 reduced visibility images from 30 original images. These images depicted an interstate highway under foggy conditions with 700, 400, 300 and 100 meter visibility. A single one of the 30 original image, along with the same image in each of the four reduced visibility conditions is shown in Figure 5. These 150 images (30 original + 120 fog) were used to test the visibility estimation algorithm. Figure 6 shows the mean and standard deviation of the algorithm's visibility estimates for each of the five visibility conditions.

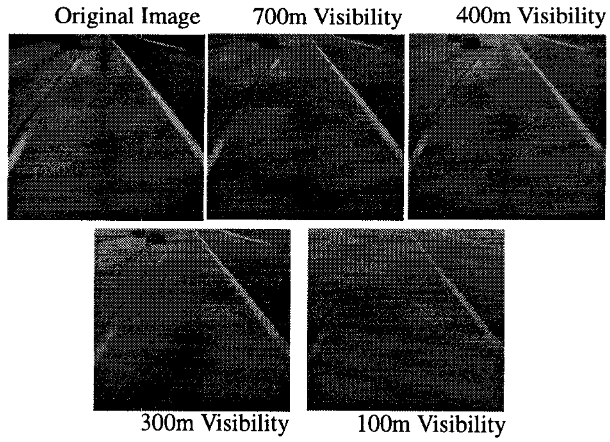

Figure 5: Original Image and four versions of the same image with simulated fog.

The first important characteristic of Figure 6 to notice is the substantial reduction in the algorithm's estimated visibility as the degree of fog increases (and hence the simulated visibility decreases). The second important attribute of Figure 6 is the large standard deviation in the algorithm's visibility estimates at each fog level (shown as the large spread in the error bars). Automatic visibility estimation with the algorithm reported here is a statistical process, since local variations in the underlying image features used to compute visibility can mask the contrast attenuation caused by ambient environmental factors. Therefore a relatively large number of images (more than 30) is required to determine visibility with certainty.

4. On-road Experiments

To overcome the problem of limited image data, and to test the algorithm under realistic conditions, a set of in-vehicle experiments were conducted using Carnegie Mellon's Navlab 8 test vehicle. Navlab 8 is an Oldsmobile Silhouette minivan equipped with a black and white video camera mounted behind the rear view mirror pointed through the windshield, and a Pentium-100 processor executing both the RALPH lane tracking algo-

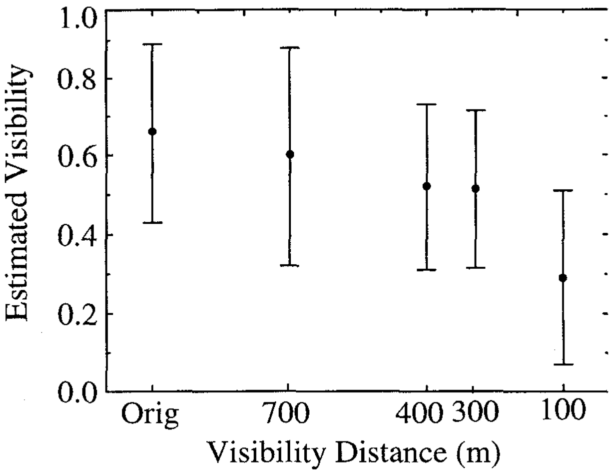

Figure 6: Mean and standard deviation of visibility estimates for the original image set, and the four reduced visibility conditions.

rithm and visibility estimation algorithm in real-time (15 frames per second).

Data on the visibility estimation algorithm's performance was collected on a 15 mile stretch of interstate highway, which offers several pavement types (concrete and asphalt) as well as a variety of lane delineating techniques, including solid and dashed white lane markings, yellow lane markings, retroreflectors, and roadside rumble strips. Data was collected on this stretch of roadway under six different conditions (See Figure 7 for example images from each condition):

- Daytime in good weather in the right lane
- Daytime good weather in the left lane
- Daytime in rainy weather
- Early morning with glare from the rising sun
- Nighttime with overhead lighting
- Nighttime without overhead lighting

The morning glare and the nighttime with overhead lighting conditions occurred on only limited stretches of the 15 mile test road. Therefore the results reported below for these two conditions were compiled over only two and three miles of testing, respectively.

Figure 8 shows the results of the experiments on the six conditions, in decreasing order of estimated visibility. First note the visibility estimates in the left and right lanes in good daytime conditions were nearly identical to each other, and were far above the estimates for the other conditions. The next best visibility was reported for the nighttime with overhead lights condition. As can be seen

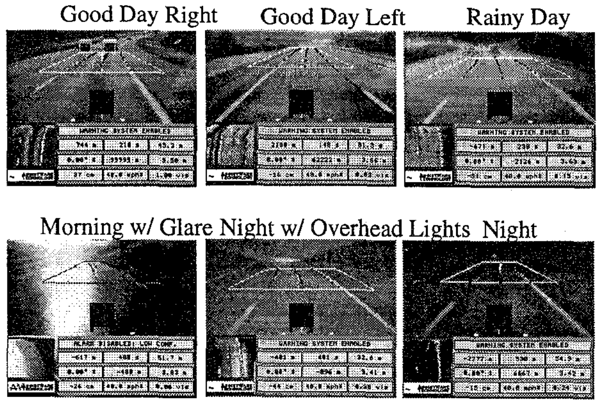

Figure 7: RALPH tracking the lane under various conditions, and estimating visibility.

from Figure 7, the overhead lights increase the range at which the road features are discernible, resulting in a corresponding increase in estimated visibility.

The nighttime condition with only headlight illumination was the situation the algorithm estimated to have the next best visibility, equivalent to approximately 30% of the good daytime visibility. Daytime rain, with significant water buildup on the windshield and substantial suspended spray in the air was determined by the algorithm to be the next to worse visibility condition tested. As Figure 7 shows, it is quite a bit more difficult to detect the road features, as well as other vehicles in this situation. However the lowest estimated visibility of the six tested was in the early morning glare condition. As is apparent in Figure 7, specular reflections off the pavement obscured the road features, and the very high ambient brightness saturated the camera, making it extremely difficult to detect the road (or anything else) anywhere except directly in front of the vehicle.

5. Discussion

The visibility estimation algorithm presented in this paper appears to perform well under a wide variety of conditions. The rank ordering of six conditions tested corresponds reasonably well to ones intuitive notion of how difficult it is to see in these situations. Note that traditional instruments for estimating visibility, which only detect suspended particles in the atmosphere, would have report less than unlimited visibility in only one of the six conditions tested, daytime rain.

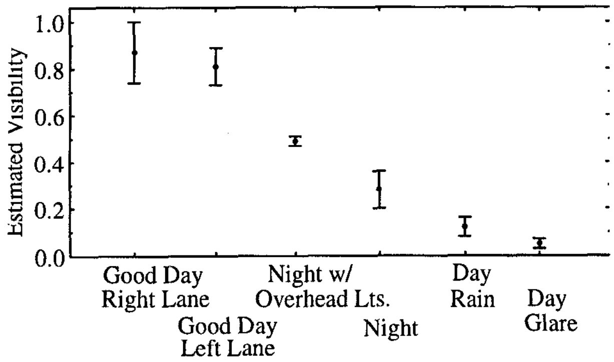

Figure 8: Mean and standard deviation of visibility estimates for the original image set, and the four reduced visibility conditions.

Interestingly, it is the very property for which vision systems are often criticized, their reduced effectiveness in adverse environmental conditions, which gives the algorithm its power. This is because the conditions in which the vision system has trouble seeing features are the same ones in which people have difficulty seeing.

One potential drawback of the visibility estimate technique presented is that it provides only a relative visibility measure, and not an absolute estimate of how far ahead road features or obstacles can be detected. However for a reduced visibility warning system, or a system to adjust the set speed and following distance of an adaptive cruise control, a consistent relative visibility measure may be sufficient. If an absolute measure of detection distance is required, it should be possible to calibrate the relative visibility estimates provided by the algorithm, although this hypothesis remains to be tested.

Live vehicle tests in fog still need to be conducted (fog is rare in Pennsylvania, particularly during the winter when these experiments were done). However, the results from the simulated fog experiments, and the live daytime tests in rainy conditions suggest that the algorithm should perform well, and report significantly reduced visibility under foggy conditions.

While all the work reported here has been done with a standard black and white CCD camara, we are investigating the potential for usingalternative sensors for improved performance. For example, a high-dynamic range camera would respond more like the human eye in extreme lighting conditions, and could therefore provide better visibility estimates.

Another possibility would be to combine this visibility estimationtechnique with a multispectral imaging device. By testing thevisibility at different wavelengths, it may be possible to select thebest wavelength(s) for operation under the current conditions.

6. Acknowledgments

The author gratefully acknowledges Larry Lazofson, Ed Kopala, and Howard Choe from Battelle for their hard work in generating the simulated fog images used during testing. This work was supported in part by the National Highway Traffic Safety Administration (NHTSA) under contract DTNH22-93-C-07023, by the Office of Naval Research (ONR) under contract N00014-95-1-0591, and by the Federal Highway Administration (FHWA) under cooperative agreement DTFH61-94-X-00001 as part of the National Automated Highway System Consortium. The content of this paper does not necessarily reflect the position or the policy of the US Government, and no official endorsement should be inferred.

References

[Najm et al., 1995] Najm, W., Mironer, M. and Fraser, L. (1995) "Analysis of Target Crashes and ITS Countermeasure Actions". *Proc. of 1995 ITS America Annual Meeting,* pp. 931-940.

[Pomerleau et al., 1995] Pomerleau, D., Kumar, B., Everson, J., Kopala, E., and Lazofson, L. (1995) "Run-Off-Road Collision Avoidance Using IVHS Countermeasures: Task 3 Report - Volume 1" NHTSA Contract DTNH22-93-C-07023.

[Pomerleau et al., 1996] Pomerleau, D. and Jochem, T. (1996) Rapidly Adapting Machine Vision for Automated Vehicle Steering. *IEEE Expert, Vol. 11, No. 2.* pp. 19-27.

[Federal Meteorological Handbook, 1996] "*Surface Weather Observations and Reports*" (1996) Federal Meteorological Handbook, 5th Edition, National Weather Service Publication FMH-1.