

An Investigation into Shadow Removal from Traffic Images

November 15, 2006

Total Length: 5,121 words + 8*250 = 7,121 words

Ryan P. Avery, PhD Candidate (Corresponding Author)
Research Assistant
Box 352700
Department of Civil and Environmental Engineering
University of Washington
Seattle, WA 98195-2700
Tel: (206) 543-7827
Fax: (206) 543-1543
Email: rpavery@u.washington.edu

Guohui Zhang
Research Assistant
Box 352700
Department of Civil and Environmental Engineering
University of Washington
Seattle, WA 98195-2700
Tel: (206) 543-7827
E-mail: zhanggh@u.washington.edu

Yinhai Wang, PhD
Assistant Professor
Box 352700
Department of Civil and Environmental Engineering
University of Washington
Seattle, WA 98195-2700
Tel: (206) 616-2696
Fax: (206) 543-5965
Email: yinhai@u.washington.edu

Nancy L. Nihan, Ph.D.
Professor and Director
Transportation Northwest (TransNow)
Department of Civil & Environmental Engineering
University of Washington
Seattle, WA 98195-2700
Tel: (206) 543-9639
Fax: (206) 543-5965
Email: nihan@u.washington.edu

ABSTRACT

Traffic surveillance cameras are becoming a viable replacement for inductive loop detectors. Their effectiveness, however, depends on video image processing algorithms that are capable of alleviating common problems such as shadows, vehicle occlusion, reflection, and camera shaking. Shadows have proven to be a large source of error in the detection and classification of vehicles. This study proposes three algorithms in increasing complexity to address the shadow problem. The algorithms each address the need to remove cast shadows from vehicles while preserving self shadows, or those areas of a vehicle that are hidden from illumination. They are also geared toward real-time analysis, which requires that they can be implemented efficiently and cannot have complex training or learning requirements. The dual-pass Otsu method of shadow removal was simplest in application but also performed the poorest. The proposed region growing technique, while showing considerable promise, failed when the pixel intensity varied widely in the shadow region. The last technique employed edge imaging to recognize shadows as areas with few edges or with edges substantially similar to the background. This method clearly out-performed the other methods and was subsequently proven in a separate paper describing a prototype vehicle detection and classification system.

Keywords: shadow removal, video image processing, region growing, traffic surveillance cameras

INTRODUCTION

Traffic surveillance video cameras have long been employed by traffic agencies to help monitor urban freeway systems. However, many agencies are beginning to recognize the additional value these cameras can provide when employed for data collection. Through video image processing, surveillance cameras can be used to collect more data than those currently provided by inductive loop detectors (1). Cameras can also be installed at a lower cost - the Washington State Department of Transportation (WSDOT) estimates the cost of loop installation to be between \$3,250 and \$5,750, including indirect costs associated with lane closure (2). Since one camera can view many lanes at once and may not even require road closure for placement, surveillance cameras are often a much cheaper alternative. Unfortunately, there are many pitfalls in video image processing – shadows, vehicle occlusion, inclement weather, camera shaking, and reflection are some of the top issues. Shadows are perhaps the most common and one of the more difficult problems.

From a camera's point of view, shadows have many of the same characteristics as vehicles – they move in similar patterns and directions, and they are considerably different from the background. Thus, they will be detected as a part of the vehicle whether a background or motion-based detection method is employed. Non-removal of these shadows can result in over-counts in adjacent lanes or under-counts by 'connecting' two vehicles in the camera field of view. Shadows can lead to bias in vehicle parameter estimates and misclassification of vehicle types. Thus, research on shadow recognition and removal is clearly necessary for accurate vehicle detection and classification.

Although shadows are easy to discern with the human eye, considerable effort has been made to help computers to do the same. An overview of the state-of-the-art in shadow removal for traffic applications follows, after which several methods of shadow removal are proposed and evaluated. Finally, this paper offers conclusions and perspectives on future work.

PREVIOUS WORK

The majority of research into removal of shadows from traffic images has been performed in the fields of computer science and electrical engineering. One of the earliest investigations in shadow removal was performed by Scanlan et al. (3). They split an image into square blocks and produced an image based on the mean intensity of each block. The median intensity of the mean values was then used as a basis for scaling all blocks below the median to the median value. The authors noted that this method is appropriate only for images where the objects of interest occupy the higher end of the intensity range. Thus, the method would not be suited for situations where the objects of interest may occupy the lower end of the intensity range (4). This method also introduces some loss of contrast and tends to cause 'blocking' (5).

Gamba et al. (6) built a shadow model based upon images from a monocular color image sequence. Noting that shadows in a scene that interact with still portions of the scene are more similar to each other than to the target objects of interest, the hue, luminosity, and saturation values were used to construct a reference image for the shadow model. The shadows present in the reference image were used as a model for moving cast shadows. However, since the reference image may not always contain enough still shadows to provide an accurate model, they also constructed a strip bitmap model to improve the shadow model. In this strip bitmap model, the image was split into a number of horizontal strips to analyze separately since they considered that luminosity values change with respect to distance from the camera (distant shadows appear lighter than closer shadows). Although the number of misclassified pixels was low, the algorithm was only tested on one scene at a supermarket parking lot. Furthermore, it has been noted that there is an implicit assumption that shadows are cast on the same kind of surface, which may not hold for a variety of outdoor scenes (4).

Gu et al. (7) implemented a biological approach to shadow removal. Noting synchronous pulse bursts in the visual cortex of cats, they implemented a Pulse Coupled Neural Network (PCNN) to simulate this effect for the removal of shadows based on optimization of the linking strength. The results indicate that shadows are satisfactorily removed for images that do not contain high degrees of noise.

Hsieh et al. (8) performed shadow removal to improve the accuracy of a person-tracking system. Their shadow removal method was based on the assumption that shadows have less variation in chromaticity and luminance than the target objects of interest. The tracked area was then decomposed by a wavelet transform and projected onto low and high frequency components to identify areas of low frequency that were considered to be shadow. The algorithm was able to perform satisfactorily even in situations where the tracked people wore colors close to the background.

Recognizing that many shadow removal algorithms produce distorted and noisy results that misrepresent the shape of the original object, Xu et al. (9) set out to fix these distortion errors. They presented a shadow removal algorithm based upon inspection of color and texture. The unique part of their work was the introduction of morphological operations upon the blobs remaining after shadow removal to reconstruct the shadow-removed object shape based on the shape of the object before shadow removal. The algorithm performs well except in cases of very large cast shadows. Correcting the brightness threshold used in the paper to account for larger shadows would improve the results but also introduce false positive shadow pixels.

Fung et al. (4) proposed a statistical shadow removal algorithm based upon construction of a probability map called the Shadow Confidence Score (SCS). The score was based on investigation of the luminance, chromaticity, and gradient density. The cast shadow was then determined to be those regions with high SCS values that were outside of the convex hull of the vehicle edge. The algorithm was tested on a variety of vehicle types and colors in different lighting conditions and viewing angles; the algorithm achieved an error rate of 14%, with motorcycles and vehicles with color similar to the background causing the highest rate of error. In the case of smaller vehicles, the error can largely be attributed to the use of a convex hull to represent the object, since smaller vehicles and motorcycles have outlines that are not very well preserved by a convex hull.

Noting that the performance of traditional Bayesian Networks deteriorates with highly varying input data, Lo et al. (10) developed an adaptive Bayesian Network to avoid the problems of their static counterparts. This was accomplished via the development of an efficient means to capture the variation in subsequent input images. This information was then utilized to adjust the network parameters. The performance was evaluated against a static Bayesian Network, and it was demonstrated that the adaptive network performed better.

Wang et al. (5) proposed a three step process to remove shadows from a foreground object obtained after subtraction of an image from a background image. The first step is illumination assessment, in which the foreground region is analyzed to determine if it contains any shadow based on pixel intensity and energy. If a shadow is suspected to exist based on aggregate statistics of bright and dark pixels, the shadow detection step is performed. The direction of illumination is found via the Otsu method (11) over the boundary pixels. Points near the boundary in the direction of illumination are sampled to derive shadow attributes. Object areas are recognized by subtracting the edge image of the background from the edge image of the foreground object. Areas with remaining edges are considered to be the object area. In the final step, the object is recovered by using information from the object area and shadow attributes to construct the object. Foreground pixels with intensity values greater than the background, or those with characteristics different from the shadow attributes, are preserved. To preserve self-shadow areas (areas where a vehicle

casts a shadow over part of itself), where pixels have similar characteristics as the cast shadow, pixels close to the object area are also preserved. Finally, any holes in the object area are filled. Scant experimental results were provided, limiting the ability to evaluate this method. Furthermore, the method used to find the direction of illumination may fail in cases where the shadow has a halo effect at the edge (pixels of high intensity at the boundary of the shadow).

An excellent survey and evaluation of many moving-shadow removal algorithms can be found in Prati et al. (12).

METHODOLOGY

As noted by Wang et al. (5), shadows may consist of two components: self shadow and cast shadow. Cast shadows are shadows cast by the object of interest upon other objects and are the type that typically comes to mind when we think of shadows. Alternatively, self shadow is the shadow an object may cast upon itself when hidden from the illumination source. These cases are illustrated in Figure 1. The distinction between these types of shadows is important for object recognition – successful shadow removal aims to remove cast shadows while recognizing self shadow as part of the object of interest and therefore preserving them. Unfortunately, the similar characteristics between cast and self shadows often complicate a computer's ability to distinguish between them.



FIGURE 1 Illustration of self shadows and cast shadows.

The unique features of traffic are also of considerable importance. Any application of these algorithms must be capable of being applied in real-time in order to provide the maximum possible benefit to agencies; this is particularly important if the information is used as a data source for Advanced Traveler Information Systems (ATIS). This real-time requirement limits not only the computational complexity of the algorithms; it also precludes the opportunity to use training data and most learning-based approaches. The transient nature of traffic also enables relatively easy construction of a background image (barring congestion) that can serve as a base condition when evaluating a scene for shadows. This is a marked difference from many shadow removal techniques in the computer science domain which operate on a still image.

Several methods were investigated in this research and will be presented in the following subsections. In each of these methods it is assumed that a region representing the vehicle and any associated shadow has already been identified via background subtraction or motion analysis, denoted as the vehicle-shadow region. Furthermore, it is assumed that a background image without vehicles is also available, such as one obtained through the processing of a series of images from the same scene and extracting an image based on median or mode pixel values (13).

Dual-Pass Otsu Method

The first proposed method attempts to remove shadows by using the Otsu automating thresholding method (11) to discern between types of shadows based on pixel intensities. The Otsu method is designed to select the optimum threshold for separation into two populations based upon maximizing the variance between them. The first application of the Otsu method separates the pixels into high and low intensity populations, with the high intensity population considered to be the vehicle of interest. However, since the cast shadow and self shadow typically are more similar in intensity than the remaining vehicle, the low intensity population likely contains both the cast shadow and self shadow. To separate these two shadow types, a second thresholding of the lower intensity population is performed. Those pixels above the resultant threshold are considered to be self-shadow pixels and included as part of the vehicle, while those pixels occupying the absolute lowest pixel range are considered to be the cast shadow. These cast shadow pixels are then replaced by the corresponding pixels in the background image:

$$\forall p_{i,j} \in S_c : fore_{i,j} = bgd_{i,j} \quad (1)$$

Where:

$p_{i,j}$ represents a pixel in the vehicle-shadow region at indices i,j

S_c is the set of all pixels belonging to the cast shadow group

$fore_{i,j}$ is the intensity of the foreground pixel at indices i,j

$bgd_{i,j}$ is the intensity of the background pixel at indices i,j

Region Growing

Region growing is a technique widely used in image segmentation applications for computer vision systems. Although we are not interested in segmenting the entire image, this concept can still be applied to segment the cast shadow and remove it while preserving the vehicle and self shadow. Shapiro and Stockman (14) state that “a region grower begins at a position in an image and attempts to grow each region until the pixels being compared are too dissimilar to the region to add them.” One characteristic of region growing is that the statistics used for determining membership in a region are updated each time a new member is added to the region. Since shadow areas are more homogeneous in terms of intensity than most vehicles, use of the region growing method to identify the shadow region cast by a vehicle appears to be a good fit.

In order to perform region growing, suitable seed pixels representing the shadow must first be chosen. In practical applications, a sample shadow can be chosen through interactive input from the users. However, here we present an automated procedure for finding seed pixels of cast shadows. This may be accomplished by choosing pixels near the identified object boundary in the direction of the cast shadow. If the orientation of the image is known, the shadow direction can be automatically computed based on the position of the sun, since it is assumed that the sun is the only source of illumination during daytime observation. The position of the sun can be calculated given the time of day and the approximate latitude and longitude of the location under study (15). Once the sun’s location is known, the pan angle of the camera view with respect to due north is all that is necessary to get the direction of the shadow in the image:

$$\alpha = \theta_c - \theta_{azim} + \frac{\pi}{2} \quad (2)$$

Where:

α is the image shadow angle in radians counterclockwise from the positive x-axis

θ_c is the camera pan angle in radians counterclockwise from due north

θ_{azim} is the sun azimuth angle in radians counterclockwise from due south

Once seed pixels have been selected, the region growing algorithm is applied. A pixel neighboring the seed pixel is included in the shadow group if it lies within a specified confidence interval of the sample population mean:

$$|p_{i,j} - \bar{x}| < \frac{Zs}{\sqrt{n}} \quad (3)$$

Where:

$p_{i,j}$ represents a neighboring pixel in the vehicle-shadow region at indices i,j

\bar{x} is the mean intensity of all pixels in the shadow region

Z is the Z-statistic corresponding to the chosen confidence level for a standard normal distribution

s is the standard deviation of the pixel intensity of all pixels in the shadow region

n is the number of pixels in the shadow region

Once a pixel is accepted as part of the shadow region, the mean and standard deviation are recalculated for the next comparison. The region growth is complete once all remaining neighboring pixels exceed the threshold, thus no longer satisfying the requirements for membership in the shadow region. The pixels in the resulting shadow group are replaced by the corresponding pixels in the background image in the same manner as the previous method.

Edge Subtraction and Morphology

The final shadow removal method investigated in this paper considers a uniquely different approach from the first two. Instead of trying to identify the shadow region based on pixel intensities, this method looks for an area with few edges or edges with high similarity to the background edges. The Canny edge detection method (16) is utilized to produce an edge image of the vehicle-shadow region. For this research, a 7x7 Gaussian kernel corresponding to a standard deviation of 1.0 was used for smoothing, and the low and high thresholds for hysteresis were 1.2% and 2.4%, respectively. Increasing the kernel size results in more smoothed intermediate image at the expense of increased processing times; altering the hysteresis thresholds affects the sensitivity with respect to detected edges.

The area with many edges is assumed to be the vehicle, whereas the remaining area is shadow. However, in cases where background features such as cracked pavement or pavement markings are still visible in the shadow area, they may result in extraneous edges that do not belong to the vehicle and result in sub-optimal shadow recognition. To account for this effect, an edge image of the background must also be obtained. Once obtained, the edge image of the background can be produced and subtracted from the foreground edge image to produce an edge image of the vehicle-shadow region as follows:

$$edge_{pxq} = Canny(img_{pxq}) - Canny(bgd_{pxq}) \quad (4)$$

Where:

$edge_{pxq}$ is the resultant edge image

img_{pxq} is the foreground edge image

bgd_{pxq} is the background edge image

The resultant binary edge image will now be black in areas where few edges are present or where the edges that were present resembled the background. This image is morphologically processed to close gaps in the edge lines and patch any holes in the lines and prepare it for shadow detection:

$$R = ((E \oplus S) \oplus S) \ominus S \quad (5)$$

Where:

R is the resulting binary edge image

E is the subtracted binary edge image of the vehicle-shadow region

S is a binary 3x3 structuring element composed entirely of ones

\oplus represents the dilation operation

\ominus represents the erosion operation

More details about morphology operations are available in Shapiro and Stockman (14). Figure 2 illustrates the resulting edge image along with images from intermediate steps. The shadow location in the image is found in the following manner. First, the angle of the shadow in the image is found in the same manner as it was in the region growing method. The centroid of the vehicle-shadow region is also found:

$$centroid = \left(\frac{\sum p_x}{A}, \frac{\sum p_y}{A} \right) \quad (6)$$

Where:

p_x is the x-coordinate of a pixel in the vehicle-shadow region

p_y is the y-coordinate of a pixel in the vehicle-shadow region

A is the pixel area of the vehicle-shadow region

A line is then drawn from the centroid in the direction of the shadow angle. The point of intersection between this line and the outer edge of the vehicle-shadow region is the beginning of the shadow area. The shadow region is then formed from the collection of all eight-connected points neighboring the initial pixel. An equivalent mathematical expression is the collection of all connected pixels where the sum of the binary value of $p_{i,j}$ itself and all its neighbors is zero:

$$\sum_{i-1}^{i+1} \sum_{j-1}^{j+1} p_{i,j} = 0 \quad (7)$$

The shadow is then removed by replacing the pixels in the foreground image that are in the shadow region with corresponding pixels from the background image as done in the previous methods.

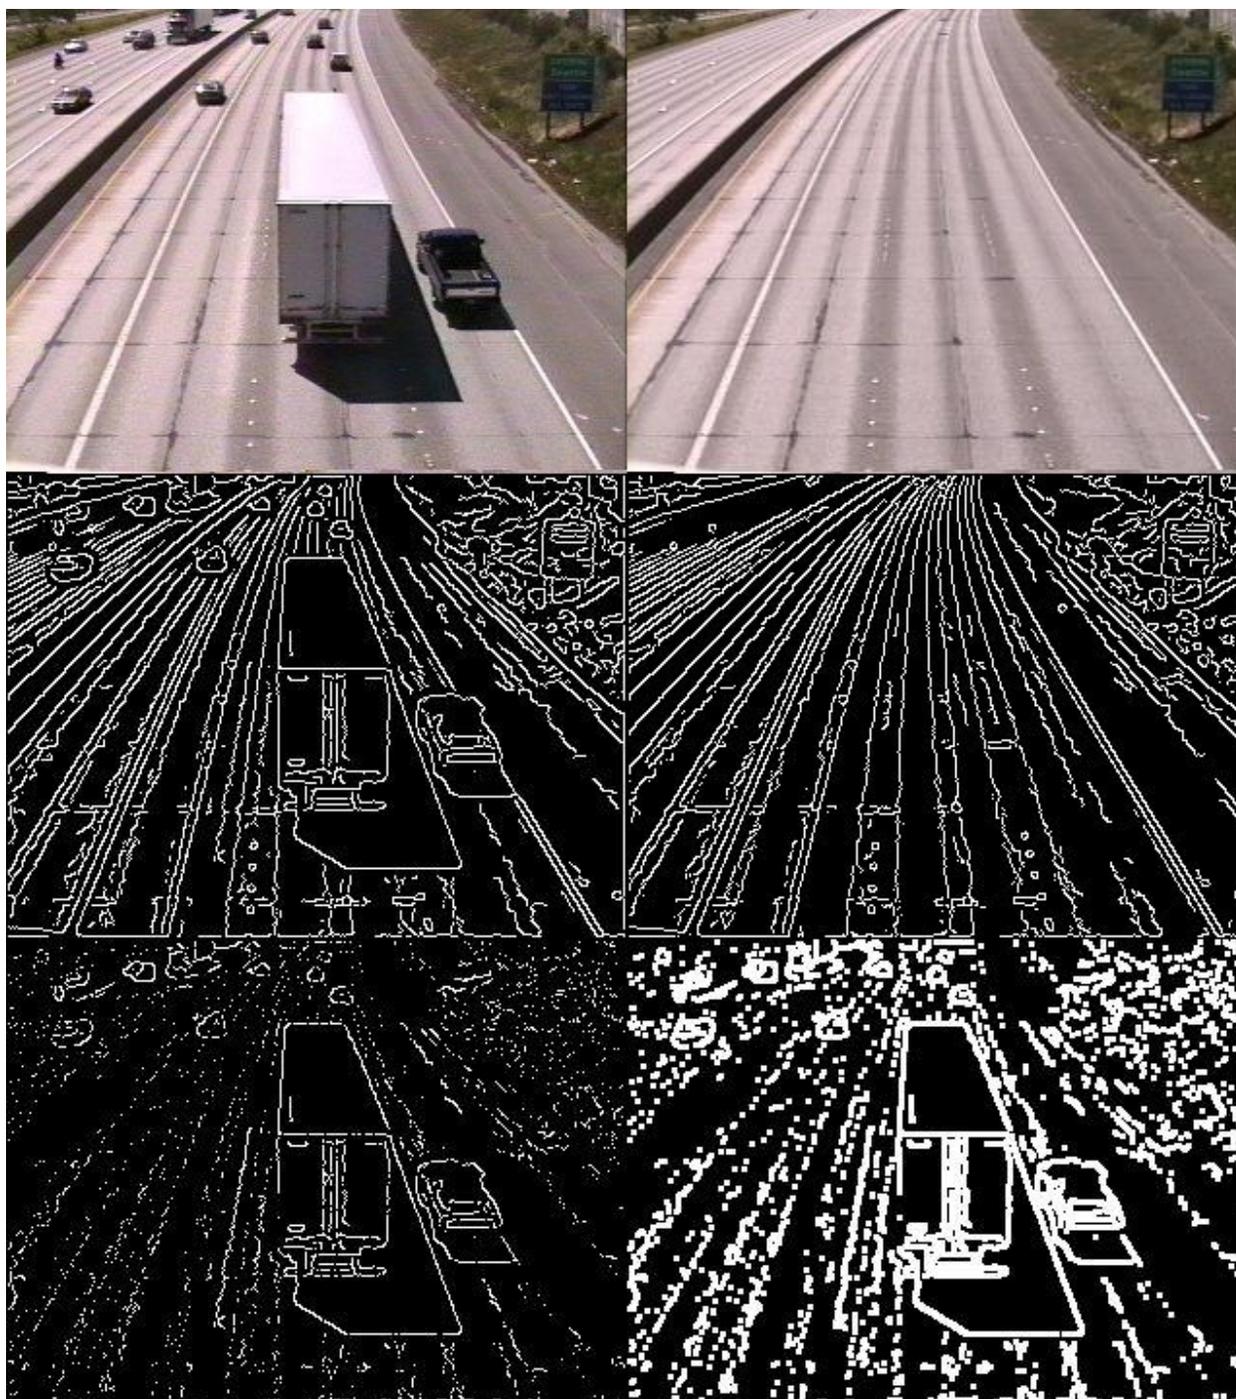


FIGURE 2 The foreground and background images (top), their respective edge images (middle), and the subtracted and dilated image (bottom).

EXPERIMENTAL RESULTS

Testing of the three methods was performed on digitized video from two locations in Seattle, Washington: I-5 southbound at 145th Street, recorded between 11:30 AM and 12:30 PM on June 11, 1999, and State Route 99 at 65th Ave N, recorded between 4:00 PM and 5:00 PM on April 22, 1999. The video was digitized at a rate of 10 frames per second, since that was more than adequate enough to ensure detection of every vehicle. A program (*I*) written in Microsoft C# was used to

extract a background image and process the digital video stream to produce classified vehicle counts by lane at each location. This processing was done in real time (10 frames per second) for the duration of the digitized video.

Each method placed different computational requirements on the test system, which was an AMD 64 3400+ processor with 1 GB PC 3200 SDRAM running Windows XP Professional with Service Pack 2. The actual amount of processing power required increased non-linearly with increasing relative size of the detected object. Thus, relative comparisons were made based upon the amount of processing time required to iterate through 100 repetitions of a detection. The dual-pass Otsu method was used as the benchmark, since it was easily the most inexpensive in terms of processing horsepower. The region growing method used 2.3-4.8 times more processing power, whereas the edge subtraction method required 3.4-7.5 times the processing time. Alternatively, memory consumption between the three methods did not appear to be an issue, as all three methods ran for extended time periods without consuming large amounts of memory.

Qualitative descriptions of the results of each method are provided below. Representative images displaying the results of shadow detection and removal are presented and grouped by the method used to detect the shadow.

Dual-Pass Otsu Method

The principal advantage of this algorithm lies in its computational inexpensiveness. Figure 3 demonstrates that although this method performs well for bright vehicles with dark cast shadows, it does not perform nearly as well when darker vehicles are considered.



FIGURE 3 Dual-pass Otsu method for a light (left) and dark (right) vehicle.

One notes that in the right image, self shadow regions of the pickup truck are classified as shadow and subsequently replaced by pixels from the background. This problem can be partially mitigated by allowing only points connected to the exterior edge of the vehicle-shadow region to be classified as shadow. The algorithm also does not perform well in lower-angle illumination cases where the cast shadow is not uniform and occupies relatively higher intensity ranges. This case is illustrated in Figure 4, where the implicit assumption of the shadow being the region with the lowest intensity is violated.



FIGURE 4 Dual-pass Otsu method in low-angle illumination.

Furthermore, this method is not spatially aware; a pixel may be classified as a shadow regardless of its neighboring pixels, thereby ignoring valuable information. Nonetheless, this method is quick and effective when a vehicle occupies a high intensity range. An intensity histogram of the vehicle-shadow area could be used to determine whether this approach would be adequate; if not, a more in depth shadow removal operation could be chosen.

Region Growing

Figure 5 exemplifies the sound performance of this method when the shadows are of uniform intensity, regardless of whether the vehicle color is light or dark.



FIGURE 5 Region growing method for a light (left) and dark (right) vehicle.

Like the previous method however, this method encounters difficulty when the shadows are non-uniform. A common cause is low-angle illumination, where the higher shadow intensity enables other features such as pavement markings, wear, and cracking to become visible, affecting the region growth. Figure 6 illustrates the poor performance in such a condition.



FIGURE 6 Region growing method in low-angle illumination.

Selection of a broader threshold for inclusion in the shadow group could compensate for this effect; in this case an automatic method to determine the appropriate threshold must be developed. Alternatively, the identification of edges in the background image (which would correspond to pavement markings and cracks) could be used to allow the region growing method to 'hop over' these pixels to continue shadow identification.

Edge Subtraction and Morphology

This method performed the best of the three methods presented. Figure 7 displays the results of this method when applied to scenes containing light and dark vehicles.



FIGURE 7 Edge subtraction and morphology method for a light (left) and dark (right) vehicle.

By not depending directly on shadow pixel intensities, this method was able to avoid many of the problems associated with non-uniform shadows and low-angle illumination, as evidenced by Figure 8. Given this method's flexibility, it was chosen for implementation in a prototype video image processing system, the Video-based Vehicle Detection and Classification (VVDC) system (17). Although the cited paper's study did not directly measure the performance of the shadow removal component, the overall success rate of detection was above 97%, while the truck detection error was less than 9% in all tests.



FIGURE 8 Edge subtraction and morphology method results in low-angle illumination.

The principal drawback to this method is that it is the most computationally expensive of all the proposed methods, as indicated earlier. Although the edge image for the background would only have to be extracted whenever the background was updated, continuous vehicle detections and production of edge images can strain a real-time system, particularly when large vehicles are detected. Thus, it is proposed that when a less powerful computer is used, a screening method may be necessary to choose the appropriate technique depending on the conditions; vehicles with uniform shadows and high average intensities can utilize the dual-pass Otsu method, while other vehicles with uniform shadows can implement the region growing method. Only the more complex shadows involving low-angle illumination would require the edge subtraction technique, thereby reducing the computing demands placed on the system.

Several challenges that are relatively independent of the method used to remove the shadow also arose. First, when a boundary between the cast and self shadows is not discernible, there is a tendency to overestimate the shadow area. This can be mitigated by limiting the allowable area of the shadow based on knowledge of the vehicle geometry and the angle of illumination. Secondly, the current work assumes that vehicles are not connected. In cases where the shadow of one vehicle is connected to another vehicle, the region growing and edge subtraction methods only remove the shadow associated with the last-encountered vehicle. This problem is best solved by separating vehicle objects before shadow removal based on geometry and lane position.

CONCLUSIONS

The attractiveness of employing traffic surveillance cameras as a data collection tool is becoming increasingly evident to many transportation agencies. The rich data they provide and their ability to monitor several lanes at once make them even more desirable to cost-conscious public agencies. Effective implementation of such systems, however, requires that a number of obstacles be overcome. Principal among these concerns is that of shadow removal – if not addressed, considerable mis-detection and misclassification will likely result. This paper proposed three methods suitable for implementation in a real-time traffic analysis scenario. Each method outperformed the previous method but also had increased computational demands. Furthermore, the edge subtraction and morphology technique has been applied in a separate program with favorable results. The selection of the appropriate algorithm should therefore be tempered by the conditions in the scene for optimum performance and efficiency.

ACKNOWLEDGMENTS

This research effort was jointly sponsored by Washington State Department of Transportation (WSDOT) and Transportation Northwest (TransNow), US Department of Transportation University Transportation Center for Federal Region 10, at the University of Washington. The authors would like to acknowledge WSDOT and TransNow for their financial support.

REFERENCES

1. Avery, R. P., Y. Wang and G. S. Rutherford. Length-Based Vehicle Classification Using Images from Uncalibrated Video Cameras. *Proceedings of the 7th International IEEE Conference on Intelligent Transportation Systems*, 2004, pp. 737-742.
2. Wang, Y. and N. L. Nihan. Can Single-Loop Detectors Do the Work of Dual-Loop Detectors? *ASCE Journal of Transportation Engineering*, 129(2), 2003, pp. 169-176.
3. Scanlan, J.M., D.M. Chabries, and R.W. Christiansen. A Shadow Detection and Removal Algorithm for 2-D Images. *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, Vol. 4, 1990, pp. 2057-2060.
4. Fung, G.S.K., N.H.C. Yung, G.K.H. Pang, and A.H.S. Lai. Effective Moving Cast Shadow Detection for Monocular Color Traffic Image Sequences. *Optical Engineering*, Vol. 41, No. 6, 2002, pp. 1425-1440.
5. Wang, J.M., Y.C. Chung, C.L. Chang, and S.W. Chen. Shadow Detection and Removal for Traffic Images. *IEEE International Conference on Networking, Sensing and Control*, Vol. 1, 2004, pp. 649-654.
6. Gamba, P., M. Lilla, and A. Mecocci. A Fast Algorithm for Target Shadow Removal in Monocular Colour Sequences. *Proceedings of the International Conference on Image Processing*, Vol. 1, 1997, pp. 436-447.
7. Gu, X., D. Yu, and L. Zhang. Image Shadow Removal Using Pulse Coupled Neural Network. *IEEE Transactions on Neural Networks*, Vol. 16, Issue 3, 2005, pp. 692-698.
8. Hsieh, C., E. Lai, Y. Wu, and C. Liang. Robust, Real Time People Tracking with Shadow Removal in Open Environment. *5th Asian Control Conference*, Vol. 2, 2004, pp. 901-905.
9. Xu, L., J.L. Landabaso, and M. Pardas. Shadow Removal with Blob-Based Morphological Reconstruction for Error Correction. *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, Vol. 2, 2005, pp. 729-732.
10. Lo, B.P.L., S. Thiemjarus, and G. Yang. Adaptive Bayesian Networks for Video Processing. *Proceedings of the 2003 International Conference on Image Processing*, Vol. 1, 2003, pp. 889-892.
11. Otsu, N. A Threshold Selection Method from Gray-Level Histograms. *IEEE Transactions on Systems, Man and Cybernetics*, Vol. 9, No. 1, 1979, pp. 62-66.
12. Prati, A., I. Mikic, M.M. Trivedi, and R. Cucchiara. Detecting Moving Shadows: Algorithms and Evaluation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 25, No. 7, 2003, pp. 918-923.
13. Zheng, J., Y. Wang, N.L. Nihan, and M.E. Hallenbeck. Extracting Roadway Background Image: a Mode-Based Approach. In Press, *Transportation Research Record*, 2006.
14. Shapiro, L. G. and G. C. Stockman. *Computer Vision*. Prentice Hall, New Jersey, 2001, pp. 289-290.
15. Gronbeck, C. *SunAngle*, 2004, <http://www.susdesign.com/sunangle/>. Accessed Dec. 9 2005.
16. Canny, J. A Computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 8, No. 6, 1986, pp. 679-698.

17. Zhang, G., R.P. Avery, and Y. Wang. A Video-based Vehicle Detection and Classification System for Real-time Traffic Data Collection Using Un-calibrated Video Cameras. Submitted to the 86th Annual Meeting of Transportation Research Board. Washington D.C., 2006.