

Vehicle Detection in Far Field of View of Video Sequences

Itzik Klein, Tomer Toledo, and Sagi Filin

Detection and tracking of vehicles from video and other image sequences are valuable in several traffic engineering application areas. Most of the research in this area has focused on detection of vehicles that are sufficiently large in the image that they can be detected on the basis of various features. As a result, the acquisition is feasible on limited sections of roads and may ignore significant parts of the available image. A method for early detection of vehicles was developed and demonstrated. This method focused on tracking motion rather than vehicle objects. With the detection of motion, the actual shape and size of the objects become less of a concern, thereby allowing detection of smaller objects at an earlier stage. One notable advantage that early detection offers is the ability to place cameras at higher vantage points or in oblique views that provide images in which vehicles in the near parts of the image can be detected by their shape or features, whereas vehicles in the far view cannot.

Detection and tracking of vehicles from video and other image sequences are valuable to many application areas such as road safety, automatic enforcement, surveillance, and acquisition of vehicle trajectories for traffic modeling. An indication of the importance of vehicle tracking is the growing number of related projects reported in recent literature. Among these are research at the University of Arizona that used video cameras, both digital and analogue, mounted on helicopter skids for the acquisition of video sequences for traffic management purposes (1); a project at the Delft University of Technology in the Netherlands, where traffic monitoring from airborne platforms was applied using digital cameras (2); and work at the Berkeley Highway Laboratory in California (3) where several cameras were deployed in a single location to form panoramic coverage of a road section. A common thread in all these projects is the application of methods that consider images with high resolution and large scale, which readily provide information on salient vehicle features (e.g., color and shape). Under these assumptions, the detection is useful in the near field of view when images are taken from an oblique view or to data from low altitudes for horizontal images. As a result, the acquisition is feasible on limited sections of roads and may ignore significant parts of the available image. These parts of the image may be valuable for early detection of vehicles that enter image frames.

Faculty of Civil and Environmental Engineering, Technion–Israel Institute of Technology, Haifa 32000, Israel. Corresponding author: I. Klein, iklein@technion.ac.il.

Transportation Research Record: Journal of the Transportation Research Board, No. 2086, Transportation Research Board of the National Academies, Washington, D.C., 2008, pp. 23–29.
DOI: 10.3141/2086-03

In this paper, a method for early detection of vehicles is presented. By contrast with the common frame-based analysis, this method focuses on tracking motion rather than vehicle objects. With the detection of motion, the actual shape and size of the objects become less of a concern, thereby allowing detection of smaller objects at an earlier stage. One notable advantage that early detection offers is the ability to place cameras at higher vantage points or in oblique views that provide images in which vehicles in the near parts of the image can be detected by their shape or features while vehicles in the far view cannot. The ability to detect vehicles earlier and to cover longer road sections may be useful in providing longer vehicle trajectories for the purpose of traffic models development (e.g., 4–6) and for improved traffic control algorithms.

RELATED WORK

Various vehicle detection methods have been reported in the literature. A general subdivision of them shows three main categories, including (a) optical flow, (b) background subtraction, and (c) object-based detection.

Optical flow is an approximation of the image motion on the basis of local derivatives in a given sequence of images. That is, it specifies how much each image pixel moves between adjacent images. Bohrer et al. (7) used a simplified optical flow algorithm for obstacle detection. De Micheli and Verri (8) applied the method to estimate the angular speed of vehicles relative to the direction of the road as part of a system to detect obstacles. Batavia et al. (9) described another system for obstacle detection, which is based on the optical flow method, although it does not explicitly calculate the flow field. This system aims to detect approaching vehicles that are in the blind spot of a car. Wang et al. (10) pointed out that for vehicle detection, optical flow suffers from several difficulties such as lack of texture in the road regions and small gray level variations that introduce significant instabilities in the computation of the spatial derivatives.

Although most optical flow–based techniques require high computational effort, object recognition via background subtraction techniques usually require significantly lower computational effort. Cheung and Kamath (11) identified the following four major steps in background subtraction algorithms: preprocessing, background modeling, foreground detection, and data validation. They also identified two main categories of background subtraction methods: recursive and nonrecursive. Nonrecursive techniques usually estimate a single background image. For example, frame differencing uses the previous frame as the background model for the current frame (12). Alternatively, the background is estimated by the median value of each pixel in all the frames in the video sequence (13).

Variations on the median method were suggested by Cutler and Davis (14) and Cucchiara et al. (15). Toyama et al. (16) proposed to estimate the background using a one-step Wiener prediction filter. Elgammal et al. (17) proposed development of a probabilistic background model by estimating a nonparametric density function of the pixel value. Cheung and Kamath (11) showed that the complexity of building this background model is similar to that of the median method, but that the foreground detection is more complex with this approach. Recursive background subtraction techniques estimate a different background for each frame. Recursive methods require less storage, but are generally more computationally complex compared with non-recursive methods. Furthermore, errors in the background model affect the results for longer periods. Among the recursive background subtraction techniques are an adaptive implementation of the median value technique, as proposed by McFarlane and Schofield (18) and implemented for an integrated traffic and pedestrian model-based vision system in Remagnino et al. (19). Another method, which was first suggested by Karmann and von Brandt (20) and was implemented for traffic by Kilger (21) uses Kalman filtering, in which the object mask is obtained by comparing evolved background images with the corresponding original images. Another alternative is the mixture of Gaussian method, which estimates the foreground and background as a mixture of Gaussian distributions that are tracked simultaneously and their parameters updated online (22). Although this method is very popular, it was shown by Cheung and Kamath (11) that it is computationally intensive, and its parameters require careful tuning. In addition, it is very sensitive to sudden changes in lighting conditions.

Background subtraction methods define objects by learning the background. By contrast, object-based detection focuses on identifying the objects themselves. With these methods, vehicles are detected using models of their features and shape. Some of the algorithms that have been proposed focus on generic pose estimation with predefined shape (23, 24), whereas others focus on high-resolution ground views (25, 26). Zhao and Nevatia (27) detected cars in aerial images by examining their rectangular shape, front and rear windshields, and shadow. They achieved a detection rate of approximately 90% with 5% false detections. Kim and Malik (28) proposed three-dimensional vehicle detection based on line features followed by a probabilistic feature grouping. Viola et al. (29) proposed a pedestrian detection system that combines appearance and motion cues extracted from image pairs. Levin et al. (30) used a co-training algorithm to improve vehicle detection with unlabeled data. Bose and Grimson (31) used scene-specific context features, such as image position and direction of motion of objects. They showed that this can greatly improve classification. Chachich et al. (32) used color as an alternative feature for detection. Vehicles are initially detected in the frame by determining the probability that an object is not the same color as the road surface. Most object-based methods require large computational effort compared with background subtraction methods.

The methods discussed previously were designed for situations in which the vehicle object was large enough in the frame to be recognized. However when the vehicles are in low-resolution or a vehicle feature is not available, those methods are not applicable. The goal of this research is to introduce a far-view vehicle motion detection method, aiming to detect vehicles as early in the image sequence as possible and in the part of the frame where most of the commonly used features will still be unnoticeable. A study by Cho and Rice (33) has

dealt with this problem; however, they were interested only in extracting speed profiles and so designed an algorithm that avoids detection and tracking of individual vehicles.

EARLY VEHICLE MOTION DETECTION

As the literature review shows, most of the reported approaches focused on the detection of well-defined vehicles and usually within a single frame. Most commonly, cameras are placed such that their view is against the direction of traffic flow. Vehicles in the far view field are characterized by a small size of only a few pixels. In this case, standard frame-based detection methods are of limited use, because objects of this size usually cannot be differentiated from noise. Therefore, the detection is limited to the near to middle field of view and vehicles are detected later compared with the time they actually enter the frame.

This paper addresses this problem in the case of oblique video cameras, which is often the viewing scheme for traffic surveillance cameras. Figure 1 demonstrates this situation, showing data acquired by the University of California at Berkeley Highway Laboratory (34). The oblique camera geometry dictates that vehicles in the far view, which is marked by a box in the upper left corner of the frame, cannot be seen properly. Although this part of the view takes up only a small portion of the frame, it amounts to approximately 50% of the actual road length covered in the frame. Therefore, algorithms that focus only on the nearer part of the image result in substantial delay in the detection of vehicles.

The proposed method focuses on early vehicle motion detection in the far view field. It is composed of five steps, as shown in Figure 2. Considering a video sequence as an input, frame differencing is applied first to obtain background and foreground estimation for each frame. Blobs that remain in the foreground frames may be noise, vehicles, or other objects. A frame summation scheme is used to distinguish blob vehicles from other objects and noise. In this method, several foreground frames are summed to observe the movement of blobs. Information on the road layout and on traffic parameters are then used to identify those blobs whose movement is consistent with that of vehicles in traffic. The algorithms used in this step will be discussed in further detail in the next section. Finally, the vehicles that were detected may be tracked between frames in a similar manner. The working assumptions underlying the proposed detection method are as follows:

1. The imaging geometry is estimable by estimating the camera external orientation parameters, through the camera-to-road homog-



FIGURE 1 Sample frame from Berkeley Highway Laboratory Video.

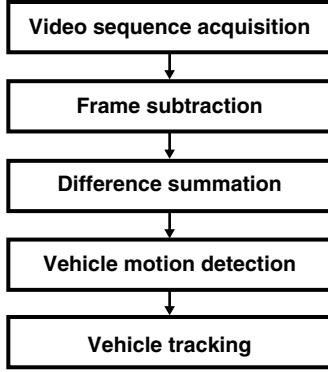


FIGURE 2 Flow chart of early vehicle motion detection method.

raphy or by vanishing point orientation estimation [e.g., Hartley and Zisserman, (35); Kim (36)].

2. The traffic flow direction in each lane is known. This will help to distinguish noise from vehicle blobs. Given the camera position and road geometry, typical vehicle dimensions can be estimated. In addition, given traffic flow characteristics, the average spacing between two vehicles can be calculated. These properties will be used in selecting the number of frames to be summed in the analysis.

As noted previously, the detection here concerns with objects whose size in the image is a few pixels, and feature-based methods are likely to fail. Therefore a frame differencing method was applied as the first step in the identification of motion. Let the video sequence consist of N frames and denote the pixel intensity as $I(x, y, k)$ where (x, y) are the spatial coordinates of a certain pixel in the frame and k is the frame index. To simplify the presentation, the spatial coordinates are removed from the intensity notation in the following. The time difference between subsequent frames is assumed to be constant and known. The difference pixel intensity image, $B(k)$, between two subsequent frames is defined by

$$B(k) = |I(k+1) - I(k)| \quad k = 1, 2, 3, \dots, N-1 \quad (1)$$

From the difference pixel intensity sequence the summed image, S , is obtained by summation

$$S = \sum_{j=1}^M B(j) \quad (2)$$

where M is the number of frames to be summed.

In S , objects that have a distinct motion will feature a linear trend, whereas random noise will not accumulate into a significant shape. The summation outcome and the quality of the detection depend on the value of M . On the one hand, if it is set too low, the summed image may contain significant noise, which will make the distinction between noise and vehicles blobs difficult. On the other hand, if it is set too high, the accumulated signatures of vehicles may over-

lap, which will result in detection of multiple vehicles as a single unit. This problem is illustrated in Figure 3. Figure 3a shows the first frame to be summed, with three vehicle blobs labeled. The images in Figures 3b through 3f show the result of the summation of a varying number of frames—2, 4, 8, 12, and 16, respectively. All three vehicles are detectable when 4, 8, or 12 frames are summed. But, Vehicle 3 is not detected when only 2 frames are summed, and when 16 frames are used, the blobs of Vehicles 1 and 3 overlap, which leads to their detection as a single vehicle.

To avoid this problem, it is necessary to determine the maximum number of frames to be summed. The value of M may be set constant for all frames in the sequence or calculated adaptively for each frame. A constant value for M offers computational benefits at the cost of potentially lower detection rate compared with an adaptive approach.

An adaptive value for M for a specific frame may be calculated as follows. First, the frame difference in question is summed with the subsequent one using Equation 2. The resulting summed image may include blobs that are vehicles and others that are noise. Under the conservative assumption that all blobs are vehicles, an upper bound on the possible number of vehicles is obtained. Using the measurements of the offset, Δx , of each blob between frames, and the elapsed time between frames, Δt , the average travel speed of these blobs can be estimated as follows:

$$\bar{v} = \frac{\sum_{j=1}^n \Delta x_j}{n \Delta t} \quad (3)$$

where, n is the number of blobs in the summed image.

Using the information on the road layout, the minimum distance between blobs in each lane is calculated. The minimum distance value among all lanes defines the overall minimum blob distance d_{\min} . The maximum permissible time interval of the frame summation to avoid overlap among any of the blobs can be calculated as follows:

$$\Delta t_{\max} = \frac{d_{\min}}{\bar{v}} \quad (4)$$

Finally, the permissible maximum number of frames to be summed is given by

$$M = \left\lfloor \frac{\Delta t_{\max}}{\Delta t} \right\rfloor \quad (5)$$

This calculation may be conducted for each frame separately, yielding different M values for each frame. Calculation of a constant value of M may be conducted in a similar fashion.

In the following, details concerning the implementation of the detection method are discussed. After its creation, the summed image, S , is enhanced using morphological operations. The methods applied are top-hat and bottom-hat filtering (37). Top-hat and bottom-hat filtering subtract the open image and the closed image from the original one, respectively. Next, image segmentation is performed using marker-controlled watershed segmentation (38) to isolate the objects in the summed image from the background. In some cases when applying the watershed transformation, oversegmentation occurs.

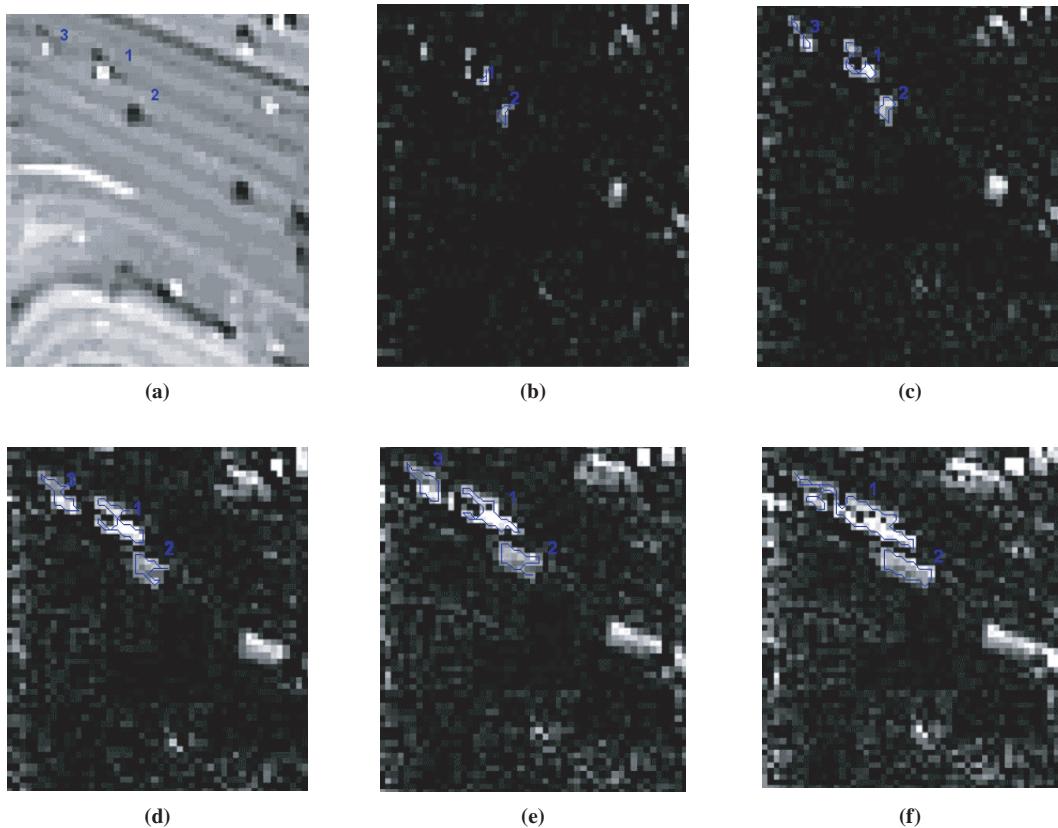


FIGURE 3 Vehicle blobs in the difference summation: (a) first frame to be summed, with three vehicle blobs labeled; (b) summation result of two frames; (c) summation result of four frames; (d) summation result of eight frames; (e) summation result of 12 frames; and (f) summation result of 16 frames.

To overcome this concern, markers (connected components) are used. There are two types of markers: internal markers, which are inside each of the objects of interest, and external markers, which are contained within the background. Various methods for computing internal and external markers, involving both linear and nonlinear filtering, have been proposed (37). In this implementation, a condition on the marked blob size is added to the watershed procedure. If the blob is smaller than a characteristic area of a vehicle, it is removed from the list of blobs that are vehicle candidates.

After the segmentation phase, blobs in the image are identified. Several criteria are used to determine those that are vehicles: The blob should be within the lane, and the direction of the blob growth should be in line with the direction of the traffic. When those criteria are met, the blob is considered a vehicle.

RESULTS AND DISCUSSION

Video data collected by the Berkeley Highway Laboratory (34) was used to demonstrate the proposed method. The early vehicle motion detection method was applied to the part of the frame in which vehicles first appear. The results reported were obtained using a constant number of frames differences that were summed in each case. The calculations to determine this value, as discussed previously, yield a value of eight frames. The sensitivity of the results to this value was also examined.

Figure 4 shows examples of the application of the method to two different frames in the video sequence, one at the top and one at the bottom of the figure. Figure 4a shows the two frames being studied. In each case, eight frames are summed. The last of these frames is shown in Figure 4b. The resulting summed images are shown in Figure 4c. The blobs in the first and last images were marked: Vehicles are circled, and noise blobs are marked with squares. The blobs that were detected were marked in the summed image. The noise blobs were identified as such and removed during the differencing procedure. The vehicle blobs were detected in the analysis. As expected, no overlapping of blobs occurred with the number of frames that were summed. The summation of 8 frames implies that given a video sequence with a rate of 24 frames per second, vehicles may be identified 0.33 s after they enter the frame. Methods that operate on a vehicle feature can be used only when the vehicle gets closer to the camera. In this sequence, the time that elapses from the time the vehicle enters the frame to the time it is detectable on the basis of its features is approximately 7 s.

The overall detection results for a sample of 185 vehicles in the video sequence, as a function of the number of frames, M , that are summed in each case are presented in Table 1 and Figure 5. Vehicles in the images may not be correctly detected in two ways: They may be missed altogether by the algorithm, or their blobs may overlap with other vehicles and so not be detected as individual vehicles. The probabilities of occurrence of the two error types vary in opposing directions when M is increased: the number of missed vehicles

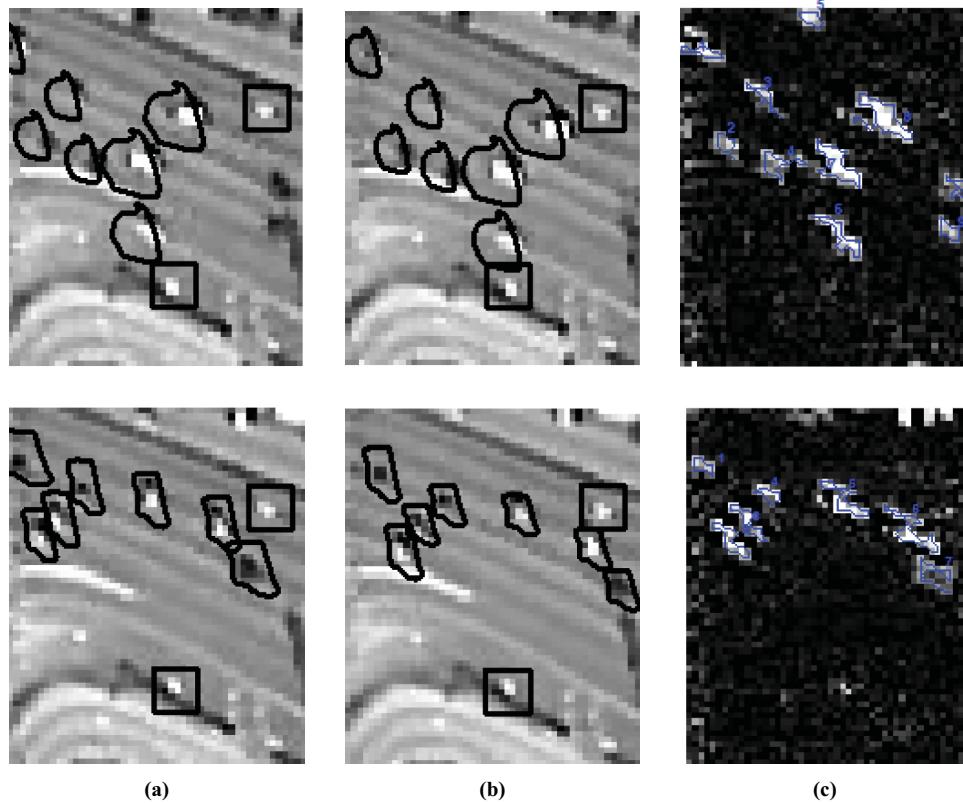


FIGURE 4 Example of application of early vehicle motion detection: (a) the two frames being studied, (b) the last summed frame, and (c) the resulting summed image.

decreases, whereas the number of overlapping vehicles increases. The selection of the parameter M should reflect the trade-off between these two errors. In this sequence, the optimal detection rate is highest (91.4%) when eight frames are summed, which is the value recommended by the analysis in the previous section. Detection rates decrease substantially when M is increased or decreased. A third source of errors is the false detection of nonvehicle blobs as vehicles. The occurrence of this error follows a similar trend to that of errors in detecting vehicles and is also lowest when eight frames are summed.

SUMMARY

This paper presented a new approach for vehicle detection in video image sequences. The method is based on detection of the movement of vehicles between frames. It does not rely on detection of vehicle features and so may be used in cases in which vehicles occupy a small number of pixels in the image. Cameras at an oblique view enable detection of vehicles in the far view whose features are not identifiable and so are not normally detectable with other methods. The early detection of vehicles can significantly increase the lengths of the sections of roads that can be monitored by cameras.

Further study of this method is underway. Among the questions that may be studied are the ability to apply the method to extract vehicle classification information, the handling of shadows, adaptation of the method to aerial video sequence, and the integration and handing over of information between this method and methods that are more appropriate for detection in the near-view of the image.

ACKNOWLEDGMENTS

This work was partially supported by grants from the Henry Ford II Transportation Research Fund, the Ran Naor Foundation, and an EC Marie Curie Grant. The authors thank the anonymous referees for comments that helped improve the quality of this paper.

TABLE 1 Summary of Detection Results

M	Vehicles Detected	Vehicles Not Detected			Total Errors
		Missed Vehicles	Overlapping Vehicles	False Detections	
1	86	99	0	16	115
3	142	42	1	31	74
5	154	21	10	9	40
8	169	6	10	4	20
11	158	3	24	19	46
15	151	0	34	48	82
18	130	1	54	53	108

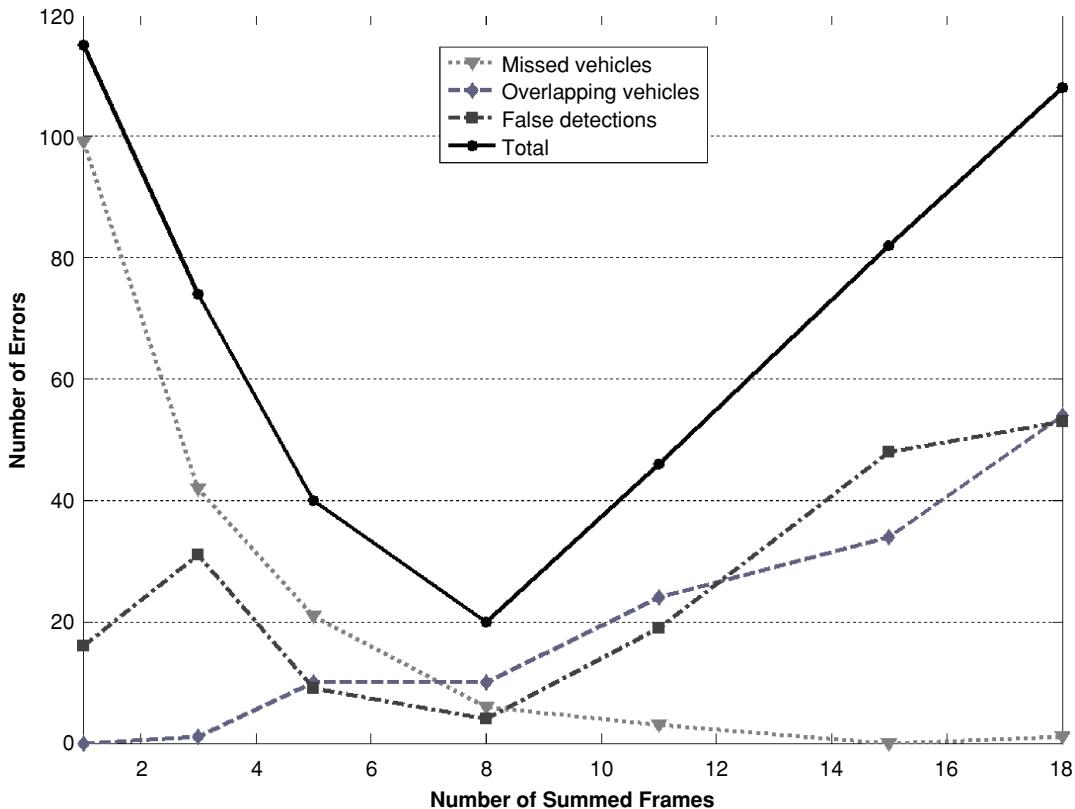


FIGURE 5 Impact of number of summed frames on detection errors.

REFERENCES

- Angel, A., A. M. Hickman, P. Mirchandani, and D. Chandnani. Methods of Analyzing Traffic Imagery Collected from Aerial Platforms. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 4, No. 2, 2003, pp. 99–107.
- Nejadasl, F. K., B. G. H. Gorte, and S. P. Hoogendoorn. Robust Vehicle Tracking in Video Images Being Taken from a Helicopter. *Proc., ISPRS Commission VII Mid-Term Symposium*, ISPRS Commission VII Mid-Term Symposium Remote Sensing: From Pixels to Processes, Enschede, Netherlands, 2006, pp. 528–533.
- Chen, C., D. Lyddy, and B. Dundar. *Berkeley Highway Lab Video Data Collection*. PATH Task Order 4129, Final Report. California Partners for Advanced Transit and Highways (PATH), University of California, Berkeley, 2005. <http://repositories.cdlib.org/its/path/papers/UCB-ITS-PWP-2004-8>. Accessed August 25, 2008.
- Beymer, D., P. McLauchlan, B. Coifman, and J. Malik. A Real Time Computer Vision System for Measuring Traffic Parameters. *Proc., IEEE Conference on Computer Vision and Pattern Recognition*, San Juan, Puerto Rico, 1997, pp. 495–501.
- Koller, D., J. Weber, and J. Malik. Robust Multiple Car Tracking with Occlusion Reasoning. *Proc., 3rd European Conference on Computer Vision*, Vol. 1, Springer, Berlin/Heidelberg, Germany, 1994, pp. 189–196.
- Kim, Z., and J. Malik. Fast Vehicle Detection with Probabilistic Feature Grouping and Its Application to Vehicle Tracking. *Proc., 9th IEEE International Conference on Computer Vision*, Vol. 1, IEEE Computer Society, Washington, DC, 2003, pp. 524–531.
- Bohrer, S., M. Brauckmann, and W. von Seelen. Visual Obstacle Detection by a Geometrically Simplified Optical Flow Approach. *Proc., 10th European Conference on Artificial Intelligence* (B. Neumann, ed.), John Wiley & Sons, Inc. New York, 1992, pp. 811–815.
- De Micheli, E., and A. Verri. Vehicle Guidance from One Dimensional Optical Flow. *Proc., IEEE Intelligent Vehicles Symposium*, Atlanta, Ga., 1993, pp. 183–189.
- Batavia, P., D. Pomereau, and C. Thorpe. Overtaking Vehicle Detection Using Implicit Optical Flow. *Proc., IEEE Intelligent Transportation Systems Conference*, Boston, Mass. 1997, pp. 729–734.
- Wang, J., J. Bebis, and R. Miller. Overtaking Vehicle Detection Using Dynamic and Quasi-Static Background Modeling. *Proc., IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Vol. 3, IEEE Computer Society, Los Alamitos, Calif., 2005, pp. 64–71.
- Cheung, S. C., and C. Kamath. Robust Techniques for Background Subtraction in Urban Traffic Video. *Proc., SPIE*, Vol. 5308, Bellingham, Wash., 2004, pp. 881–892.
- Puri, A., H. M. Hang, and D. Schilling. An Efficient Block-Matching Algorithm for Motion-Compensated Coding. *Acoustics, Speech and Signal Processing. Proc., IEEE International Conference on Acoustics, Speech and Signal Processing*, Vol. 12, IEEE, Washington, D.C., 1987, pp. 1063–1066.
- Lo, B. P. L., and S. A. Velastin. Automatic Congestion Detection System for Underground Platforms. *Proc., International Symposium on Intelligent Multimedia, Video and Speech Processing*, IEEE, Washington, D.C., 2001, pp. 158–161.
- Cutler, R., and L. Davis. View-Based Detection and Analysis of Periodic Motion. *Proc., 14th International Conference on Pattern Recognition*, IEEE, Washington, D.C., Vol. 1, 1998, pp. 485–500.
- Cucchiara, R., M. Piccardi and A. Prati. Detecting Moving Objects, Ghosts, and Shadows in Video Streams. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 25, IEEE, Washington, D.C., 2003, pp. 1337–1342.
- Toyama, K., J. Krumm, B. Brumitt, and B. Meyers. Wallower: Principles and Practice of Background Maintenance. *Proc., 7th IEEE Conference on Computer Vision*, Vol. 1, IEEE Computer Society, Los Alamitos, Calif., 1999, pp. 255–261.
- Elgammal, A., D. Harwood, and L. Davis. Non-Parametric Model for Background Subtraction. *Proc., 6th European Conference on Computer Vision—Part II*, Springer-Verlag, London, 2000, pp. 751–767.

18. McFarlane, N., and C. Schofield. Segmentation and Tracking of Piglets in Images. *Machine Vision and Applications*, Vol. 8, No. 3, 1995, pp. 187–193.
19. Remagnino, P., A. Baumberg, T. Grove, D. Hogg, T. Tan, A. Worrall, and K. Baker. An Integrated Traffic and Pedestrian Model-Based Vision System. *Proc., 8th British Machine Vision Conference*, British Machine Vision Association, University of Essex, Colchester, United Kingdom, 1997, pp. 380–389.
20. Karmann, K. P., and A. von Brandt. Moving Object Recognition Using an Adaptive Background Memory. In *Time-Varying Image Processing and Moving Object Recognition* (V. Cappellini, ed.), Elsevier, Amsterdam, Netherlands, 1990, pp. 289–296.
21. Kilger, M. A Shadow Handler in a Video-Based Real-Time Traffic Monitoring System. *Proc., IEEE Workshop on Applications of Computer Vision*, IEEE Computer Society Washington, D.C., 1992, pp. 1060–1066.
22. Stauffer, C., and W. Grimson. Learning Patterns of Activity Using Real-Time Tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 22, 2000, pp. 747–757.
23. Koller, D., K. Daniilidis, and H. H. Nagel. Model-Based Object Tracking in Monocular Image Sequences of Road Traffic Scenes. *International Journal of Computer Vision*, Vol. 10, No. 3, 1993, pp. 257–281.
24. Worrall, A., G. Sullivan, and K. Baker. Pose Refinement of Active Models Using Forces in 3D. *Proc., 3rd European Conference on Computer Vision*, Springer-Verlag, Secaucus, N.J., 1994, pp. 341–352.
25. Schneiderman, H., and T. Kanade. A Statistical Method for 3D Object Detection Applied to Faces and Cars. *Proc., IEEE Conference on Computer Vision and Pattern Recognition*, Vol. 1, IEEE Press, New York, 2000, pp. 746–751.
26. Bensrhair, A., M. Bertozi, A. Broggi, P. Miché, S. Mousset, and G. Toulminet. A Cooperative Approach to Vision-Based Vehicle Detection. *Proc., IEEE International Conference on Intelligent Transportation Systems*, Oakland, Calif., 2001, pp. 209–214.
27. Zhao, T., and R. Nevatia. Car Detection in Low Resolution Aerial Image. *Proc., IEEE International Conference on Computer Vision*, Vol. 1, Elsevier Science Inc. New York, 2001, pp. 710–717.
28. Kim, Z., and J. Malik. Fast Vehicle Detection with Probabilistic Feature Grouping and Its Application to Vehicle Tracking. *Proc., 9th IEEE International Conference on Computer Vision*, Vol. 1, IEEE Computer Society Washington, D.C., 2003, pp. 524–531.
29. Viola, P., M. Jones, and D. Snow. Detecting Pedestrians Using Patterns of Motion and Appearance. *Proceedings of IEEE International Conference on Computer Vision*, Vol. 2, IEEE Computer Society, Washington, D.C., 2003, pp. 734–741.
30. Levin, A., P. Viola, and Y. Freund. Unsupervised Improvement of Visual Detectors Using Co-Training. *Proc., IEEE International Conference on Computer Vision*, Vol. 1, IEEE Computer Society, Washington, D.C., 2003, pp. 626–633.
31. Bose, B., and E. Grimson. Improving Object Classification in Far-Field Video. *Proc., IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Vol. 2, IEEE Computer Society, Los Alamitos, Calif., 2004, pp. 181–188.
32. Chachich, A., A. Pau, A. Barber, K. Kennedy, E. Olejniczak, J. Hackney, Q. Sun, and E. Mireles. Traffic Sensor Using a Color Vision Method. *Proc., SPIE*, Vol. 2902, IEEE Computer Society, Los Alamitos, Calif., 1997, pp. 156–165.
33. Cho, Y., and J. Rice. Estimating Velocity Fields on a Freeway from Low-Resolution Videos. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 7, No. 4, 2006, pp. 463–469.
34. Next Generation Simulation Community Website. <http://ngsim.camsys.com>. Accessed July 31, 2007.
35. Hartley, R., and A. Zisserman. *Multiple View Geometry in Computer Vision*, (2nd ed.). Cambridge University Press, Cambridge, United Kingdom, 2004.
36. Kim, Z. Real-Time Road Detection by Learning from One Example. *Proc., 7th IEEE Workshop on Application of Computer Vision*, Vol. 1, IEEE Computer Society, Washington, D.C., 2005, pp. 455–460.
37. Gonzalez, R. C., R. E. Woods, and S. L. Eddins. *Digital Image Processing Using MATLAB*. Prentice Hall, Upper Saddle River, N.J., 2003.
38. Meyer, F., and S. Beucher. Morphological Segmentation. *Journal of Visual Communication & Image Representation*, Vol. 1, 1990, pp. 21–46.

The Intelligent Transportation Systems Committee sponsored publication of this paper.