Processing Traffic Data Collected by Remote Sensing

Victor L. Knoop, Serge P. Hoogendoorn, and Henk J. van Zuylen

Video data are being used more often to study traffic operations. However, extracting vehicle trajectories from video by current methods is a difficult process, typically resulting in many errors. The process requires extensive labor to correct the trajectories manually. This paper proposes a method to process video data from traffic operations. Instead of detecting a vehicle in each picture of the video separately, the video data are transformed so that the trajectories of the vehicles (their position over time) become visible in a single image. In this single image, the trajectories can be found by detecting lines. The difference from other methods is that trajectories rather than vehicles are detected. Trajectory (line) detection is more robust than vehicle (rectangle) detection; with this method, about 95% of the trajectories are detected correctly and, more important, the segments of each trajectory are much longer compared with results from other methods in the literature. Also, the detection is a quick process because only a single image is required to be analyzed. For a data set 5 min long, transforming costs several minutes, and automatically detecting and tracking costs 40 to 50 min per lane. Manual correction is then necessary, which costs about 10 min per lane. In contrast, with a different method the total processing time for analyzing traffic operations costs about 1 week for all lanes together.

Traffic congestion is a growing problem in modern society. Traffic engineers try to improve the situation, but knowledge of traffic characteristics in general and at bottleneck locations are crucial. In early days, most traffic-related studies were carried out with macroscopic data. The first attempt to find a relation between density and speed was made 75 years ago by Greenshields et al. (*1*). Chandler et al. argued about the necessity of data on individual driving behavior in 1958, using wired vehicles on a test track for their experiments (*2*). Treiterer and Myers chose a different approach and were among the first to capture traffic operations on video in 1974 (*3*). In that study, the data were processed manually.

With modern techniques, it is possible to gain better knowledge of traffic on this detailed level. Behavior of individual drivers can be observed. A detailed understanding of the individual driving behavior of road users can help improve traffic flow and thus reduce delays. The data for individual drivers (i.e., microscopic data) lead to insights into, for example, headway and traffic dynamics from which the capacity or stability of traffic flow can be derived. There exist different ways to collect microscopic data. This paper focuses on video data; the word "video" indicates a consecutive series of pictures showing traffic operations with a fixed time interval, regardless of picture size and frequency. The word "picture" here indicates a single time frame of the video.

The time-dependent coordinates of all vehicles give an overview of all traffic operations. From the video, these coordinates over time are known, which gives insights into the dynamics of driving. This paper discusses the process of extracting this information from given video images. The newly proposed method gives more robust results than existing ones, because other methods detect independent vehicles, leading to difficulty linking the vehicles in consecutive pictures to obtain vehicle trajectories. When the program fails to recognize a vehicle in a later picture, that particular vehicle needs to be indicated by the user. These manual corrections are time-consuming. In the method proposed here, the duration of the parts of found trajectories are much longer, which is an advantage. Therefore, in using the proposed method, much less time is needed to complete the set of trajectories manually. Another advantage of the newly proposed method is that it provides insight into traffic operations without actually tracking vehicles; that is, relevant properties of traffic are put directly into a single image, making the remote sensing data more easily accessible.

The next section gives an overview of data collecting and processing methods for individual vehicles reported in the literature. After this literature review is a brief description of how data captured from a (moving) helicopter are transformed into a series of stable images. The main section of the paper then describes how the video can be transformed and shows which operations can be applied to combine all relevant information in a single image. Then there is a discussion of how trajectories can be detected automatically from this (transformed) image. The focus in this process is on obtaining trajectory data. The final section summarizes the research and gives concluding remarks.

LITERATURE REVIEW

Microscopic Traffic Data

Many studies require traffic data at the level of individual vehicles, which can be collected in various ways. An exhaustive overview as well as the applications can be found in Appendix C in the thesis of Ossen (*4*). Only the most frequently used methods are briefly discussed here.

One way to collect microscopic data is to record the information obtained from (double) loop detectors in the road at the vehicle level, which provides passing times (and therefore headways) and speeds at a specific location (*5*). However, because the cars are not followed over time, information is not provided about the dynamics of traffic.

Faculty of Civil Engineering, Transport, and Planning, Delft University of Technology, P.O. Box 5048, 2600 GA Delft, Netherlands. Corresponding author: V. L. Knoop, v.l.knoop@tudelft.nl.

Transportation Research Record: Journal of the Transportation Research Board, No. 2129, Transportation Research Board of the National Academies, Washington, D.C., 2009, pp. 55–61.
DOI: 10.3141/2129-07

A second way to collect data is by the Global Positioning System (GPS), which can be used to record information on the speed and location of an instrumented vehicle over time. This system provides the dynamics for a particular instrumented car (*6*) but does not give insights into the interaction of vehicles. Only when there is a platoon of equipped vehicles is there insight into the relative positions and speeds of the vehicles. This situation is realized only in an experiment at a test track or in an artificially created situation on the road.

A third way to collect data is with an instrumented vehicle. For example, vehicles equipped with radar can provide a more representative view on vehicle dynamics (*7*). The radar can measure the distance between two vehicles (distance to leading vehicle, following vehicle, or both). Radar can be used to monitor the behavior of the driver of the equipped vehicle, which is called the active mode, or to monitor the behavior of the following vehicle, which is called the passive mode. Only a few drivers can be studied in the active mode, but every driver who follows the instrumented vehicle can be studied in the passive mode. The disadvantage of the passive mode is that the test vehicle influences the traffic. Also, the distances measured by radar can be combined with a GPS signal to observe the influences of the location on car-following behavior. The disadvantage of using an instrumented vehicle is that observations for different cars are made under different conditions because the instrumented vehicle is moving. Therefore, all observations are made at different locations, or—if the instrumented vehicle returns to the same location—observations are made at different times and under different traffic conditions.

A fourth way to collect data is to capture video from a high point of view, which provides insights into the behavior of drivers both longitudinally (e.g., car following) and laterally (e.g., lane selection). Consequently, one can analyze headway choice—as with using an instrumented vehicle—as well as the influence of the distance to the leader on acceleration. One advantage is that the information has both a spatial and a temporal dimension. Another advantage is that many drivers can be observed under similar driving conditions.

Collecting and Processing Video Data

A few years ago, projects started using remote sensing data on a larger scale—for example, tracing congestion dynamics at the Delft University of Technology or the Next Generation Simulation (NGSIM) project (*8*), which provides data on individual vehicles. These projects are used in the United States for a wide variety of applications from analyzing signalized intersections (*9*) to lane-change behavior (*10*) and emergency behavior (*11*). In Europe, these U.S. data are used in several projects, for example, to show properties of the fundamental diagram (*12*) or to calibrate car-following models (*13*).

Before these video data can be used, they need to be processed. Hoogendoorn et al. (*14*) proposed a method of first finding vehicles by comparing pixel values in a picture with the median value of that pixel (i.e., that position) over time, which can be interpreted as the brightness of the roadway background. Then, clusters of pixels deviating from the values identified for the background are marked as vehicles. This process is repeated for each picture. In the final step, vehicles are tracked by comparing the location of vehicles in two later images. A detection of more than 90% is claimed.

This method had some flaws when it was applied it to a data set with different shades of tarmac and different light intensities. In practice, it turned out that the trajectory of a vehicle was interrupted over time, making it necessary to combine partial short trajectories. A normal filter procedure would discard too-short trajectories, which could be interpreted as noise. Because the trajectories of vehicles were interrupted and short, such a filter procedure could not be applied now. Also, objects other than vehicles were detected (e.g., road lines), which could not be ignored and presented serious trouble with regard to the combination of detected parts as it was unclear whether the parts were vehicles or other objects. Thus, the main problem was tracking.

A similar methodology was stated by Angel et al. (*15*). In each picture of video taken from a helicopter, intensity was compared with a predefined pixel value. A concentration of deviating pixels is marked as a vehicle. In contrast to Hoogendoorn et al. (*14*), where only the information on the pictures was used to find reference points, Angel et al. (*15*) used added information collected during the flight: the position (GPS coordinates) and the angle of the helicopter were saved together with the video data. The main problem indicated here was in matching vehicles, that is, to connect the same detected vehicle over different frames. Angel et al. also reported a matching percentage of about 90% (*15*).

In the NGSIM project (*16*), a subproject deals with extracting the trajectory information from the NGSIM video (*17*). There, vehicles are found in each image, but origins and destinations are also provided. In the tracking process, this information on flow direction is used. The tool has the potential to track both forward and backward in time. No reliability value is given for the automated procedure.

Given the comparable methodology (*14, 15*), it is likely that similar problems occur. Also, the manual states that there is a need to postprocess the tracks manually, for which a tool is provided. This work is time-consuming: "Experienced NGSIM-VIDEO operators can effectively process video at a rate of 40–100 seconds of tracking per work day for a 500-meter freeway section" (*17*).

Cho and Rice (*18*) presented a method to estimate the velocity field from video data based on the image intensity in each lane. To find a vehicle, they took the maximum value for the image intensity across a single lane. This method detects bright vehicles, but darker vehicles cannot be identified because the maximum intensity on a cross section is the intensity of the pavement. However, that is not a problem for determining the velocity field, which can be derived from bright vehicles only. For the same reason, unmatched parts of a trajectory do not constitute a problem, because it is sufficient to use the parts that are available. A more detailed analysis requires the full trajectories of all vehicles.

There is therefore a need for a method that gives the trajectories of vehicles from remote-sensing data reliably for longer time intervals without much manual postprocessing.

TRANSFORMATIONS

This paper proposes reducing the problem with one dimension and capturing all relevant information in one image. This section explains which transformations can be made to construct these images.

The video needs to be stabilized if captured from a moving helicopter and rectified if the viewing angle is not perpendicular to the road surface. Because of space restrictions, the process is not discussed here. The starting point is a series of stabilized and rectified pictures. Usually, vehicles are identified in each picture, and the process is repeated for each picture. The word picture is used for a time slice of the video (snapshot of traffic conditions in one moment in time), whereas the word image is used for graphic representations that are constructed otherwise.

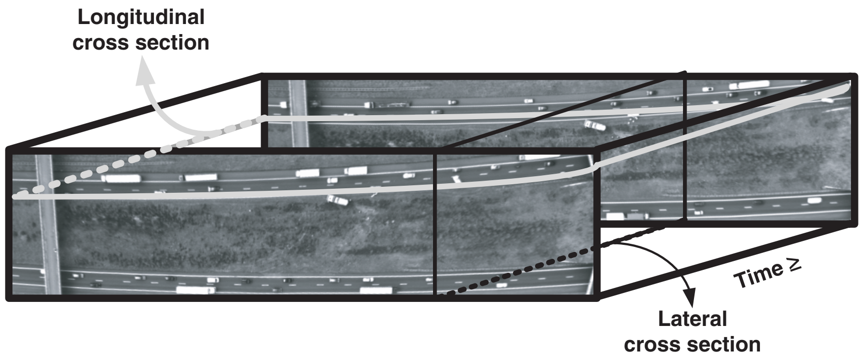

FIGURE 1 Three-dimensional box of pictures with lateral (red) and longitudinal (green) cross section.

The time series of pictures can be combined to a three-dimensional box, with two space dimensions and a time dimension (Figure 1). It is proposed that a cross section (i.e., a slice or cross cut of this box) be taken in one of the space dimensions. The first subsection shows the transformation in the lateral direction, perpendicular to the driving direction, indicated by a red frame in Figure 1. For one point along the road, the image shows the occupancies over time and the flows. The second subsection shows the transformations in the longitudinal direction, along the driving direction, indicated by a green frame in Figure 1. This transformation gives the occupancy over space and time for one lane, directly indicating the vehicle trajectories.

The video used here as an example and to quote computation times originates from an experiment described elsewhere (*19*). The pictures, with a size of 1,392 × 1,040 pixels, were recorded at a frequency of 15.1 Hz.

Lateral Cross Section

Figure 2 shows a lateral cross section of a series of pictures; it is a sectional plane perpendicular to the driving direction, like the lateral plane in Figure 1. The size of the image is a distance, 50 m from bottom to top, times a time, 75 s from left to right. Solid lines indicate a time step of half a minute, and dotted lines indicate a time step of 15 s.

This image can be interpreted as the occupancy of a virtual loop detector over time. It essentially shows when, on the horizontal axis, a vehicle passes a point along the road. Data from detection loops can lead to a similar image. However, this image also has a space component. Therefore, it shows where on the lateral *y*-position a vehicle passes the point of the cross section. This image can be constructed in a few minutes.

The image shown here is easily confused with a normal picture of traffic, because the structure of the road appears to be well visible. However, the length of the vehicles at the horizontal axis varies strongly, as in this image the length of a vehicle is the time it would occupy the virtual detector and therefore faster-moving cars are shorter.

In this image, cars in both directions have the nose on the left and the back on the right. Each vehicle first reaches the cross-section location with the nose of the vehicle. Because left in the image is an early point in time, the nose of the vehicle is drawn on the left-hand side, which is most easily recognized in articulated lorries with the truck on the left and the lorry on the right.

The transformation proposed in this subsection is ideal for obtaining time headways and flow values from remote sensing data. The time headway can be read out as the distance between two passing cars as indicated in the figure. To obtain the flow, it is necessary to count the number of vehicles that pass in a time interval, either manually or automatically, which is preferred for a longer data set.

This type of image is very informative for the lateral position on the road. One can directly observe whether cars in a platoon drive at the same lateral position or whether that depends on their position in the platoon or the headway or the type of car they are following. One can also determine whether the headway depends on the type of car (passenger car, truck) drivers are following.

There is one pitfall. Vehicles are not detected if they are upstream of the virtual detector in one picture and downstream in the next.

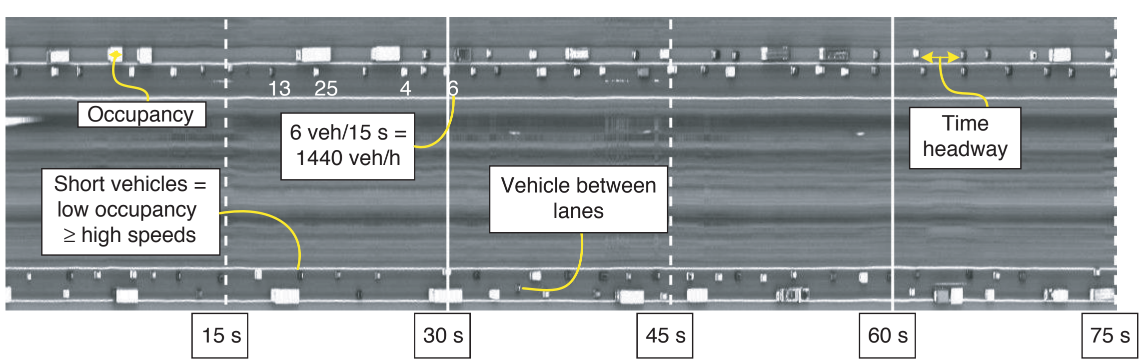

FIGURE 2 Lateral cross section, 1.25 min × 50 m.

However, when recording pictures at a time interval of (in this case) 15 images per second or more, that situation is very unlikely. In the worst case, a small [4-m (13-ft)] vehicle does not quite reach the cross-section point when the first image is taken. When the next picture is taken, it has passed the point completely. So it should have covered 4 m in 1/15th of a second and thus should drive 60 m/s, which equals more than 200 km/h (120 mph). For a frame rate of 10 images per second, the critical speed is 145 km/h (90 mph). For longer cars, these speeds are proportionally higher.

From this image, one cannot derive the dynamics of the car-following behavior. It can be derived from a longitudinal cross section discussed in the next subsection.

Longitudinal Cross Section

Figure 3 shows a longitudinal cross-section image along the axis of the lane of a road, as indicated with a green frame in Figure 1. On the horizontal axis, time progresses from left to right. Vertical lines indicate 15 s (dotted lines) and 30 s. On the vertical axis, one finds the distance along the road with the traffic flowing from bottom to top. The cross section is not necessarily taken along a straight line. A curved road requires one to take a section along a curved line, as shown in Figure 1.

Cho and Rice (*18*) proposed a method of filtering one lane out of a set of pictures. However, they take out one lane, which is several

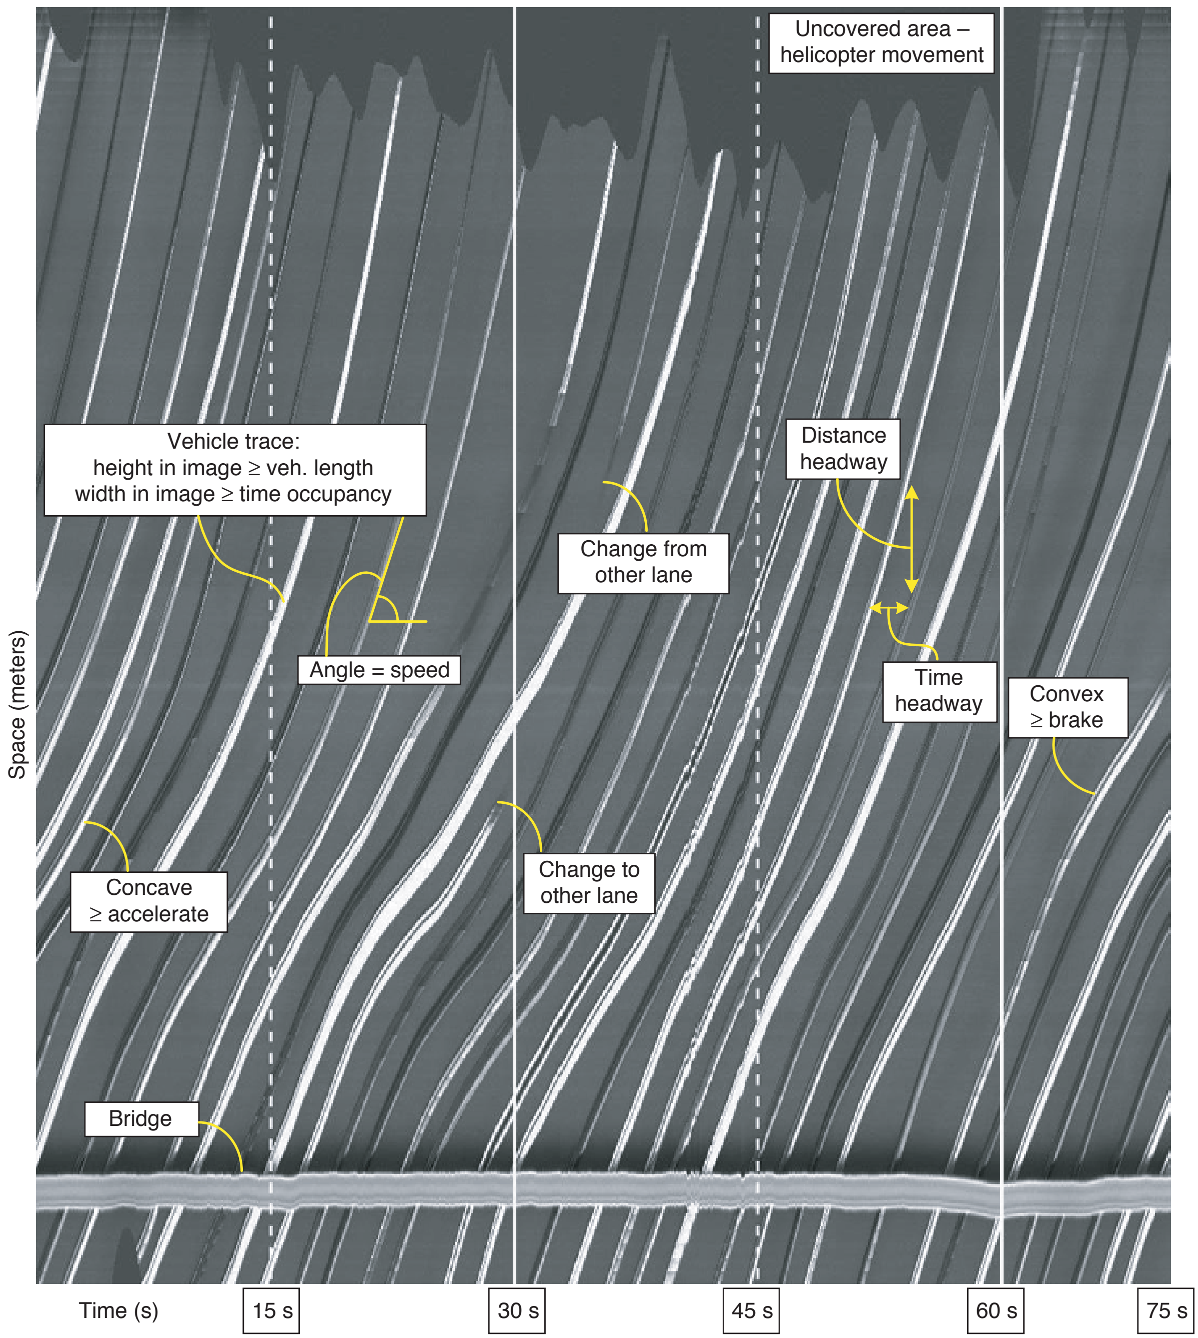

FIGURE 3 Longitudinal cross section.

tens of pixels wide instead of one pixel row as proposed. Then the maximum of the pixels in that lane across the road is taken, along a lateral cross section of the lane, which gives a clear image of the bright cars, but the dark cars disappear because the maximum intensity of the pavement is higher than that of a dark car. It is proposed that the data be directly narrowed to a one-pixel-wide cross section, which is—to the best of the authors' knowledge—the first time such an approach has been taken in the field of traffic flow.

Figure 3 is constructed from video data in a few minutes. It shows the intensity values of the (grayscale) images. The background of the image is the gray shade of the pavement. Vehicles move from the bottom (downstream) part of the image to the upper part. A bright line indicates that a bright vehicle is moving, a dark line shows the passing of a dark vehicle. Figure 3 is a graphic representation of the trajectory that is indicated by the word "trace" in this paper. The length of the vehicle determines the width of the line in this image; the occupancy time at one location determines the height of the line.

At the bottom part of Figure 3, a bridge that is present during the whole observation, blocks the view on the road. Therefore, it forms a line from left to right. Because the bridge is higher than the road surface, the part of the road that is made invisible by the bridge changes slightly when the helicopter moves and therefore the line is not straight.

The angle that the trace of a vehicle makes with the horizontal line is its speed. In this way, the image shows the speeds of all cars at all times. Also the changes in speed are visible, as concave traces equal a reduction in speed and therefore braking and convex traces equal an increase in speed and therefore accelerating.

Lane changes are represented by the appearance or disappearance of the trace of a vehicle in Figure 3. The following vehicle accelerates and closes the gap to the predecessor. Another vehicle then merges into the gap that is created behind the accelerating vehicle. The lane-changing vehicle enters this stream and thus appears in the image.

With this longitudinal cross-section image, the dynamics of traffic operations can be made visible without advanced detection and tracking algorithms. The image is typically constructed in several minutes with a personal computer equipped with a 1.5-GHz core processor. The constructing time is determined by the speed of reading the images because no advanced image processing algorithms are required. In this way, the dynamics of vehicle movements are made visible in a few minutes without fitting a model.

Contrary to the lateral transformation, this longitudinal transformation provides no information about the lateral position of vehicles other than the lane in which they are driving. In most cases, however, the only lateral information that is needed is the lane in which a vehicle is driving. Apart from that, each lane is shown in a different image. Nevertheless, one can track vehicles that change lanes by combining the tracks of the vehicles in different lanes.

VEHICLE DETECTION

This section describes how vehicles can be detected and tracked in the cross-sectional images. The longitudinal cross-section image is used for an example, because the trajectory information is the information that is more often required in studies with remote sensing data. The method of finding the cars (described in subsection on threshold values) is applicable for the lateral cross-section image as well. However, smoothing is needed only for the trajectories found in the longitudinal cross-section images. With a longitudinal cross-section image, a single vehicle could be tracked for a longer time than if a procedure were used that detects and tracks vehicles in each picture separately. Both the automatic part and the manual corrections are much quicker than the methods proposed so far. An image of a data set 5 min long can be processed (detecting and tracking) in 40 to 50 min. The manual corrections cost around 10 min.

The images provide a lot of information on traffic flows, which can easily be extracted manually. The sequel of this section describes how it can be done automatically. However, fitting a polynomial function through the lines of a longitudinal cross-section image gives the trajectories very accurately. In fact, the extent to which the polynomial curve follows the trace can be influenced by the number of points one uses to fit the trace. Generally, the number of points is higher for vehicles with more fluctuation in speed. Therefore, choosing an automatic detection or a manual identification is a trade-off between programming time and optimization of the algorithm on the one hand and working time (clicking points in the image) on the other hand.

Threshold Values

The images need to be changed from a gray value to a Boolean value (true or false) indicating vehicle presence. Compared with existing approaches (*14, 15*), two changes are made in the process of detecting vehicles. First, one does not take the values of one pixel over the whole time as a basis for the threshold values because it presents problems with the change in light intensity due to passing clouds. It works much better if a shorter time (10 s) is considered. However, this procedure also has a drawback, as, for example, a truck could be at the same spot for more than half the time. In that case, with a median for the background, the intensity value of the truck would actually be considered to be the background and therefore the pavement would be seen as a vehicle. To avoid this problem, the pixels were clustered in blocks of about 50 m (about 150 pixels), which gave many more pixels to extract the distribution of intensities. This process gives "blocks" of about 150 pixels (10 s) × 150 pixels (50 m) to find the background intensity. The second difference with existing methods (*14*) is that the background was not determined and a deviation of it was found. In contrast, upper and lower percentile values were used, which gave more reliable results.

For each block mentioned in the previous paragraph, a lower and upper value were computed, as shown in Figure 4, where the intensity along one line is plotted. The background intensity profile varies along

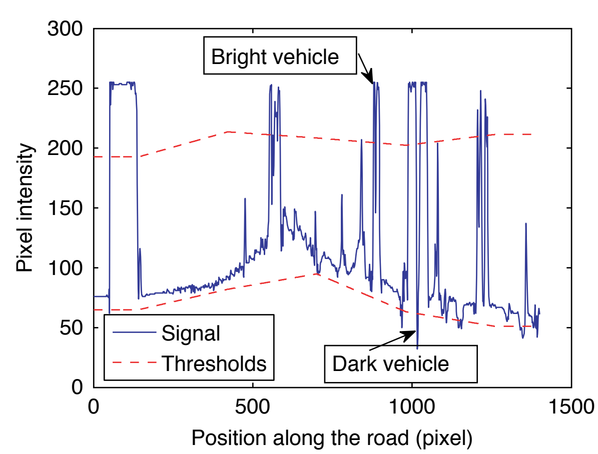

FIGURE 4 Intensity of image over space.

the road, which could be caused by reflections of the road at various angles. The threshold values for each of the blocks are interpolated linearly to obtain the threshold values for each pixel individually. This pattern is also plotted in Figure 4.

The percentage of pixels marked as a car is fixed by choosing the percentile values. Moreover, it also determines whether the dark or the light deviations will be marked as a vehicle. Therefore, the choice for the best percentile values depends, among others, on traffic conditions and the brightness of the pavement relative to the brightness of the vehicles. For example, if there is much traffic, the lower percentile value should be chosen low and the higher percentile value should be chosen high to ensure that a large part of the pixels are marked as vehicles. If the pavement is light, most vehicles will be darker than the pavement and therefore both the lower and the higher percentile values can be increased. In practice, it is best to try different settings of percentile values. In this case, it was best to choose 15% for the lower percentile value and 90% for the higher percentile value.

Special attention should be paid to shadows, as they might interfere with vehicle detection. For instance, if shadows lie over the adjacent lane, it is advisable to choose a pixel line of the video more to the side of the lane where no the shadows lie. This procedure was possible for the test data set. If it is unavoidable to take a pixel line with shadows, the percentile values have to be changed. In particular, the lower percentile value has to be increased to avoid detecting the shadows as cars. Similarly, shadows could be detected as part of a car if they lie along the road, which could increase the length of the area being detected as different from the background and therefore increase vehicle length and thus intervehicle spacing. To avoid this problem, the lower threshold value should be adjusted.

Another problem arises when there are vehicles with a color that gives a pixel intensity similar to or equal to the pavement. Vehicles with the exact same reflection as the pavement cannot be detected. However, in the test case there were no such vehicles, as could be derived from the shadows. Distinguishing vehicles from the pavement is easier if color video is used.

Vehicle presence now has to be filtered. First, the image with traces (in the space–time plane of the longitudinal cross section) is filtered by applying an exponential moving average filter with an exponent of 2 pixels (2/15 of a second) in the time direction and 1 pixel (30 cm) in the space dimension to filter out the small areas that deviate from the neighboring pixels. The result is a smoother image. Then the image is filtered based on the properties of cars. Any vehicle has to have a length of 3 to 20 m (10 to 66 ft). If a vehicle is too small, the signal is ignored because it is noise. Additionally, a gap between two vehicles has a minimum size of 3 m (10 ft). If a gap is too small, the two parts are combined because they are either noise or, for example, an articulated lorry.

Now, the vehicles found in each time step should be matched to create trajectories. In two successive lines, the expected displacement of the vehicle can be calculated based on speed; if speed is not known—for example, in the first time step—a fixed estimate can be used. If the match of size and location is better than a threshold value, the detected vehicles are matched to make a trajectory.

This process will give the trajectory in distinct parts. In some time instants, vehicles are not detected properly, which means there are many "partial" trajectories that have to be coupled in the next step. Of trajectories that end at a location other than the edge of the image, most have to be coupled to another partial trajectory; the vehicle disappears from the image only if it changes lanes. The coupling of this partial trajectory works as follows. Each partial trajectory is extrapolated based on the position and speed slightly before it ends.

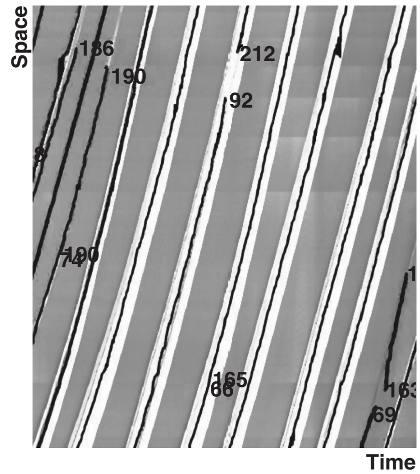

FIGURE 5 Trajectories on top of image of longitudinal cross section.

In this process, it is better not to use the very last position detected because the last measurement gets noisy before it is not detected anymore. This extrapolated trajectory is compared with every starting partial trajectory. If one of them matches in space and speed, both partial trajectories are coupled. The maximum allowed error in this matching process can be chosen. If one chooses a high value for this maximum, more partial trajectories are coupled, but that increases the probability of coupling partial trajectories that do not belong to the same vehicle; a lower threshold leaves some partial trajectories uncoupled, but there are fewer errors in the coupling. In this paper a relatively strict option was chosen.

This choice means that several partial trajectories remain uncoupled; on average, each vehicle trajectory now consists of two partial trajectories. They are now corrected manually. To this end, the partial trajectories and a unique number identifying the trajectory are plotted (Figure 5). An Excel sheet is used to indicate which partial trajectories should be combined. This process takes around 10 min for a video 5 min long. From this point, further processing is automatic.

Smoothing

The procedure described in the preceding subsections gives the trajectory for each vehicle. Nevertheless, these trajectories are not smooth, as Figure 5 shows. This situation is caused mainly by the spatial and temporal resolution of the images but also by detecting and tracking errors. To represent reality, they need to be smoothed without smoothing out the dynamics of traffic. To that end, the method of Toledo et al. is applied (*20*). It averages the position at a time step with the position at neighboring time steps, with a larger weight for time steps that are closer. The smoothing filter takes time frames into account with times that differ at a maximum of 1.5 s from the current time frame. An example of the end result—the trajectories—is presented in Figure 6.

CONCLUSIONS

Video data from traffic operations contain much valuable information. However, extracting, for example, trajectory data by existing methods is a difficult and often erroneous process. The approach proposed here

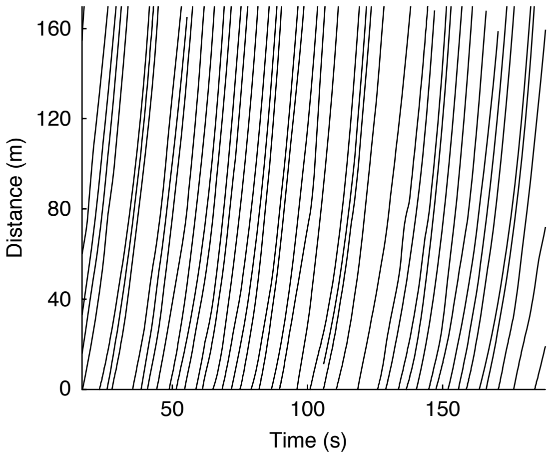

FIGURE 6 Trajectories for several vehicles.

increases the length of the parts of the trajectories, overcoming the shortcomings of existing methods. To that end, instead of analyzing each picture in a video sequence separately, the data of the video are transformed in such a way that the vehicle trajectories (vehicle positions over time) become visible within a single image.

For example, the cross section of the three-dimensional box along the axis of the road provides quick and easy insight into the dynamic car-following behavior. Within a few minutes, traces of the cars (i.e., the vehicle trajectories) are plotted in a space–time diagram. From the constructed image, many dynamic properties of the traffic are apparent—such as shockwaves, reaction times, and driver over-reaction. With image detection techniques, the vehicle positions can be quantified—a very reliable means of constructing trajectories.

In other methods, the vehicle trajectories are often divided into many different segments. In the method presented here, the segments can be longer and more easily coupled. Close to 95% of the vehicle presence (in the space–time frame) is detected, while the remaining 5% can be manually reconstructed very quickly. The complete process consisting of transforming, detecting, tracking, and manually correcting a single lane can be performed in about 1 h for a 5-min data set with a personal computer equipped with a 1.5-GHz core processor.

Another cross section of the three-dimensional box, perpendicular to the driving direction, shows other traffic properties. Headways and occupancies can be derived directly from this single image. For these purposes, it is not necessary to detect all vehicles in all pictures; only the vehicles currently passing the location of the virtual detector need to be detected.

In these analyses, black and white images were used. With color images, the contrast between vehicles and pavement increases, making it easier to detect the vehicles. The tracking process also benefits from color as vehicles in consecutive pictures can be identified and linked by color. Thus, in the future, the use of color images is expected to increase further the length of the automatically discerned vehicle trajectories.

ACKNOWLEDGMENTS

This research was supported by the research program Next Generation Infrastructures. The authors thank the reviewers for their helpful comments, which improved the paper.

REFERENCES

1. Greenshields, B. D., J. R. Bibbins, W. S. Channing, and H. H. Miller. A Study of Traffic Capacity. *Highway Research Board Proceedings,* Vol. 14, HRB, National Research Council, Washington, D.C., 1935, pp. 448–477.
2. Chandler, R. E., R. Herman, and E. W. Montroll. Traffic Dynamics: Studies in Car Following. *Operations Research,* Vol. 6, 1958, pp. 165–184.
3. Treiterer, J., and J. A. Myers. The Hysteresis Phenomenon in Traffic Flow. In *Proc., 6th International Symposium on Transportation and Traffic Theory* (D. J. Buckley, ed.), Elsevier, Amsterdam, Netherlands, 1974, pp. 13–38.
4. Ossen, S. J. L. *Longitudinal Driving Behavior: Theory and Empirics.* Delft University of Technology, Delft, Netherlands, 2008.
5. Hoogendoorn, S. P. Unified Approach to Estimating Free Speed Distributions. *Transportation Research B,* Vol. 39, No. 8, 2005, pp. 709–727.
6. Brockfeld, E., R. D. Kühne, and P. Wagner. Calibration and Validation of Microscopic Traffic Flow Models. In *Transportation Research Record: Journal of the Transportation Research Board, No. 1876,* Transportation Research Board of the National Academies, Washington, D.C., 2004, pp. 62–70.
7. Brackstone, M., B. Sultan, and M. McDonald. Motorway Driver Behaviour: Studies on Car Following. *Transportation Research F,* Vol. 5, No. 1, 2002, pp. 31–46.
8. *NGSIM Overview. Next Generation SIMulation Fact Sheet.* FHWA-HRT-06-135. FHWA, 2006.
9. Kyte, M., M. P. Dixon, A. Abdel-Rahim, V. Nayak, and T. Urbanik. What's New in the Queue: Discovering More about Queue Discharge Characteristics and Their Effect on Signal Timing Using the New NGSIM Data Set. Presented at 87th Annual Meeting of the Transportation Research Board, Washington, D.C., 2008.
10. Choudhury, C., M. E. Ben-Akiva, T. Toledo, A. Rao, and G. Lee. *Cooperative Lane Changing and Forced Merging Model: NGSIM Final Report.* FHWA, U.S. Department of Transportation, 2006.
11. Hamdar, S. H., and H. S. Mahmassani. From Existing Accident-Free Car-Following Models to Colliding Vehicles: Exploration and Assessment. In *Transportation Research Record: Journal of the Transportation Research Board, No. 2088,* Transportation Research Board of the National Academies, Washington, D.C., 2008, pp. 45–56.
12. Chiabaut, N., C. Buisson, and L. Leclercq. Fundamental Diagram Estimation Through Passing-Rate Measurements in Congestion. Presented at 87th Annual Meeting of the Transportation Research Board, Washington, D.C., 2008.
13. Kesting, A., and M. Treiber. Calibrating Car-Following Models by Trajectory Data: A Methodological Study. In *Transportation Research Record: Journal of the Transportation Research Board, No. 2088,* Transportation Research Board of the National Academies, Washington, D.C., 2008, pp. 148–156.
14. Hoogendoorn, S. J., H. J. van Zuylen, M. Schreuder, B. Gorte, and G. Vosselman. Microscopic Traffic Data Collection by Remote Sensing. In *Transportation Research Record: Journal of the Transportation Research Board, No. 1855,* Transportation Research Board of the National Academies, Washington, D.C., 2003, pp. 121–128.
15. Angel, A., M. Hickman, and P. Mirchandani. Methods of Analyzing Traffic Imagery Collected from Aerial Platforms. *IEEE Transactions on Intelligent Transportation Systems,* Vol. 4, No. 2, 2003, pp. 99–107.
16. www.ngsim.fhwa.dot.gov/. Accessed July 29, 2008.
17. Zhang, L., V. Kovvali, N. Clark, D. Sallman, and V. Alexiadis. *NGSIM-VIDEO User's Manual.* FHWA-HOP-07-009. FHWA, U.S. Department of Transportation, 2007.
18. 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.
19. Knoop, V., P. Hoogendoorn, and J. van Zuylen. Capacity Reduction at Incidents: Empirical Data Collected from a Helicopter. Presented at 87th Annual Meeting of the Transportation Research Board, Washington, D.C., 2008.
20. Toledo, T., H. N. Koutsopoulos, and K. I. Ahmed. Estimation of Vehicle Trajectories with Locally Weighted Regression. Presented at 86th Annual Meeting of the Transportation Research Board, Washington, D.C., 2007.

The Intelligent Transportation Systems Committee sponsored publication of this paper.