PERFORMANCE EVALUATION AND ERROR SEGREGATION OF VIDEO-COLLECTED TRAFFIC SPEED DATA

Paul Anderson-Trocmé, Graduate Research Assistant
Department of Civil Engineering and Applied Mechanics, McGill University
Room 391, Macdonald Engineering Building, 817 Sherbrooke Street West
Montréal, Québec, Canada H3A 0C3
Email: paul.anderson-trocme@mail.mcgill.ca

Joshua Stipancic, Corresponding Author, Graduate Research Assistant
Department of Civil Engineering and Applied Mechanics, McGill University
Room 391, Macdonald Engineering Building, 817 Sherbrooke Street West
Montréal, Québec, Canada H3A 0C3
Email: joshua.stipancic@mail.mcgill.ca

Luis Miranda-Moreno, Associate Professor
Department of Civil Engineering and Applied Mechanics, McGill University
Room 268, Macdonald Engineering Building, 817 Sherbrooke Street West
Montréal, Québec, Canada H3A 0C3
Phone: (514) 398-6589
Fax: (514) 398-7361
Email: luis.miranda-moreno@mcgill.ca

Nicolas Saunier, Associate Professor
Department of Civil, Geological and Mining Engineering
Polytechnique Montréal, C.P. 6079, succ. Centre-Ville
Montréal, Québec, Canada H3C 3A7
Phone: (514) 340-4711 x. 4962
Email: nicolas.saunier@polymtl.ca

Word count: 5614 words + 7 tables/figures x 250 words (each) = 7364 words

November 15th, 2014

ABSTRACT
Validating the accuracy of sensors is an essential step in the collection of traffic speed data. The accuracy of automated speed data has been evaluated in small- and large-scale tests using multiple technologies and methods. While inductive loops are standard, video-based detectors have demonstrated the ability to substitute conventional detection devices. Though existing literature documents several issues associated with extracting vehicle speeds from video, the analysis of speed data, especially at the microscopic or individual level, has been limited. The purpose of this paper is to evaluate the accuracy of a video-based detection system, comprised of commercially available video cameras and an open-source computer vision software system. Several camera orientations were tested along an urban arterial and a highway in Montreal, Canada. A semi-automated vehicle tracking process was used to extract the vehicle speeds, which were compared to manually observed speeds. Although the traditional mean relative error approach led to unacceptable results, a new approach was proposed for the evaluation of traffic detection technologies. The segregated error approach divides simplistic mean error into separate values for accuracy and precision. In doing so, several of the camera orientations exhibited precision error values within the accepted range for speed data quality (5%). Even with large errors, the potential exists to calibrate video-based speeds, by removing the over- or underestimation bias, to acceptable performance levels as long as precision error is minimized through appropriate selection of camera position and orientation.

Keywords: traffic detector, sensor, video data, object tracking, computer vision, performance, evaluation, accuracy, precision

INTRODUCTION

The collection and analysis of vehicular speed data are essential for urban transportation systems. Sensors that collect accurate and consistent data are necessary to guide engineering decisions and treatments towards desired impacts in planning, construction, or operations (*1*). The greatest challenge in any data collection campaign or large-scale test is the creation of "ground truth" data, or a "reference data set that represents the actual history of the traffic" (*2*). The need for accurate data is critical as errors in this early stage will compound through analysis, skewing study outcomes and misleading decision making (*3*). Data quality is paramount, and sensors must be sufficiently accurate for the specific data needs of a given project (*3*).

Traditional automated vehicular data collection was limited to the use of inductive loops at fixed locations (*1*), to the point that loops became the "de facto standard" in many jurisdictions and are still widely used today (*4*). Despite the performance of these systems, it is impractical and costly to maintain an adequate network of permanent collection locations in an urban road network (*5*). Accordingly, the use of non-intrusive traffic data collection technologies has become increasingly popular. Non-intrusive devices do not require access to the travel lane for installation, are often installed outside the right of way, and are safer to install and operate compared to other technologies (*1*). Video-based traffic sensors are among the most promising non-intrusive technologies. Simple video cameras have the ability to substitute conventional detection devices (*6*), provide flexibility in mounting locations, enable multiple lane detection, and provide rich positional data beyond counts and speed (*1*).

As manually processing video is resource demanding, "there is a high demand for automation of this task" (*7*). Numerous systems have been developed for the automated extraction of traffic data from video footage using computer vision techniques (*8; 9*). These systems are able to provide a wide variety of data, from conventional traffic parameters such as flow and velocity (*1; 10*), to new parameters such as trajectories which provide information on manoeuvring and traffic conflicts (*11*). Unlike other detectors, the raw sensor signal (video) contains rich information (the entire series of events as they occurred during data collection) that can be extracted and verified manually, with the potential to be extracted automatically as technology evolves.

While video data has many advantages, before any system is relied on for traffic data collection, the accuracy of the system must be verified to ensure data quality is maintained. With respect to existing literature, this research provides several key contributions. Most attempts to verify data quality have quantified error based on aggregate data, with little research focused on reasonable accuracy for individual vehicle data. Moreover, existing literature provides little guidance on acceptable accuracy for microscopic speed data. Methods for evaluating error have considered simplistic mean error without consideration for more robust analyses. The purpose of this study is to evaluate the accuracy of a video-based detection system, comprised of commercially available consumer video cameras and the open source Traffic Intelligence video analysis software system (*11*). The objectives of this research are to; evaluate the accuracy of automated vehicular speed extraction, recognizing detection as a necessary precursor; and, to propose a technique for evaluating separately the precision and accuracy of collected traffic data.

LITERATURE REVIEW

Vehicle tracking through computer vision is a powerful tool that has seen implementation in several areas of transportation and safety research. Vehicle tracking techniques provide information in the form of vehicle trajectories (*10*), or the sequence of positions indexed by time of an object of interest, such as a vehicle or pedestrian, from which the velocity vector can be derived. The nature of video detection allows vehicles to be tracked continuously over a road segment rather than being detected at a single location. The analysis of vehicle trajectories can be used to automate otherwise resource-intensive studies, including vehicle manoeuvring (*10*), lane-changing, queuing patterns (*6*), automated incident detection (*4*), driver behaviour (*7*), and conflict analysis (*11*). This increased sophistication does not prevent video-based sensors from substituting or complementing traditional traffic detectors. Along with trajectories, video has the ability to capture traditional traffic parameters including flow, speed, headway, and density at both the individual vehicle (*10*) and corridor level (*6*). This vehicle data has applications in the calibration of traffic flow models (*10*) and surrogate safety analysis. As detecting and extracting the position and speed of vehicles provides the basis for surrogate safety analysis, verifying the parameters themselves will lend additional credibility to the technique.

To extract traffic data, vehicles must first be successfully detected. In a study of video-based vehicle classification, Gupte et al. (*12*) achieved a detection rate of 90%, while the system utilized by Messolodi et al. (*13*) exhibited an average count error of -5.2% operating in real time. Analysis of video-extracted speed data has been limited, especially at the individual vehicle level, where the amount of video that must be processed is high and analysis is resource intensive (*6*). Additionally, methods used to assess speed data quality have focussed on simplistic approaches for quantifying error. Coifman (*10*) extracted data for velocity, flow, and density collected in real-time on a freeway facility. Data was temporally aggregated (averaged) to 5-minute intervals, resulting in 514 samples. When compared to ground truth from calibrated inductive loops, 100% of the samples exhibited speed error less than 10%, while 95% exhibited error less than 5%. Malik et al. (*14*) showed improvements using post-processing. Across the study, detection rates varied between 75% and 95%. Using 5-minute aggregation periods, nearly all samples showed speed error of 5% or less, though error varied with lane position relative to the camera. MacCarley et al. (*2*) did not strictly consider the accuracy of each observation, but claimed that 95% of all extracted speeds were "reasonable" when compared to speeds determined manually.

Dailey (*15*) utilized individual vehicle data, but considered only 190 vehicles from 40 seconds of video. While able to demonstrate that the mean error of speed across all observations was zero, the individual relative error for each observation varied between -40% and 80%. Though estimating mean traffic speed was possible, the technique is impractical for microscopic data. Schoepflin (*16*) compared speed distributions created with manual data and video data, noting that they were approximately equal in mean and distribution by visual comparison. The authors claimed this indicated "certain equivalence" between the video-based speeds and actual events. Although errors of up to 20% were possible, averaging individual speeds over 20-second intervals reduced variability by a factor of 10. These studies exemplify the issue with data aggregation. Using temporally aggregated data for speed analysis effectively eliminates the influence of the highest and lowest recorded speeds and obscuring the performance of a device exhibiting compensating errors (cancellation of high and low outliers) (*5*), behaviour clearly present in existing video-based systems.

Several notable issues exist with regard to extracting vehicle speeds from video, including false or missed detections. False detections involve the detection of any object that does not exist. Shadows cast by vehicles in adjacent lanes are particularly problematic (*2*). Vehicles can be missed if they are partially or fully occluded by other vehicles. Occlusions can also disrupt trajectories, creating inaccurate trajectories and speed estimations (*11*). Vehicle position relative to the camera may affect accuracy. Vehicles that are further from the camera occupy a smaller number of pixels, and may be difficult to identify and track, leading to potential variability in speed data (*6*). Vehicle tracking may be inhibited by an overestimation of derivative values, in which the "distance between two measured points is systematically biased towards longer distances, which results in speed overestimation" (*7*).

If video detection is to be considered a reasonable alternative to other devices, data must meet similar quality and reliability requirements. Bahler (*1*) indicated that inductive loops exhibiting count errors less than 4% over 1-hour aggregation periods was of sufficient quality (*5*). The same study demonstrated that most commercially available non-intrusive traffic detectors, including video, were able to provide counts within 3% of actual, and speeds within 8% (*1*). However, most attempts in existing literature to verify data quality have quantified error based on aggregated data, with little research focused acceptable accuracy for individual vehicle data. Additionally, there has been no consideration for evaluating device precision (repeatability of speed measurement) and accuracy (general over- or underestimation bias) separately. In general, researchers should endeavour to find detectors that "approach the ideal, but fall within some level of tolerance" given the specific application (*3*).

METHODOLOGY

Site Selection, Instrumentation, and Data Collection

The quality of video-extracted traffic speeds was evaluated in a highway and arterial environment to incorporate variation in geometry and traffic parameters such as speed and volume. Autoroute 15 (A15) in Montreal, Quebec, Canada is a major north-south corridor with an AADT of approximately 90,000 (*17*). Data was collected in September 2013 at midday in clear conditions. At the test location near Boulevard Henri Bourassa, four lanes are present in each direction, and the posted speed limit is 100 km/h. Boulevard Taschereau in Brossard, Quebec, Canada runs east-west on the South Shore of Montreal, connecting two major bridges and servicing important highway connections to the island of Montreal. The section chosen for this study, near Boulevard Lapinere, featured five lanes in each direction with a posted speed limit of 50 km/h. Data was collected in October 2013 during the morning peak in clear and overcast conditions.

A GoPro Hero 3 camera was used to collect video at both sites, and was set to record 720p video at 30 frames per second. The camera is capable of recording up to six hours of video on a single battery charge, is highly portable, and provides flexibility in mounting location. Sites were chosen with existing roadside infrastructure for camera installation. At the A15, a pedestrian overpass structure was utilized, allowing the camera to be mounted closer to the roadway compared to other potential locations. The video camera was attached to the guardrail on the overpass, shown in Figure 1a. At Taschereau, the camera was mounted to a 20-foot telescoping fibreglass surrogate pole, which was subsequently fixed to the base of existing luminaire poles, shown in Figure 1b. The close proximity of adjacent poles allowed for the collection of simultaneous video data from multiple orientations.

(a) (b)

FIGURE 1 Mounting configurations for freeway (a) and arterial (b) environments

In addition to multiple sites, multiple camera orientations were utilized to analyze the effect of orientation on reported accuracy. Knowledge of accuracy with respect to the distance between the object and camera is vital because the ability to utilize multiple orientations improves flexibility and provides more mounting options, which may benefit data collection in urban environments. Three camera orientations were used at each site. For the first orientation (Perpendicular) the camera was positioned over the roadway orthogonal to the direction of travel. This orientation may provide the most accurate speeds, as vehicles are relatively close to the camera and the effects of perspective are minimized. Two parallel orientations, with the camera positioned parallel to the direction of travel, were also tested. A parallel orientation is beneficial if information, such as vehicle trajectory, is required over a longer segment of the road. This orientation was used with one speed extraction zone approximately 10 m from the camera (Parallel Close) and one approximately 20 m from the camera (Parallel Far). At least 30 consecutive minutes of video were recorded for every orientation at each site. The locations of the cameras and their data extraction zones are provided in Figure 2. These orientations allowed for speed data to be collected at three different camera-to-object distances.

(a) (b)

FIGURE 2 Perpendicular (A), Parallel Close (B), and Parallel Far (C) study areas (white and black dots denote the camera positions corresponding to the color of the study area letters) for (a) arterial and (b) highway locations (aerial images from *Google Maps*. https://www.google.ca/maps/)

Feature-Based Tracking Algorithm

Data extraction was automated using an open-source computer vision software system, Traffic Intelligence (*18*), developed at Polytechnique Montreal, Canada. The program enables users to analyze video, extract vehicle trajectories, and evaluate trajectory data using several libraries and tools. The primary tool is a feature-based tracking algorithm that outputs trajectories for all moving objects in the video frame, which are mapped to real-world measurements using a homography matrix to convert object positions from the image (pixels) to the road surface (meters). The extraction and grouping of trajectories into corresponding vehicles is a crucial step. First, moving points, or features, are identified and tracked between consecutive frames. Features are grouped into objects based on several criteria, and are stored in a database with their two-dimensional coordinates and instantaneous velocity values for each video frame.

The main issues with feature-based tracking are over-segmentation and over-grouping of trajectories. Over-segmentation occurs when a single object is assigned multiple trajectories. Over-segmentation can lead to inflated vehicle counts and may skew speed distribution, but does not affect speed accuracy since grouped features belong to a single vehicle. Over-grouping occurs when multiple vehicles are represented by a single trajectory, due to the proximity of neighbouring features. The over-grouping of objects will lead to inaccurate speed calculations within a range (if several objects are grouped, they must have similar speed by construction) and false volume counts (*19*). Given these issues, the most important parameters within the tracking software are related to the criteria used for grouping features into objects. In order to accurately apply the software, the key parameters were calibrated by ensuring extracted counts matched manual counts within 2% over a sample of the collected video.

Data Extraction

The data sets were compiled following a semi-automated approach. The extraction of speeds was completely automated through the computer-vision software. Virtual speed boxes were added to the video frame, where extracted trajectories were evaluated and instantaneous object speeds were averaged to obtain the mean speed over the box length of 10 to 12 meters, measured using standardized pavement marking lengths. Speed boxes were created for two lanes (Lane 2 and Lane 3, with Lane 2 being the rightmost lane), for each orientation, at both sites, yielding a total of 12 study areas. The video output provides an object number and trajectory overlaid on the corresponding vehicle. Ground truth comparison speeds were determined manually for those vehicles that had been tracked. Using the length of the speed box and the travel time (in frames, converted using the video frame rate), the speed of the vehicles could be calculated. This method results in a discrete set of possible ground truth speeds, as only an integer value of frames is possible for manually computing travel time. This method results in a margin of error of approximately 2 % at 50 km/h and 4% at 100 km/h. Manual speeds were matched one-to-one with the extracted speeds using the corresponding object number. Vehicles that were over-segmented or over-grouped were removed from the data set.

Data Analysis

Analysis of the extracted speeds was completed in three steps. First, the mean relative error was calculated for every orientation at each site. The use of mean error is a simplistic quantification method for speed measurement and is consistent with analyses demonstrated within existing detector testing literature. A sample of 100 consecutive vehicles was selected in each of the 12 study areas, for which the mean relative error was calculated for the extracted speeds. 1200 total observations represent a fairly substantial effort for manual verification. The error was calculated for each individual record by normalizing the difference between the automatically extracted and manually observed speed. These individual errors were averaged across the sample to yield the mean relative error, according to

$$Mean\ Relative\ Error\ = \frac{1}{100}\sum\frac{|V_{extracted}-V_{observed}|}{V_{observed}} \qquad (1)$$

To better understand the behaviour of the detector and the characteristics of the errors, and to observe trends across lanes and orientations, the extracted speeds were plotted against the observed speeds in a data visualization exercise. The plots utilize a diagonal line to indicate ideal detector performance (that is, data from an ideal detector follows the line $y=x$). Data points above the line indicate overestimation of speed, while points below the line indicate underestimation. A fitted line with slope equal to 1 can be added to the data which allows for the observation of accuracy (the distance between the line of best fit and the line y=x) and precision (the residual errors between the data points and the line of best fit) as separate phenomena, as illustrated in Figure 3.

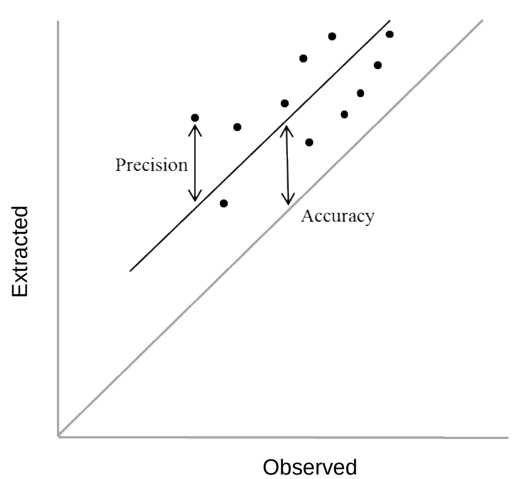

FIGURE 3 Illustration of error segregation

This study contends that mean error is insufficient at capturing the true behaviour of detectors and other measures are necessary to define device precision and accuracy separately. While data plots provide this information visually, it is beneficial to compute the relative error for precision and accuracy. Utilizing relative error values matches the approach utilized in existing literature, and provides an intuitive and communicable comparison between sites and camera orientations. The y-intercept of the fitted line quantifies the difference between the detector and ideal behaviours and indicates the magnitude of difference between the mean of the extracted speeds and the mean of the observed speeds. The precision error is quantified similarly to the mean error, with the subtraction of a correction factor equal to the y-intercept of the fitted line. To normalize the intercept value consistent with the relative mean error, the y-intercept is evaluated at every data point (divided by the harmonic mean of observed speed). Values for relative precision and accuracy error are calculated as

$$Relative\ Precision\ Error\ = \frac{1}{100}\sum\frac{|(V_{extracted}-y\ intercept)-V_{observed}|}{V_{observed}} \quad (2)$$

$$Relative\ Accuracy\ Error = \frac{1}{100}\sum\frac{|y\ intercept|}{V_{observed}} \quad (3)$$

RESULTS

Mean Error Approach

Mean relative errors are reported in Table 1. Extracted speeds exhibit important difference when compared to manually observed speeds at many of the study areas. At Taschereau, the mean error in five of the six cases exceeded 10%, and variation between the lanes is observed for a single camera orientation. For the A15, the mean error values are consistently lower with less variation between the lanes (between 3% and 12%). The results of the mean error approach indicated that video extracted speeds do not fall within acceptable limits for traffic detectors at Taschereau. For the A15, the video-based speed data is of acceptable quality (<5%) for exactly half of the study areas. In general, if performance were quantified using only the mean relative error, then the video data would fail to meet the performance requirements for traffic speed data and would be unacceptable for future studies.

TABLE 1 Mean Relative Error Values for Video-Extracted Speeds

	Mean Relative Error	
	Lane 2	**Lane 3**
Taschereau		
Perpendicular	0.16	0.08
Parallel close	0.22	0.12
Parallel far	0.15	0.15
A15		
Perpendicular	0.08	0.03
Parallel close	0.05	0.05
Parallel far	0.10	0.12

Data Visualization

To better understand device behaviour, the data visualization exercise was completed for all study areas, the results of which are presented in Figure 4 and Figure 5. The detected speeds were plotted against the observed speeds, and a fitted line was added with the slope held equal to one. Fixing the slope to be equal to 1 indicates an assumption that over- or underestimation is independent of operating speed; an assumption that holds as no contradictory patterns of error were observed in the data and the lines appear to provide a sufficient fit. Using this technique, the intercept of the trend line can indicate the direction and magnitude of the general estimation error (the bias), while the coefficient of determination, R^2, can indicate the precision of the extracted speeds (the "spread" around the trend line). Based on the intercept values alone, 11 of the 12 study areas exhibit overestimation, consistent with previous research (*7*).

While the intercept provides insight into the accuracy of the extracted speeds, the R^2 value for each trend line reveals the precision, or repeatability of each extracted speed measurement, with a higher R^2 indicating a higher degree of repeatability. In general, the R^2 values for the perpendicular camera orientation were the highest in all cases, ranging between 0.77 and 0.89. The parallel close orientation had the second highest repeatability, with R^2 values between 0.26 and 0.63. The worst results were for the parallel far orientation, with R^2 values near zero. R^2 was calculated as

$$R^2 = 1 - \frac{residual\ sum\ of\ squares}{total\ sum\ of\ squares} \tag{4}$$

Therefore, the negative R^2 value in Figure 4e indicates that the fitted line (with slope equal to 1) fits the data more poorly than the mean value of speed (residual sum of squares is greater than total sum of squares).

Segregated Error Approach

While the visualization exercise proved to be powerful by allowing observation of both precision and accuracy as separate phenomenon through the use of the R^2 and intercept values, it is beneficial to formally compute the relative errors using the presented formulae. The segregated errors for all sites and orientations are presented in Table 2. No patterns are observable from the mean error alone. There is no clear indication that one orientation will consistently provide the lowest mean error, and the video data fails to meet the performance standards required for collection of disaggregate traffic speeds. When the error is segregated, more interesting results are observed. Accuracy similarly appears to be independent of orientation and lane, but also of the initial mean error.

TABLE 2 Segregated Relative and Absolute (km/h) Error Values for Video-Extracted Speed

	Lane 2 Relative (Absolute) Error			**Lane 3 Relative (Absolute) Error**		
	Mean	**Accuracy**	**Precision**	**Mean**	**Accuracy**	**Precision**
Taschereau						
Perpendicular	0.16 (8.7)	0.17 (8.7)	0.04 (2.2)	0.08 (4.2)	0.07 (3.7)	0.04 (2.3)
Parallel close	0.22 (11.1)	0.20 (10.2)	0.10 (5.0)	0.12 (6.1)	0.05 (2.8)	0.10 (5.1)
Parallel far	0.15 (7.8)	0.11 (5.9)	0.11 (5.7)	0.15 (8.2)	0.05 (2.7)	0.13 (7.3)
A15						
Perpendicular	0.08 (7.7)	0.08 (7.2)	0.03 (3.3)	0.03 (3.3)	0.02 (2.1)	0.03 (2.7)
Parallel close	0.05 (4.5)	0.04 (3.6)	0.05 (4.1)	0.05 (4.7)	0.03 (3.0)	0.04 (3.6)
Parallel far	0.10 (9.5)	0.04 (3.8)	0.09 (8.6)	0.12 (10.6)	0.08 (7.1)	0.10 (8.9)

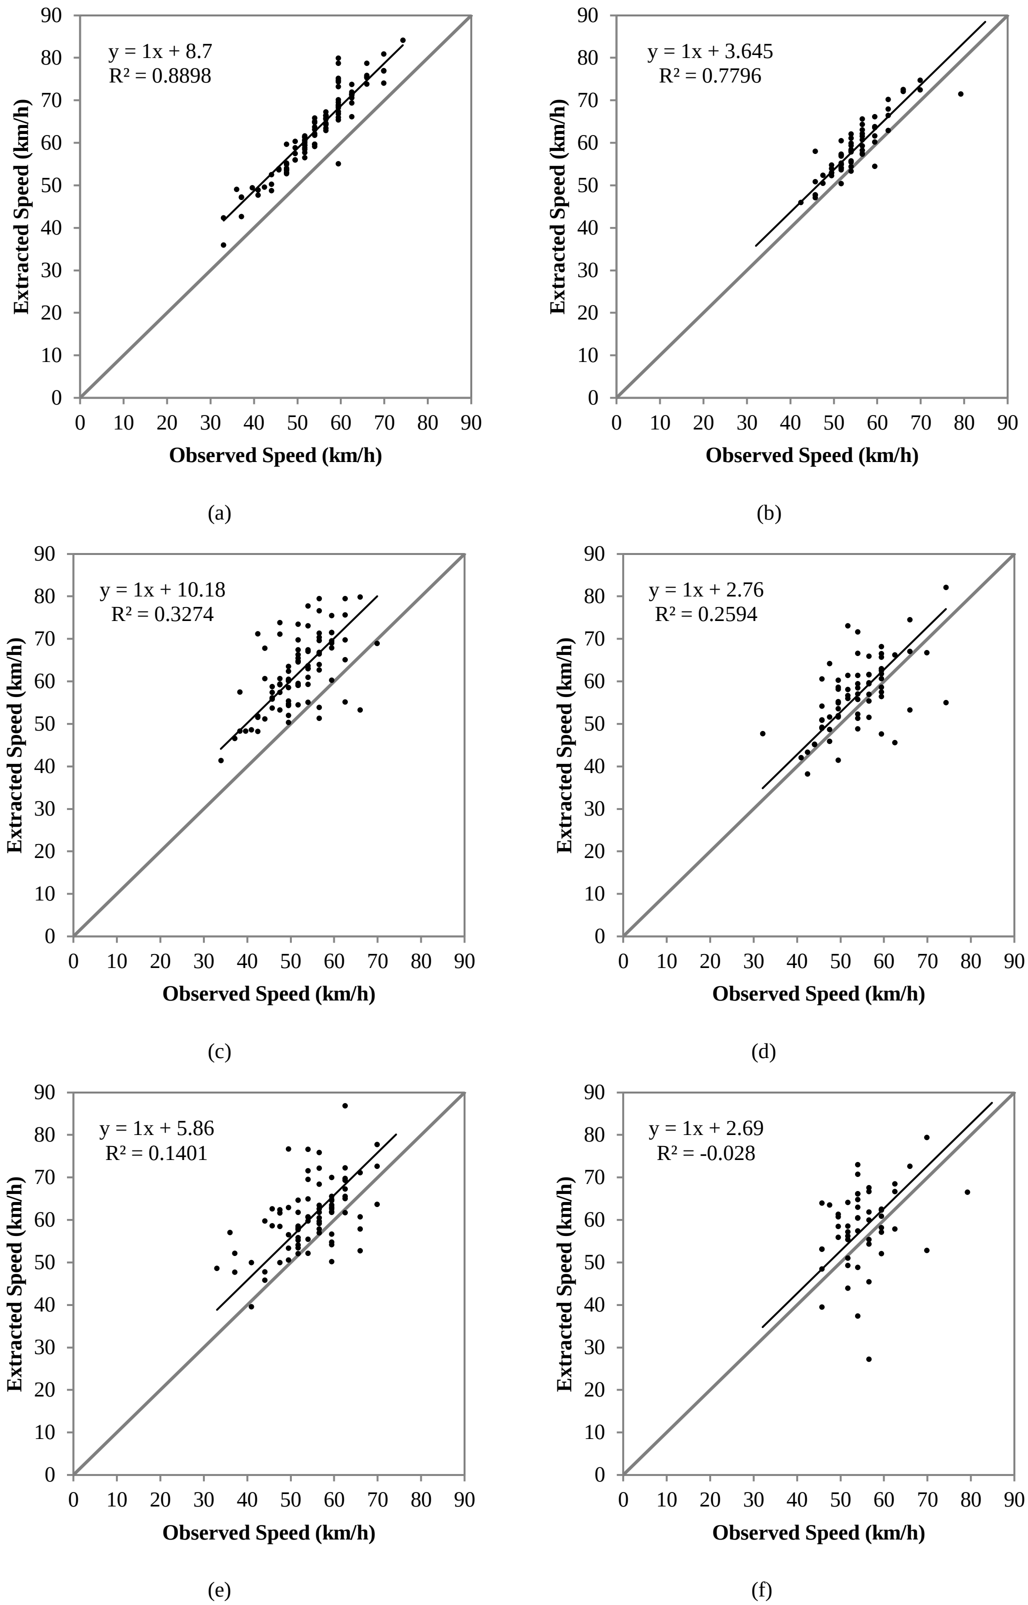

FIGURE 4 Extracted and observed speed for Taschereau, perpendicular Lane 2 (a) and Lane 3 (b), parallel close Lane 2 (c) and Lane 3 (d), and parallel far Lane 2 (e) and Lane 3 (f)

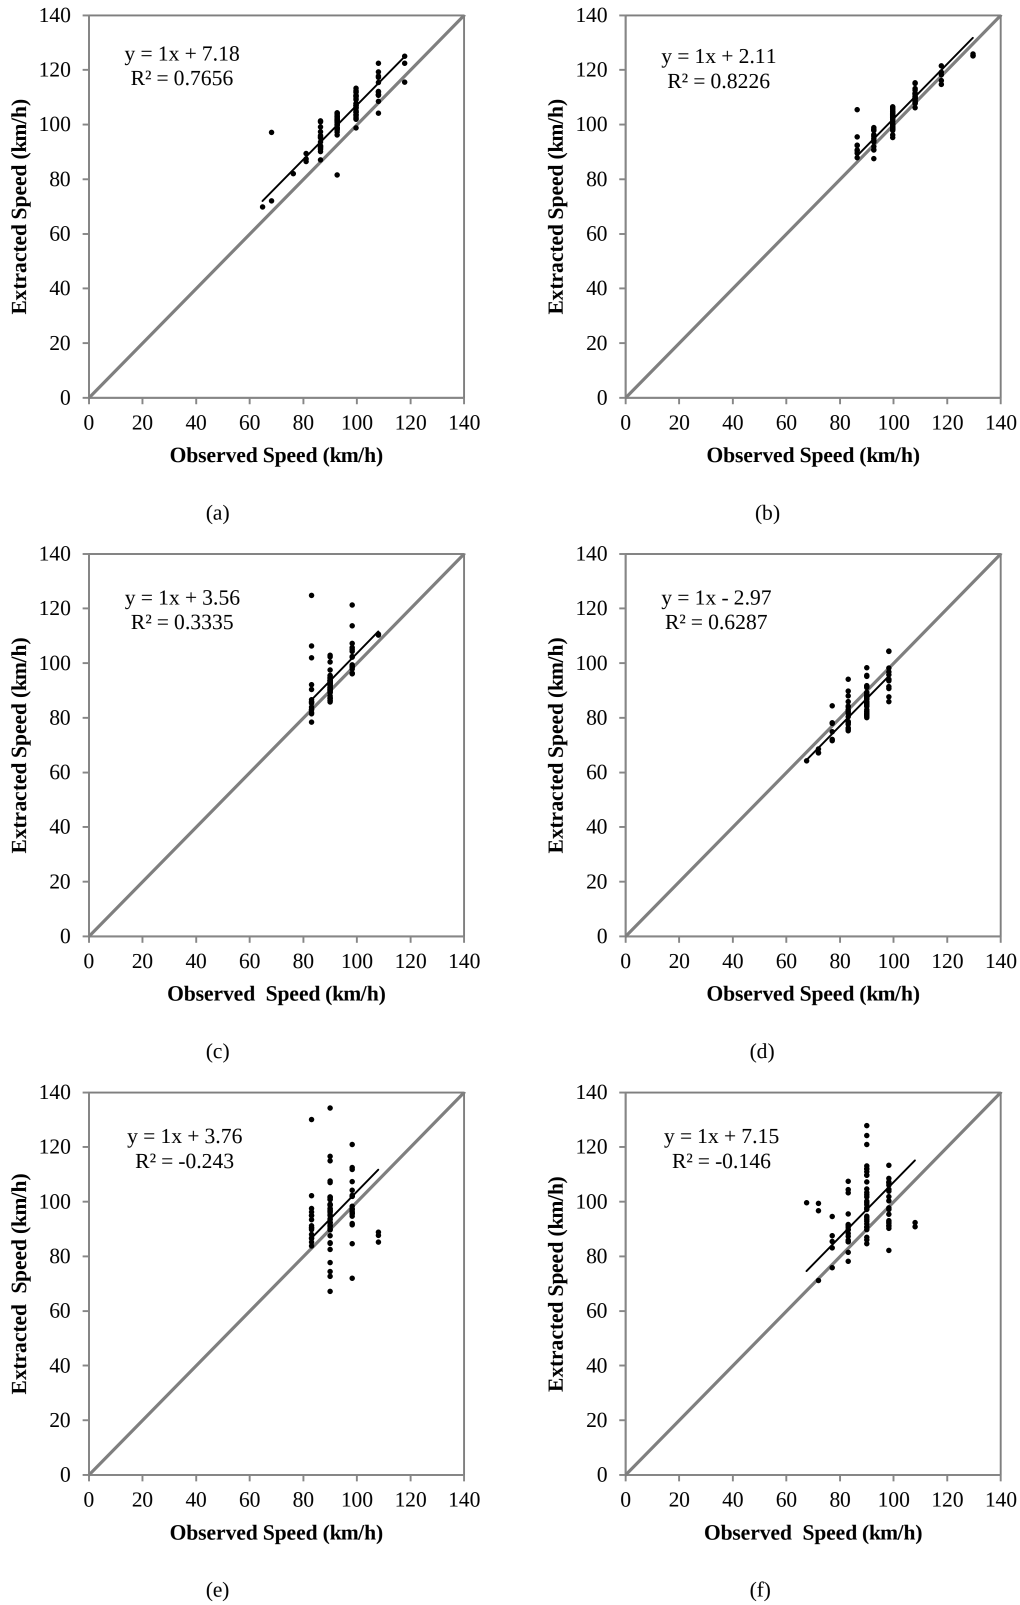

FIGURE 5 Extracted and observed speed for A15, perpendicular Lane 2 (a) and Lane 3 (b), parallel close Lane 2 (c) and Lane 3 (d), and parallel far Lane 2 (e) and Lane 3 (f)

However, the precision error does yield promising results. The perpendicular camera orientation consistently produced the lowest precision error (3% to 4%), followed by the parallel close orientation (4% to 10%) and the parallel far orientation (9% to 13%). This indicates that minimizing the distance between object and camera will improve the precision of speed measurements and that the repeatability of speed measurements is improved when the speed extraction zone is closer to the camera. Precision is consistent across multiple lanes for a given camera orientation. At Taschereau, the data was collected simultaneously for each lane and each orientation. Therefore, the samples at Taschereau contain the same vehicles and differences in observed speeds across study areas are negligible. Collecting data simultaneously and using the same vehicles across each sample data set further validates the results. The precision is always less than the total mean relative error. This is significant, because, if the general over- or underestimation bias can be removed from the data, the overall performance of video-based speed extraction can be improved.

Some differences were observed between the test sites. The precision error at the A15 was consistently lower than the errors reported at Taschereau. While the errors for the perpendicular orientation were approximately equal, the errors for both parallel configurations at Taschereau were between 3% and 5% higher. Some variation is likely explained by difference in camera setup, distance to speed extraction zone, and speed limit. At Taschereau, the camera location on the side of the road skewed the orientation, as it was impossible to install the camera exactly parallel to the lanes, while at the A15, mounting to the overpass structure allowed for true parallel orientations. The camera mounting height also varied between sites. The pole used along Taschereau has a maximum height of 20 feet whereas the overpass height along the A15 was approximately 25 feet above the road surface. Operating speeds at Taschereau were approximately 55 km/h, while at the A15 operating speeds were approximately 100 km/h. At least a portion of the variation can be attributed to these differences.

Discussion

The results above indicate that precision is dependent on camera orientation and may be predictable in nature. Precision is dependent only on the ability of the software to recognize and group features and to track objects throughout the video, an ability that is constant for a given camera orientation. For example, using a perpendicular orientation, the relative ease of object tracking is high, because objects are closest to the camera, features are distinct, and pixels represent smaller real-world distances. In this situation repeatability is high, and even if errors are made in calibrating the tracking software or homography, the error will manifest in all extracted speeds (even inaccurate speeds will be consistently inaccurate). In contrast, accuracy does not appear to be predictable. Accuracy is dependent on the calibration of the software and creation of the homography, and is less dependent on the ability to recognize and track features. Homography and software calibration must be completed for each camera installation, resulting in errors that are unique for each data set. These errors are independent of the tracking method itself, and are therefore less predictable.

Although the perpendicular orientation yields more precise vehicle speeds compared to the parallel close angle, both orientations serve a purpose in different data collection applications. The availability of roadside structures suitable for video data collection is a deciding factor in many situations. One of the most important issues with the perpendicular orientation is the occlusion that occurs when vehicles in separate lanes travel through the extraction zone simultaneously. If the recording system cannot be installed high enough above the road surface, a substantial number of vehicles can be missed, resulting in a false representation of the traffic environment. Alternately, the parallel orientation is necessary when vehicle interactions and lateral movements need to be observed. Ideally, both orientations should be complementary, offering a more detailed picture of the traffic environment.

CONCLUSIONS

This study evaluated the quality of disaggregate vehicle speed data obtained from video-based detection and tracking software. Several factors potentially affect data quality, as observed from multiple camera orientations that were tested along two lanes of arterial and freeway facilities. The traditional mean relative error approach indicated that the video extracted speeds failed to meet the performance requirements for traffic speed collection. In order to better understand the nature of the errors, a technique was proposed for evaluating separately the precision and accuracy of collected data. The proposed segregated error approach divides the mean error into separate values representing accuracy and precision error. In doing so, several of the camera orientations exhibited precision error values less than 5% for disaggregate data, which approaches previously established thresholds for aggregate speed data. In the past, a single error value was used as a threshold for acceptable data quality. Under the proposed scheme, it is clear that the use of a single value is a narrow-minded approach to data collection. The results of this research show that, in general, precision error is small (less than relative mean error) and predictable, while accuracy error may be neither.

This result indicates that if the general over- or underestimation bias can be removed from the data, then the overall quality of video-extracted speeds may be improved. The calibration of video data could be achieved by comparing extracted speeds to a set of manually observed speeds, and subtracting the intercept of the fitted line as demonstrated herein. This would reduce the error of the video-based speed to the magnitude of the precision error, which is often within reasonable standards defined by existing literature. While this method may improve the quality of the extracted speeds, the necessity of 100 manual speed observations may be impractical in some cases. Therefore, future work should explore other methods of calibrating video data which require less manual intervention, using, for example, training and test data sets from other ground truth detectors to estimate the transferability of the calibration (whether the extracted speeds can be consistently corrected across sites to reach acceptable performance). If the intercept calibration approach is preferred, then the Pearson's correlation coefficient should be used to verify the linearity of video-based speeds with respect to the manually observed speeds. In the future, additional work should be completed to verify the linearity of video-based speeds across multiple sites and camera orientations

The segregated error approach should be applied in selecting data collection devices with high levels of precision, recognizing that there are opportunities to eliminate or reduce the over- or underestimation bias. This result has implications for the testing of new traffic detection technologies and the selection of technologies for the process of traffic data collection. Even with large errors, the potential exists to calibrate video data to acceptable levels of performance, so long as precision error is minimized through appropriate selection of camera position and orientation. The greatest benefit of the segregated error approach is that it allows for data collection by devices that might be dismissed as inaccurate by traditional approaches. In the future, the behaviour of more detection technologies should be evaluated using the segregated error approach. Although the proposed approach proved effective, the method should be compared to other statistical approaches in future work.

ACKNOWLEDGMENT

Funding for this project was provided in part by the Ministry of Transportation Quebec and by the Natural Science and Engineering Research Council.

REFERENCES

1. Bahler, S. J., J. M. Kranig, and E. D. Minge. Field test of nonintrusive traffic detection technologies. *Transportation Research Record: Journal of the Transportation Research Board,* Vol. 1643, No. 1, 1998, pp. 161-170.
2. MacCarley, C. A., and J. Slonaker. Automated Consensus-Based Data Verification in Caltrans Detector Testbed. *Transportation Research Record: Journal of the Transportation Research Board,* Vol. 1993, No. 1, 2007, pp. 124-130.
3. Coifman, B. Vehicle level evaluation of loop detectors and the remote traffic microwave sensor. *Journal of transportation engineering,* Vol. 132, No. 3, 2006, pp. 213-226.
4. Freeway detector assessment: aggregate data from remote traffic microwave sensor. *Transportation Research Record: Journal of the Transportation Research Board,* Vol. 1917, No. 1, 2005, pp. 149-163.
5. Kranig, J., E. Minge, and C. Jones. Field Test of Monitoring of Urban Vehicle Operations Using Non-instrusive Technologies: Final Report.In, 1997.
6. Cho, Y., and J. Rice. Estimating velocity fields on a freeway from low-resolution videos. *Intelligent Transportation Systems, IEEE Transactions on,* Vol. 7, No. 4, 2006, pp. 463-469.
7. Laureshyn, A., and H. Ardö. Automated video analysis as a tool for analysing road user behaviour.In, Proceedings of ITS World Congress, London, 8-12 October 2006., 2006.
8. Michalopoulos, P. G. Vehicle detection video through image processing: the Autoscope system. *IEEE Transactions on Vehicular Technology,* Vol. 40, No. 1, 1991, pp. 21-29.
9. D. Koller, K. D., T. Thórhallson, and H.-H. Nagel. Model-based object tracking in traffic scenes. *Computer Vision - ECCV'92,* Springer Berlin Heidelberg, Vol. 588 of Lecture Notes in Computer Science, 1992.
10. Coifman, B., D. Beymer, P. McLauchlan, and J. Malik. A real-time computer vision system for vehicle tracking and traffic surveillance. *Transportation Research Part C: Emerging Technologies,* Vol. 6, No. 4, 1998, pp. 271-288.
11. N. Saunier, T. S., K. Ismail. Large Scale Automated Analysis of Vehicle Interactions and Collisions. . *Transportation Research Record: Journal of the Transportation Research Board*, No. 2147, 2010, pp. 42-50.
12. Gupte, S., O. Masoud, R. F. Martin, and N. P. Papanikolopoulos. Detection and classification of vehicles. *Intelligent Transportation Systems, IEEE Transactions on,* Vol. 3, No. 1, 2002, pp. 37-47.
13. Messelodi, S., C. M. Modena, and M. Zanin. A computer vision system for the detection and classification of vehicles at urban road intersections. *Pattern analysis and applications,* Vol. 8, No. 1-2, 2005, pp. 17-31.
14. Malik, J., and S. Russell. Traffic surveillance and detection technology development: new traffic sensor technology final report. 1997.
15. Dailey, D. J., F. W. Cathey, and S. Pumrin. An algorithm to estimate mean traffic speed using uncalibrated cameras. *Intelligent Transportation Systems, IEEE Transactions on,* Vol. 1, No. 2, 2000, pp. 98-107.
16. Schoepflin, T. N., and D. J. Dailey. Cross-correlation tracking technique for extracting speed from cameras under adverse conditions. *Transportation Research Record: Journal of the Transportation Research Board,* Vol. 1867, No. 1, 2004, pp. 36-45.
17. *Atlas des Transport*. http://transports.atlas.gouv.qc.ca/. Accessed July 15, 2014, 2014.
18. Saunier, N. *Traffic Intelligence*. https://bitbucket.org/Nicolas/trafficintelligence/wiki/Home. Accessed Jun 24, 2014.
19. Saunier, N., T. Sayed, and K. Ismail. An Object Assignment Algorithm for Tracking Performance Evaluation.In *11th IEEE International Workshop on Performance Evaluation of Tracking and Surveillance (PETS 2009), Miami, Fl, USA, June, No. 25*, 2009. pp. 9-17.