Performance Evaluation and Error Segregation of Video-Collected Traffic Speed Data

Paul Anderson-Trocmé, Graduate Researcher, McGill University
Joshua Stipancic, Graduate Researcher, McGill University
Luis Miranda-Moreno, Associate Professor, McGill University

Paper prepared for presentation
at the Measuring Congestion Session

of the 2014 Conference of the
Transportation Association of Canada
Montreal, Quebec

ABSTRACT

Validating the accuracy of sensors and methods is an essential step in the collection of traffic speed data. The accuracy of automated speed data has been evaluated in both small- and in large-scale testing efforts using multiple technologies and methods, as documented in existing literature. In these studies, an important challenge is the creation of a ground truth speed data set that represents actual traffic history. Though inductive loops are standard for data collection, the use of non-intrusive traffic data collection technologies has become increasingly popular. Video-based detectors have demonstrated the ability to substitute conventional detection devices. Computer vision systems and video tracking software provide a wide variety of data, including conventional traffic parameters such as flow and velocity, while preserving a complete record of events. 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 the Traffic Intelligence video analysis software system. To provide robust calibration, several camera orientations were tested along two types of facilities in Montreal, Canada. Video was collected on an urban arterial and a highway section, with cameras oriented both perpendicular and parallel to traffic direction. After calibrating the feature-tracking software, a semi-automated vehicle tracking process was used to extract the vehicle speeds. Comparison to manually observed speeds was undertaken to evaluate the quality of the extracted speeds. Although the traditional mean error approach led to unacceptable results, a new approach was proposed for the evaluation of traffic detection technologies. The proposed segregated error approach divides the mean error values into separate values representing accuracy and precision errors. In doing so, several of the camera orientations exhibited precision error values within the accepted range speed data quality (5%). Even with large errors, video data can be calibrated to acceptable levels of accuracy so long as precision error is minimized through appropriate selection of camera position and orientation.

INTRODUCTION

The collection and analysis of vehicular speed data are essential for any urban transportation system. Systems 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 any given project [3].

Traditional 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, and enable multiple lane detection [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. These systems are able to provide a wide variety of data, from conventional traffic parameters such as flow and velocity [1, 8], to new parameters such as object trajectories which can lead to information on vehicle manoeuvring and traffic conflicts [8, 9]. Perhaps the greatest benefit of video detection is data preservation. Unlike other detectors, video captures and stores the entire series of events as they occurred during data collection. As data needs evolve, video footage can be reanalyzed and data can be refined.

While video detection has many advantages, before any system is relied on for traffic data collection, the accuracy of the system must be comprehensively 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 aggregated 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. Finally, the methods of evaluating error have used a simplistic mean error approach without consideration for more robust analysis. The purpose of this study is to evaluate the accuracy of a video-based detection system, comprised of commercially available video cameras and the Traffic Intelligence video analysis software system. The objectives of this research are to; evaluate the accuracy of automated vehicular speed extraction, recognizing detection is a precursor to speed extraction; to propose a technique for evaluating separately the precision and accuracy of collected data; and to consider the required quality of video-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. Compared to conventional detection devices, vehicle tracking techniques provide additional information in the form of vehicle trajectories [8]. Trajectories are representations of objects and their characteristics over both time and space. The nature of video detection allows vehicles to be tracked continuously over a road segment rather than being detected at a singular location. The analysis of vehicle trajectories can be used to automate otherwise resource-intensive studies, including vehicle manoeuvring [8], lane-changing, queuing patterns [6], automated incident detection [4], and driver behaviour [7]. One application pertinent to this study is conflict analysis. The use of surrogate safety measures has become increasingly popular in order to diminish reliance on crash data. Traffic conflicts are interactions between road users that are sufficiently close to real crashes and are physically and predictably related to crash events. For analysis, additional interpretation modules may be used to complement video detection and object tracking [9]. Once trajectories are extracted, interactions can be categorized as normal operation or conflict. Such techniques can be used to verify the safety improvements of engineering countermeasures without crash data using a before-after approach [10].

The increased sophistication of video detection does not preclude its use in substituting or complimenting traditional traffic detectors. Along with trajectories, video has the ability to extract traditional traffic parameters including flow, speed, headway, and density at both the microscopic, or individual vehicle, level [8] and the macroscopic, or corridor, level [6]. This vehicle data has applications in the calibration of traffic flow models [8] and conflict analysis. As detecting and extracting the position and speed of vehicles provides the basis for conflict analysis, verifying the parameters themselves will lend additional credibility to the technique.

To determine speed, vehicles must first be successfully detected and tracked. In a study of video-based vehicle classification, Gupte et al. [11] achieved a detection rate of 90%, while the system utilized by Messolodi et al. [12] 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 level. This is potentially due to the amount of video that must be processed and the resources required to perform the analyses [6]. Additionally, methods used to assess speed data quality have lacked robustness. Coifman [8] extracted data for velocity, flow, and density collected by video in real-time on a freeway facility. Data was temporally aggregated 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%. Using the same data, Malik et al. [13] showed improvements when post-processing was used in place of real time analysis. Across the study, detection rates varied between 75% and 95%. In daylight conditions and using 5-minute aggregation periods, nearly all samples showed speed error of 5% or less. Data quality was found to vary 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 [14] utilized individual vehicle data, but considered only 190 vehicles from 40 seconds of video data. While able to demonstrate that the mean error in speed detection across all observations was zero, if disaggregate records are considered, the individual mean error varied between -40% and 80%. While the results are sound for estimating mean traffic speed, the technique is clearly not applicable to individual vehicles. Schoepflin [15] utilized a comparison of speed distributions created manually and by video data, noting that the distributions were approximately equal in mean and distribution. The authors claimed this indicates “certain equivalence” between the video data 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 using aggregated data. Using temporally aggregated data for speed analysis effectively eliminates the influence of higher and lower speed vehicles, and aggregation can obscure the performance of a device exhibiting compensating errors [5]. The reduction in error due to aggregation clearly demonstrates the compensating errors present in video extracted speeds.

Several notable issues exist with regard to extracting vehicle speeds from video. Vehicle detection may be hindered by with false or missed detections. False detections involve the tracking of any object that is not a road user. 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 feature tracks, creating inaccurate trajectories and speed estimations [9]. Vehicle position relative to the camera may affect data quality. Vehicles that are further from the camera occupy a smaller number of pixels, may not show distinctive features, and may be difficult to identify and track, leading to potential variability in speed data [6]. Vehicle tracking is inhibited by an overestimation of derivative values. This means that the "distance between two measured points is systematically biased towards longer distances, which results in speed overestimation" [7].

Regardless of its successes, if video detection is to be considered a reasonable alternative to other devices, the data must meet similar accuracy requirements. Bahler [1] indicated that inductive loops exhibiting count errors less than 4% over 1 hour aggregation was of sufficient quality [5]. The same study demonstrated that most commercially available non-intrusive traffic detectors were able to provide counts within 3% of actual, and speeds within 8% [1]. Several issues exist with regards to existing literature on data requirements. Existing literature provides little guidance on acceptable accuracy for microscopic speed data. Most attempts to verify data quality have quantified error based on aggregated data, with little research focused on reasonable accuracy for individual vehicle data. Additionally, there has been no consideration for evaluating device precision and accuracy separately. With this considered, researchers should endeavor to find detectors that "approach the ideal, but fall within some level of tolerance that may vary from application to application" [3].

METHODOLOGY

Site Selection, Instrumentation, and Data Collection

The quality of video-extracted traffic speeds was evaluated in both a highway and arterial environment. Arterials and highways provide variation in geometry and, importantly, in traffic parameters such as speed and volume. Autoroute 15 (A15) in Montreal, Quebec, Canada was the selected freeway facility. The A15 is a major north-south corridor and one of the most heavily travelled highways in Montreal, with an AADT of approximately 90,000 [16]. At the testing location, four lanes are present in each direction, and the posted speed limit is 100 km/h. Boulevard Taschereau in Brossard, Quebec, Canada was the selected arterial facility. Taschereau 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 featured 6 lanes in the westbound direction with a posted speed limit of 50 km/h.

A GoPro Hero 3 video 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 6 hours of video on a single battery charge. The camera is highly portable and provides flexibility in mounting location. Site selection ensured the presence of 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 using a specialized mounting system, 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, as shown in Figure 1b. The close proximity of adjacent poles allowed for the collection of simultaneous video data from multiple orientations.

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 orientation is vital because the ability to utilize multiple orientations improves flexibility and provides more mounting options, which may be highly beneficial for data collection in the urban environment. Three camera orientations were used at each site. The first was a perpendicular orientation, were the camera was positioned over the roadway looking directly across the lanes of travel. It was believed that this orientation would provide the most accurate speeds, as vehicles are as close as practically possible to the camera. Two parallel orientations were also tested. A parallel orientation, were the camera is positioned looking down the roadway, is beneficial if more information, such as vehicle trajectory, is required. This orientation was used with a speed extraction zone approximately 10 m from the camera (parallel close) and with a speed extraction zone approximately 20 m from the camera (parallel far). At least 30 consecutive minutes of video was recorded for every orientation at each site. The locations of the cameras at and their respective study areas are provided in Figure 2**.**

Feature-Based Tracking Algorithm

Data extraction was automated using computer-vision software Traffic Intelligence [17], an open source project developed at Ecole Polytechnique in Montreal, Canada. The program enables users to analyze video, extract vehicle trajectories, and evaluate trajectory data using several built-in 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 a 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. This primarily occurs when vehicle geometry creates several distinct feature groups. Over-segmentation can lead to inflated vehicle counts but does not affect speed accuracy since grouped features still belong to a single vehicle. Over-grouping occurs when multiple vehicles are defined by a single trajectory, due to the proximity of neighbouring features. The over-grouping of objects will lead to inaccurate speed calculations and false volume counts [18]. Given these issues, the most important parameters within Traffic Intelligence are related to the criteria used for grouping features into objects. In order to accurately apply Traffic Intelligence, the key parameters were calibrated by ensuring extracted counts accurately matched manual counts over a sample of the collected video.

Data Extraction and Analysis

Data was extracted 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 a single speed for each object. Speed boxes were created for two lanes, 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. Vehicles were tracked through the distance corresponding to the virtual speed box. Using the length of the speed box and the video frame rate, the speed of the vehicles could be calculated. Manual speeds were matched one-to-one with the extracted speed using the corresponding object number.

Analysis of the extracted speeds was completed in two steps. The mean error for the extracted speeds was calculated for every orientation at each site. The use of mean error is consistent with analysis demonstrated previously within existing detector testing literature. However, this study contends that mean error is insufficient at capturing the true behaviour of detectors. Therefore, to fully understand the behavior within the error present in the extracted data, total error should be divided into precision and accuracy error. A simple method for segregating error in this way is described in the results below. Figure 3 demonstrates the differences between precision and accuracy error for speed data and is further explained in the following section.

RESULTS

Mean Error Approach

A sample of 100 consecutive vehicles was selected in each of the 12 study areas, for which the mean error was calculated for the extracted speeds. The mean error is the simplest way to quantify error in speed detection and is widely used in existing studies. The error was calculated for each individual record by normalizing the difference between the extracted and observed speed. These individual errors were averaged across the sample to yield the mean error, according to

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

Mean error results are presented in Table 1. Extracted speeds exhibit significant difference when compared to manually observed speeds at many of the study areas. At Taschereau, the mean error in five of the six cases exceeds 10%, with the sixth case exhibiting an error of 8%. For a single camera orientation, variation between the lanes is observed. 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 indicate that video extracted speeds do not fall within acceptable limits for traffic detectors at Taschereau. For the A15, the video data is of acceptable quality for exactly half of the study areas. In general, it appears that the quality of video extracted speed data is not acceptable.

Data Visualization

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. These plots utilize a diagonal line to indicate the ideal detector performance (that is, data from an ideal detector follows the line $y=x$). Data points over the line indicate overestimation of speed, while points below the line indicate underestimation. This data visualization is potentially powerful in identifying general trends within the data.

The results for all 12 study areas are presented in Figure 4 and Figure 5.Visually, it was determined that the mass of the observations tended to fall above the ideal diagonal line, indicating a general overestimation, consistent with previous research [7]. To formally investigate this observation, a trend line was added to the plots. The slope of the trend line was held to be 1, such that it would be parallel with the ideal diagonal line. Using this technique, the intercept of the trend line can indicate the direction and magnitude of the general estimation error, while the R-squared value can indicate the precision of the extracted speeds independent of any over- or under-estimation. Based on the intercept values alone, 11 of the 12 study areas exhibit overestimation.

While the intercept provides insight into the accuracy of the extracted speeds, the R-squared value for each trend line reveals the precision, or repeatability of each extracted speed measurement, with a higher R-squared indicating a higher degree of repeatability. In general, the R-squared 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-squared values between 0.26 and 0.63. The worst results were for the parallel far orientation, with R-square values near zero. 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 the lines provide sufficient data fit.

Segregated Error Approach

The visualization exercise proved to be powerful by allowing observation of both precision and accuracy as separate phenomenon through the use of the R-squared and intercept values. However, it would be extremely beneficial to transform these values into values of precision and accuracy error. Utilizing normalized error values maintains the notation utilized in existing literature, and provides an intuitive and communicable comparison between sites and camera orientations. The y-intercept of the fitted lines quantifies the difference between the line-of-best-fit and the ideal detector line, 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 addition of a correction factor equal to the y-intercept of the fitted line. The remaining error can be attributed to device accuracy, as the normalized error between the line-of-best fit and the ideal line at every data point. These values can be calculated simply as

$$Presicion\ Error\ = \frac{1}{100}\sum \frac{|(V_{extracted} - y\ intercept) - V_{observed}|}{V_{observed}}$$

$$Accuracy\ error = Mean\ error - Precision\ error$$

The segregated errors for all sites and orientations are presented in Table 2. If only the mean error values are considered, patterns are difficult to observe. For Lane 3 at both sites, the perpendicular orientation produced the lowest error, the parallel close exhibiting more error, and parallel far exhibiting the most. However, this pattern does not hold for Lane 2 at either site. Based only on the mean error, the quality of video extracted speed data would not be acceptable as discussed previously.

However, patterns in the data do emerge for the segregated errors, especially the precision error. The perpendicular camera orientation consistently produced the lowest precision error, followed by the parallel close orientation and the parallel far orientation, with the highest precision error. Additionally, the precision error is observed to be nearly equal for both lanes within each orientation. The perpendicular orientation resulted in the most precise extracted speeds. 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. The results from Taschereau mirror the trends in precision data, and error was lowest for the perpendicular orientation (4%), followed by parallel close (10%), and parallel far (12% and

13%). Collecting data simultaneously and using the same vehicles across each sample data set further validates the results.

Patterns for accuracy error are not as clear. The accuracy error is typically not equal across lanes at a single site. When considering camera orientation, it is clear that the parallel far orientation typically produced the lowest accuracy error. This indicates that the repeatability of speed measurements is much better when the speed extraction zone is close to the camera (i.e. either the perpendicular or parallel close orientations). However, the speed that is averaged across multiple consecutive vehicles will be much closer to the true mean speed when the parallel far orientation is utilized. In other words, the overestimation problem is less prevalent for more distant objects.

Discussion

While these general trends are easily observed, differences between sites are also of interest. The precision error at the A15 was consistently lower than the errors from 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 and distance to speed extraction zone. At Taschereau, the camera's location on the side of the road limited skewed the orientation to be not exactly parallel to the lanes. 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. At least a portion of the variation can be attributed to these differences.

The results above indicate that precision error is highly dependent on camera orientation and is predictable in nature. Precision error is dependent only on the ability of the software to recognize and group features and to track objects throughout the video. This ability 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 a smaller real-world distance. In this situation repeatability is high, and even if errors are made in calibrating the tracking or speed extraction modules, the error will manifest in all speeds. In other words, even inaccurate speeds will be consistently inaccurate. In contrast, accuracy does not appear to be predictable. Accuracy is more dependent on the parameterization of the software and less dependent on object tracking. Homography and software calibration must be completed for each camera installation, leading to errors that are unique for each individual orientation. This helps to explain why patterns are less visible for accuracy error.

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 used to provide complementary data, offering researchers 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 feature-tracking software. Multiple camera orientations were tested along two lanes of arterial and freeway facilities. The first objective of this study was to evaluate the accuracy of automated extracted vehicular speeds. The traditional mean error approach indicated that the video extracted speeds were not within an acceptable range for use as a data collection method. In order to better understand the nature of the errors, the second study objective was to propose a technique for evaluating separately the precision and accuracy of collected data. Although the traditional mean error approach did not lead to acceptable results, a new approach was developed for the evaluation of traffic detection technologies. 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 within the accepted range for data collection technologies (5%).

Both of these results are especially important from the perspective of the third study objective; to consider the required quality of video-collected traffic 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 and predictable, while accuracy error may be neither. This is crucial, because accuracy can be calibrated for in video-extracted speeds. Although video data displays a total mean error much greater than 5% in most cases, the precision, or repeatability, of video extracted speeds is often within reasonable limits for data quality. This method provides the ability to compensate for the over-estimation problem present in many video-based detection systems. Importantly, even devices exhibiting a high accuracy error can be calibrated to provide consistent and accurate speeds, if the y-intercept of a fitted line is used to adjust all speed measurements.

The segregated error approach can guide selection of devices with high levels of precision regardless of the level of accuracy. 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, video data can be calibrated to acceptable levels of accuracy, 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 future studies, the validity of calibrating detection devices using this approach will be considered. Importantly, more testing sites should be utilized to ensure that the patterns demonstrated herein apply universally. The behaviour of more detection technologies should be evaluated with respect to the segregated error approach. Although the proposed approach proved effective, the method should be compared to other statistical approaches, and Pearson’s correlation coefficient should be used to confirm linearity of video extracted speeds.

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, 1998. **1643**(1): p. 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, 2007. **1993**(1): p. 124-130.
3. Coifman, B., *Vehicle level evaluation of loop detectors and the remote traffic microwave sensor.* Journal of transportation engineering, 2006. **132**(3): p. 213-226.
4. Coifman, B., *Freeway detector assessment: aggregate data from remote traffic microwave sensor.* Transportation Research Record: Journal of the Transportation Research Board, 2005. **1917**(1): p. 149-163.
5. Kranig, J., E. Minge, and C. Jones, *Field Test of Monitoring of Urban Vehicle Operations Using Non-instrusive Technologies: Final Report*. 1997.
6. Cho, Y. and J. Rice, *Estimating velocity fields on a freeway from low-resolution videos.* Intelligent Transportation Systems, IEEE Transactions on, 2006. **7**(4): p. 463-469.
7. Laureshyn, A. and H. Ardö. *Automated video analysis as a tool for analysing road user behaviour*. 2006. Proceedings of ITS World Congress, London, 8-12 October 2006.
8. Coifman, B., et al., *A real-time computer vision system for vehicle tracking and traffic surveillance.* Transportation Research Part C: Emerging Technologies, 1998. **6**(4): p. 271-288.
9. Saunier, N. and T. Sayed, *Automated analysis of road safety with video data.* Transportation Research Record: Journal of the Transportation Research Board, 2007. **2019**(1): p. 57-64.
10. Sayed, T., et al., *Feasibility of Computer Vision-based Safety Evaluations.* Transportation Research Record: Journal of the Transportation Research Board, 2012. **2280**(1): p. 18-27.
11. Gupte, S., et al., *Detection and classification of vehicles*. Intelligent Transportation Systems, IEEE Transactions on, 2002. **3**(1): p. 37-47.
12. 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, 2005. **8**(1-2): p. 17-31.
13. Malik, J. and S. Russell, *Traffic surveillance and detection technology development: new traffic sensor technology final report.* 1997.
14. 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, 2000. **1**(2): p. 98-107.
15. 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, 2004. **1867**(1): p. 36-45.
16. COMB. *Seeking highest traffic-volume intersectionsfor your OOH campaigns?* 2013 [cited 2013 June 24]; Available from: comb.org/site/ywd_comb/assets/pdf/seeking_highest_traffic-volume_intersections_for_your_ooh_campaigns.pd.
17. Saunier, N. *Traffic Intelligence*. 2013 [cited 2013 Jun 24]; Available from: https://bitbucket.org/Nicolas/trafficintelligence/wiki/Home.
18. 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*. 2009.

TABLE 1 Mean Error Values for Video-Extracted Speeds

	Mean 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

TABLE 2 Segregated Error Values for Video-Extracted Speeds

	Lane 2 Error			Lane 3 Error		
	Mean	**Precision**	**Accuracy**	**Mean**	**Precision**	**Accuracy**
Taschereau						
Perpendicular	0.16	0.04	0.12	0.08	0.04	0.04
Parallel close	0.22	0.10	0.12	0.12	0.10	0.02
Parallel far	0.15	0.12	0.03	0.15	0.13	0.02
A15						
Perpendicular	0.08	0.03	0.05	0.03	0.03	0.00
Parallel close	0.05	0.05	0.00	0.05	0.05	0.00
Parallel far	0.10	0.09	0.01	0.12	0.10	0.02

(a) (b)

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

(a) (b)

FIGURE 2 Camera orientations and study areas for (a) arterial and (b) highway locations

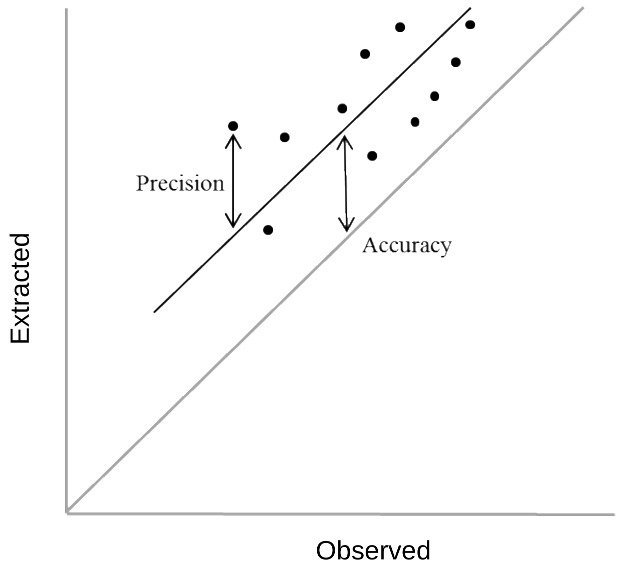

FIGURE 3 Demonstration of error segregation

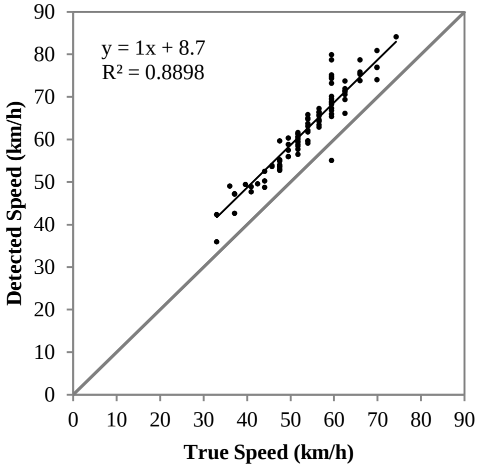

(a)

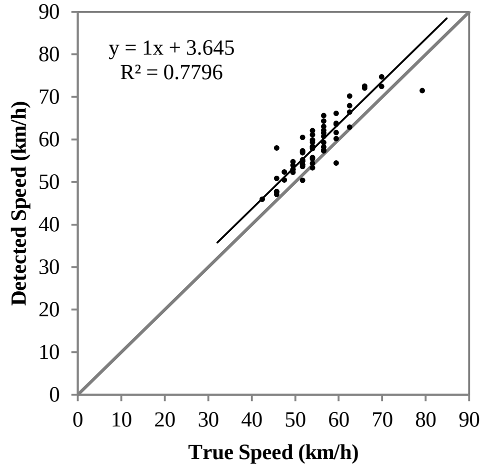

(b)

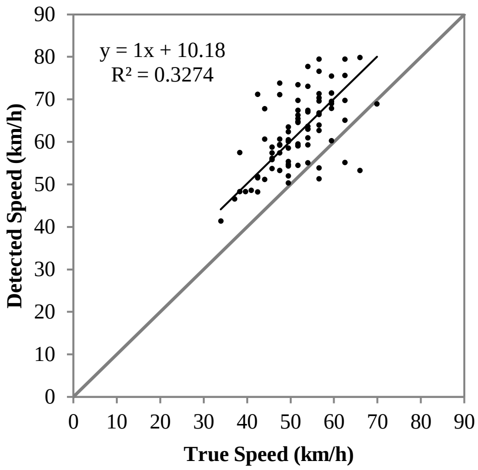

(c)

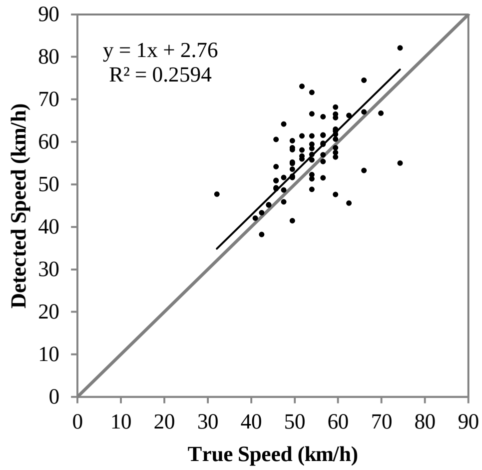

(d)

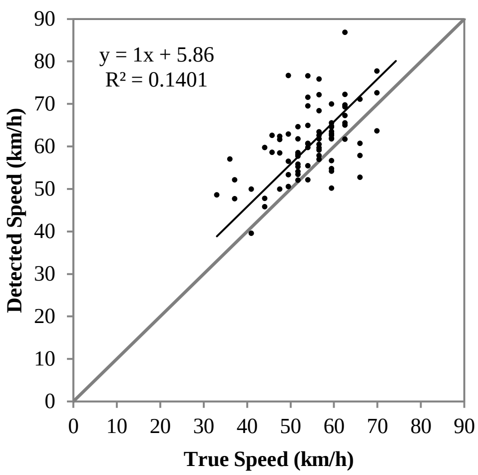

(e)

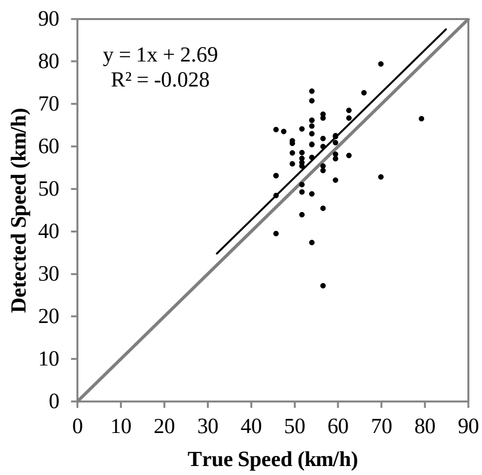

(f)

FIGURE 4 Detected and true 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)

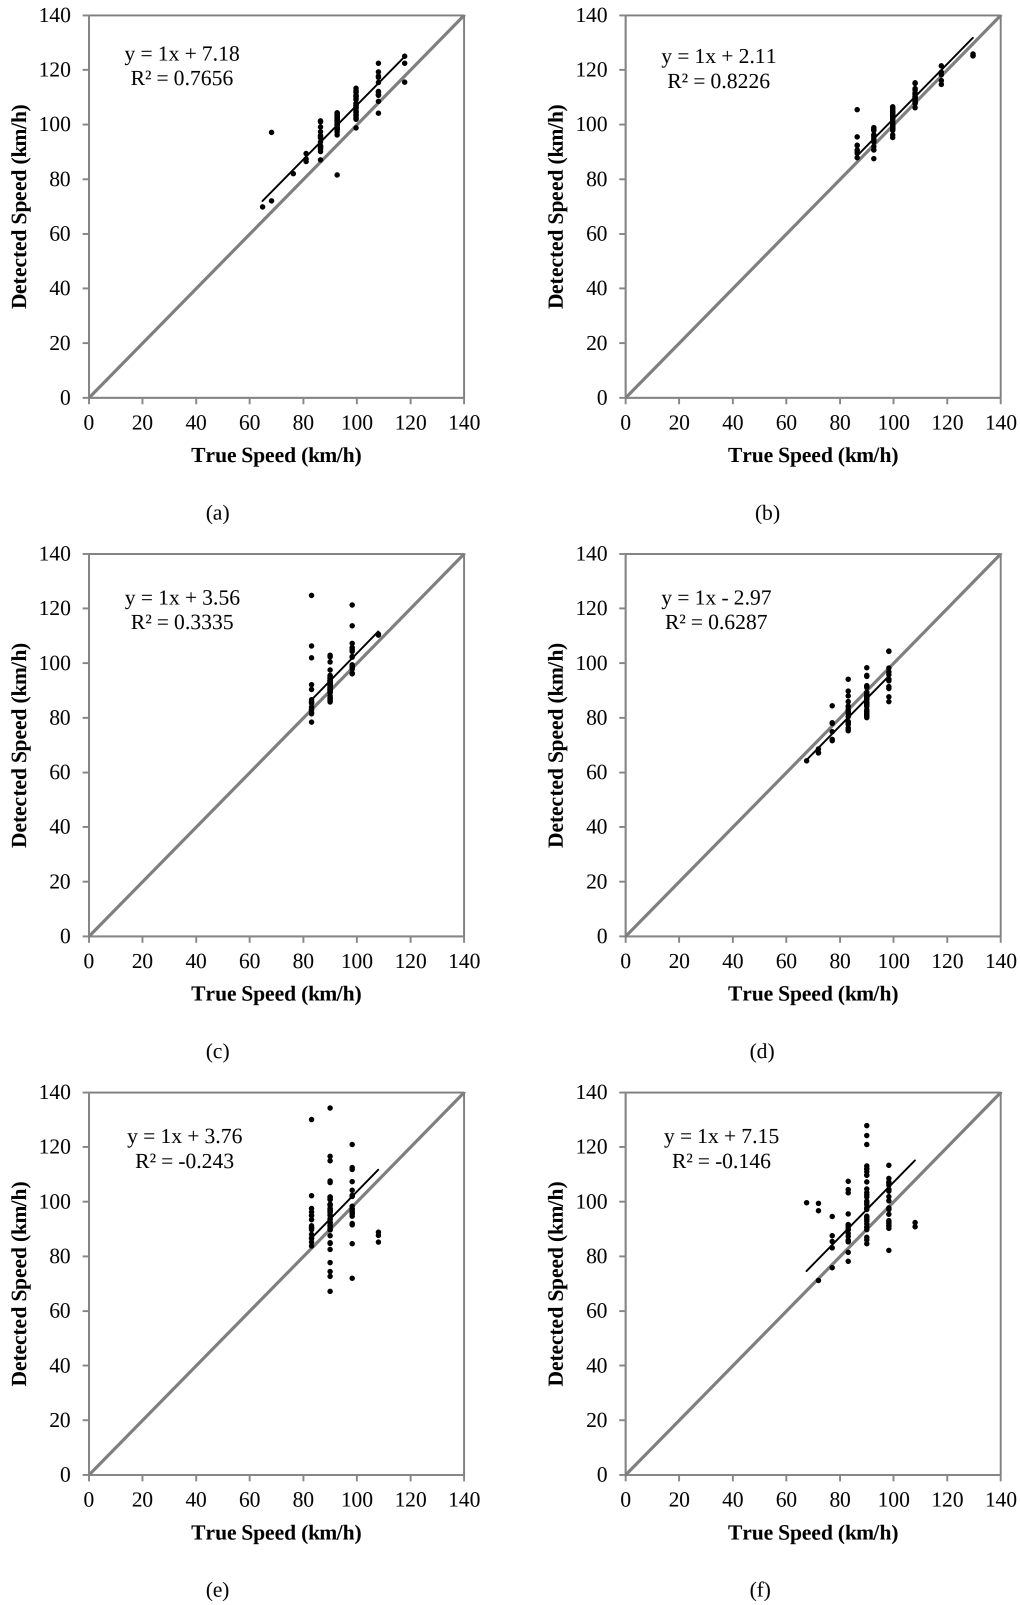

FIGURE 5 Detected and true speed for Autoroute 15, perpendicular Lane 2 (a) and Lane 3 (b), parallel close Lane 2 (c) and Lane 3 (d), and parallel far Lane 2 (e) and Lane 3 (f)