Available online at www.sciencedirect.com

ScienceDirect

Transportation Research Procedia 00 (2026) 000–000

Transportation Research Procedia

www.elsevier.com/locate/procedia

Road Safety and Simulation 2026 – RSS2026

Conflict detection and analysis in urban arterial roads of Brasília, Brazil, using HD-CCTV monitoring cameras and the YOLO model.

Julia Alves Porto[a]*, Stella Roussou[a], Michelle Andrade[b], Apostolos Ziakopoulos[a], George Yannis[a]

[a]*National Technical University of Athens, Department of Transportation Planning and Engineering, 5 Iroon Polytechneiou Street, 15773, Athens, Greece*
[b]*University of Brasília, Department of Civil and Environment Engineering, Anex SG-12, University Campus Darcy Ribeiro, 70910-900, Federal District, Brazil*

Abstract

Traffic conflict analysis has been advocated in the past years as a surrogate measures of safety, complementing or even substituting crash-based assessments by enabling proactive evaluation of road safety. Although traffic conflicts occur more frequently than crashes, the require formal definition through indicators, either based on individual movements (e.g. harsh braking or harsh swerving) or on spatiotemporal proximity measures. This work evaluates the feasibility of automated traffic conflict detection at four major urban intersections in Brasília, Federal District, Brazil, using pre-existing surveillance footage. Video data were collected with HD-CCTV cameras from the regional monitoring system. Motorized road users were detected using YOLOv11m, while object tracking was implemented in Python following state-of-the-art approaches from the literature to estimate frame-by-frame trajectories. These trajectories were subsequently vectorized and used to calculate Time to Collision (TTC) between pairs of vehicles. The results demonstrate a substantial improvement over previous automated safety monitoring efforts conducted in the same area. However, steep camera angles, limited image quality, and the erratic motion of motorcyclists still hinder robust tracking performance. The tracking pipeline performs reliably for buses and passenger cars, although occasional duplicate detections occur when temporal matching is disrupted. In contrast, motorcyclists are not consistently detected under real-world conditions, and trucks tend to be over detected. Despite these limitations, the conflict detection framework provides meaningful insights into interactions among road users and highlights the effectiveness of complete traffic flow separation in reducing hazardous encounters, even in the presence of motorcycles. Future work will focus on improving object matching criteria to enhance trajectory continuity and prediction stability. More broadly, this research demonstrates the potential of integrating tracking algorithms, surveillance systems, and surrogate safety measures for scalable and continuous monitoring of urban road safety.

Peer-review under responsibility of the scientific committee of the Road Safety and Simulation 2026 – RSS2026
Keywords: traffic conflict, surveillance footage, TTC, video data, road safety

* Corresponding author. Tel.: +30-694-996-1497
E-mail address: julia_porto@mail.ntua.gr

Peer-review under responsibility of the scientific committee of the Euro Working Group on Transportation Annual Meeting 2025 - EWGT2025

1. Introduction

Road Safety remains a major concern worldwide, particularly in urban environments where interactions between different types of road users are frequent and complex. According to the international road safety research, traditional safety assessment methods have historically relied on crash data to identify hazardous locations and evaluate safety performance (Elvik, 2008). However, such events are very rare and usually involve a combination of random and systematic factors. This makes statistical analysis difficult, especially when small time periods and geographical areas are considered (Lord and Mannering, 2010).

In addition to this, the issue of data incompleteness arises mainly because minor crashes and those involving vulnerable road users, such as pedestrians and cyclists, tend to be underreported (Yannis et al., 2014). This leads to a delay in the assessment of safety issues on the road, considering that safety problems tend to be identified only when a large number of crashes have already occurred (Tarko et al., 2009). As such, road safety research has, over time, increasingly adopted a more proactive approach to safety, in contrast to the traditional reactive approach that relies on data collected from crashes (Gettman and Head, 2003).

To overcome the limitations of crash-based safety analysis, researchers are using Surrogate Measures of Safety (SMoS) that are based on conflicts and interactions among road users rather than actual crashes (Tarko et al., 2009). Using SMoS, safety conditions can be evaluated proactively because they are capable of detecting potential danger before a crash happens (Sayed and Zein, 2018).

A traffic conflict is generally defined as an observable situation in which two or more road users approach each other in space and time to such an extent that a collision would occur unless at least one of the road users performs an evasive maneuver (Archer, 2005). Traffic conflicts occur much more frequently than crashes, which allows researchers to analyse safety conditions using shorter observation periods and smaller datasets compared to crash-based analysis (Gettman and Head, 2003).

Past studies have established the potential of using traffic conflicts as a measure of safety, especially in intersections and urban arterial roads, which experience high traffic conflicts between vehicles, pedestrians, and cyclists (Sayed & Zein, 2018; Laureshyn et al., 2017). Consequently, the application of SMoS has been widespread in road safety, traffic simulation, and video safety analysis.

The most frequently used Surrogate Measures of Safety are Time to Collision (TTC), Post Encroachment Time (PET), Deceleration Rate (DR), and Gap Time (GT) (Laureshyn et al., 2017). Time to Collision is defined as “the total time before a potential conflict occurs between two road users, assuming that both road users continue on their current course and at their current speed” (Zheng et al., 2019). The measure is widely applied in detecting conflicts, especially in rear-end and longitudinal conflicts (Laureshyn et al., 2017).

One of the main challenges in traffic conflict analysis is the collection of accurate trajectory data for road users. Traditionally, traffic conflicts were identified through manual observation of video recordings, a process that is time-consuming, labour-intensive, and subject to observer bias (Saunier and Sayed, 2010).

In addition, many cities do not have detailed trajectory data available, and crash databases often lack information about vehicle speeds, trajectories, and interactions before the crash occurred (Lord and Mannering, 2010). This lack of detailed behavioural data makes it difficult to understand the mechanisms that lead to crashes and limits the ability to perform proactive safety analysis (Yannis et al., 2014).

With the increasing availability of traffic monitoring cameras and video surveillance systems, video data has become an important source of traffic information. Video recordings allow the extraction of trajectories, speeds, accelerations, and interactions among the road users, which can be used to compute the surrogate safety measures. However, challenges have been identified in the use of video data, including occlusion, camera angle, lighting, and camera calibration, which may affect the accuracy of the extracted information (Ismail et al., 2009).

In the past few years, advancements have been achieved in the automation of traffic analysis with the help of computer vision and machine learning approaches. The first automated approach for traffic analysis was based on the use of background subtraction and feature tracking for the detection and tracking of vehicles from video recordings (Coifman et al., 1998). Later approaches have utilized automated trajectory extraction methodologies for the calculation of surrogate safety metrics (Saunier and Sayed, 2010).

Recent advancements in the field of traffic analysis have utilized the effectiveness of deep learning approaches for the detection of objects in traffic scenes. CNN models, as well as object detection models like Faster R-CNN, SSD, and YOLO (You Only Look Once), have been used for the detection and classification of road users from video recordings (Redmon et al., 2016). These models enable the real-time detection of vehicles, pedestrians, and cyclists from video recordings, which can then be used along with trajectory extraction models like the optical flow approach for the calculation of safety metrics (Wojke et al., 2017). However, real-world applications still face challenges, especially because of occlusions and adverse weather conditions, which reduce road-user detection performance using vision-based data alone. In addition, perspective effects introduce substantial scale variation within images, causing the apparent size of objects to change according to their position in the scene and further complicating the detection task (Abdel-Aty et al., 2023).

The use of automated video analysis based on deep learning models has become a key factor in the evaluation of road safety, as it has the potential to deal with a large amount of information regarding road users and provide in-depth information regarding their interaction (Ismail et al., 2009). As a result, computer vision-based traffic conflict analysis is increasingly used for proactive safety assessment, infrastructure evaluation, and traffic management applications.

In this research, we assess the use of video footage from operational traffic monitoring systems applied currently in Brasília, Brazil, for collecting road safety data relying on automated object detection and trajectory extraction. An early attempt has been made based on the computer vision tracking tool developed by Saunier and Sayed (2010). Since then, updates in computer vision, especially the advances in object detection and tracking, have allowed for a higher performance for safety monitoring than what was originally achieved in Porto and Andrade (2023). Besides discussing the framework to allow data extraction, this paper discusses the limitations of the current monitoring system of Brasília for safety assessment and provides a profile of risk behavior on the monitored intersections.

2. Data and methods

2.1. Data collection

The data used in this analysis was collected by the traffic organization *Departamento de Estradas de Rodagem do Distrito Federal* ("Federal District's Department of Driving Roads"), DER, and it was collected on October of 2021. For each one of the four locations, three hours of footage were provided, collected through the camera surveillance at the nearest intersection, which are al HD-TV type for the surveyed locations, with 1080 x 1920 pixel size and 30.0 frames per second (fps) rate. An excel file with the cameras type, denomination and geographical locations was also provided, allowing for an easy check of their position. Figure 01 presents a frame of each location.

(a) (b) (c) (d)

Fig. 1. Video frame for each location, respectively (a) Local 01: DF; (b) Local 02: EPNB; (c) Local 03: EPTG and (d) Local 04: PS

2.2. Pre-processing

To make the video size shorter and avoid possible discontinuities or corruptions in the file, each video was divided into a 15-minute segment with a 4 second overlap. Therefore, for each hour of video there were four parts: (i) 00:00:00 – 00:15:02; (ii) 00:14:58 – 00:30:02; (iii) 00:29:58 – 00:45:02; and (iv) 00:44:58 – 01:00:00.

To make a ground-based projection of the trajectories and, therefore, a realistic metric of speed, time and trajectory lengths, a homography matrix was applied using OpenCV’s “findHomography” function. The aerial image for the projection was fetched from Google Earth for each data location. Furthermore, the projection local distortion was calculated and, based on the results, a mask area was established to define the video region where the object detection would be applied. The homography calculation and mask definition was performed once for each location and stored to eliminate possible pixel distortion.

2.3. Object detection and tracking

The overall tracking strategy follows closely previous literature on the topic (Saunier and Sayed, 2010; Ventura et al., 2025; Wojke et al., 2017), written in Python and relying mainly on OpenCV library and YOLO object detection model. Algorithm 1 describes the structure of the algorithm.

Algorithm 1. Object detection and tracking

```
While video is valid, for every frame:
    If first frame:
        Detect road users
        Create track objects
    End if
    Propagate trajectory
    Project coordinates to aerial image
    If frame index is divisible by detection rate:
        Detect road users
        Associate new road users with previous tracked objects
        Update track information
        Update lost tracks list
        Create new tracks for new detections
    End if
End for
```

Object detection was performed using YOLOv11m (Redmon et al., 2016), selected for its balance between accuracy and computational cost. The model was restricted to four road user classes (car, bus, truck, and motorcycle), with confidence and IoU thresholds of 0.50. All thresholds were selected empirically. Each track stored the information required for object association and trajectory reconstruction, including object class, bounding box coordinates, feature points, and image- and projection-plane trajectories.

The trajectory propagation is done with OpenCV’s optical flow estimation with the Lucas-Kanade method, which is dependent on the detection of good features to track, also done with OpenCV’s library, using Shi-Tomasi Corner Detector method. For every frame, the corners would be updated and the trajectory propagated both on the image plane and on the projection plane.

To optimize the speed of the tracking algorithm, we performed ID matching every two frames. On the empirical trials, it showed the same accuracy for matching IDs than performing it every frame, at a much faster rate. We also tested different features to guide the matching algorithm (colour histogram, deep features from the YOLO model) but found that the object’s centroid provided the most consistent matching guide for our dataset. The new detected bounding box’s centroid was compared with the last propagated centroid, based on the detected corners, if any, or on the previous bounding box centroid. The modified Jonker-Volgenant algorithm from SciPy was used to perform linear sum assignment and match new and previous detections, creating new tracks or recording a missed frame, if needed.

On GPU NVIDIA GeForce RTX-2080, CUDA version 12.6, and Intel® Core™ i7-8700 CPU, this tracking algorithm was capable to process 15 minutes of video in 50 minutes.

2.4. Speed and Time to Collision

Since the trajectory tracking did not include a smoothing filter, smoothing was performed during the displacement and speed calculation to minimize bias caused by lateral spikes and frame-by-frame noise. To address that, a 15-frame rolling window was used to calculate displacement, and the speed was calculate per 10 frames, therefore roughly every 0.33s. The projection coordinates were first multiplied by the meter per pixel rate to provide a real-world displacement and speed rates. Furthermore, the maximum speed was clipped at 40 m/s to avoid unrealistic speed measurements.

Finally, Time to Collision (TTC) was calculated per frame, and per interaction for every two vehicles appearing simultaneously at a frame. The shape of each vehicle was included, yet the relative speed and distance relative to each object's centroid were used for the TTC calculation. To avoid unrealistic detections due to tracking issues on the borders of the region of interest, the first and last 30 frames (1 second) of each object was discarded prior to TTC calculation. Lastly, a potential conflict (triggered interaction) was flagged if the interaction had a TTC value ≤ 1.5s.

3. Results

On Table 1, it is identified the number of vehicles tracked automatically and manually, the mean and median speeds, the outlier rate, the approximate speed autocorrelation for a random track at 20 and 60 lagged frames, and the total and triggered number of interactions, provided by the automated pipeline, at 15-min intervals. On Table 2, the traffic composition with automated/manual count is presented.

Table 1. Volume, speed and interaction results

Local	Starting time	Traffic volume (manual count)	Traffic volume (automatic count)	Mean speed km/h	Median speed km/h	**Outlier rate (%)**	**Speed Autocorrelation (20 - 60)**	Total interactions	Triggered interactions
01	16:45	820	893	24.51	19.64	7.80	0.2 – 0.1	4578	132
01	17:00	839	857	19.03	11.84	8.70	0.2 – 0.1	4915	170
01	17:15	804	919	18.44	9.79	8.85	0.2 – 0.1	4548	123
01	17:30	824	901	19.06	11.63	8.83	0.2 – 0.1	4720	163
01	17:45	689	942	17.98	10.26	8.17	0.2 – 0.1	5145	191
01	18:00	789	948	18.53	9.45	8.84	0.2 – 0.1	4254	133
02	17:00	940	977	15.39	12.67	0.00	0.6 – 0.5	3795	17
02	17:15	905	1006	13.26	10.91	2.65	0.5 – 0.4	4572	23
02	17:30	895	926	14.44	12.68	2.43	0.5 – 0.4	3890	12
02	17:45	912	935	14.08	12.16	3.39	0.5 – 0.4	4391	20
02	18:00	941	932	15.31	14.05	2.66	0.5 – 0.4	4013	12
02	18:15	944	919	16.17	15.22	1.95	0.7 – 0.5	3720	13
02	18:30	949	978	17.73	16.91	0.70	0.7 – 0.6	3812	11
02	18:45	901	936	15.14	13.83	0.99	0.7 – 0.6	4121	17
02	19:00	890	934	17.60	17.46	0.52	0.7 – 0.6	3394	11
02	19:15	656	639	22.37	19.72	0.57	0.7 – 0.7	1007	1
02	19:30	875	862	38.76	41.22	5.99	0.4 – 0.4	1069	5
02	19:45	736	813	13.24	8.68	5.91	0.8 – 0.7	2767	9
03	16:30	783	1096	24.11	21.77	3.05	0.3 – 0.2	4507	81
03	16:45	756	970	25.34	24.24	2.74	0.4 – 0.3	3798	79
03	17:00	771	996	27.87	27.50	3.39	0.3 – 0.2	3799	81
03	17:15	788	1091	24.60	24.36	2.76	0.4 – 0.4	4728	79
03	17:30	762	938	28.31	27.84	3.18	0.3 – 0.2	3542	67
03	17:45	765	973	28.33	27.19	4.29	0.3 – 0.2	3265	82
03	18:00	788	1050	25.02	24.38	3.49	0.3 – 0.3	3225	65
03	18:15	798	1116	28.94	27.80	4.87	0.2 – 0.2	2343	53
03	18:30	759	1160	26.78	26.33	4.14	0.3 – 0.2	1869	30
03	18:45	771	1142	26.60	26.11	3.94	0.3 – 0.2	1600	26
03	19:00	841	1194	25.86	25.43	4.07	0.3 – 0.2	2092	22
03	19:15	830	1197	28.44	27.06	5.07	0.2 – 0.1	2050	27

Table 1. Continuation. Volume, speed and interaction results

Local	Time	Traffic volume (manual count)	Traffic volume (automatic count)	Mean speed km/h	Median speed km/h	**Outlier rate (%)**	**Speed Autocorrelation**	Total interactions	Triggered interactions
04	16:30	917	1232	34.01	34.31	0.96	0.5 – 0.4	2216	4
04	16:45	1091	1505	38.96	35.81	8.08	0.5 – 0.3	2633	4
04	17:00	1126	1524	31.72	32.92	1.34	0.6 – 0.3	2984	6
04	17:15	1201	1587	33.87	34.38	7.42	0.5 – 0.3	3181	1
04	17:30	1119	1492	33.89	34.34	1.00	0.5 – 0.2	2894	2
04	17:45	1164	1606	33.03	34.33	0.91	0.5 – 0.3	3079	4
04	18:00	1224	1571	32.60	34.38	0.92	0.6 – 0.3	3314	3
04	18:15	1296	1587	32.69	32.93	9.98	0.5 – 0.3	2402	3
04	18:30	1145	1419	33.45	32.67	4.96	0.4 – 0.3	1326	3
04	18:45	1020	1256	34.76	34.97	7.95	0.4 – 0.2	1109	3
04	19:00	1000	1294	33.41	33.48	6.70	0.4 – 0.2	1007	1
04	19:15	1009	1222	34.36	34.67	7.67	0.4 – 0.2	1094	3

Table 2. Traffic Composition (%)

Local	Bus	Car	Motorcycle	Truck	Bus	Car	Motorcycle	Truck
	Manual count				Automatic count			
01: DF	4.3	74.1	15.4	6.2	3.9	77.5	6.2	12.4
02: EPNB	3.4	79.1	14.4	3.1	1.6	85.2	8.8	4.4
03: EPTG	1.0	88.7	9.4	0.8	1.1	88.7	7.6	2.6
04: PS	0.5	85.8	11.4	2.4	0.4	88.5	5.0	6.1

Automated counts generally overestimated traffic volumes, primarily due to fragmented trajectories that caused some vehicles to be counted multiple times. The bias was observed at all sites except Local 02, suggesting that camera geometry played a larger role in tracking performance than traffic volume, as can be seen by the different configurations on Figure 1.

Table 2 indicates that model performance is lower for motorcycles and trucks. Motorcyclists are difficult to detect in footage data due to their susceptibility to occlusion and erratic movements. Besides that, automated motorcycle detections decreased disproportionately relative to manual counts at Locations 02 and 04 after sunset, indicating reduced detector performance under deteriorating lighting conditions. Truck volumes, on the other hand, were generally overestimated, mainly because of fragmented trajectories and duplicate detections.

Table 3 compares trajectory accuracy between the proposed pipeline and Porto and Andrade (2023), based on manual trajectory inspection.

Table 3: Tracking accuracy comparison with previous pipeline

Local	Previous Pipeline (%)	Current Pipeline (%)	Improvement (%)
01	36	40	4
02	75	91	16
03	74	94	20

Trajectory accuracy improved substantially at Locations 02 and 03, increasing from approximately 75% to over 90%. The largest gain was observed at Local 03, where YOLO-based detections reduced trajectory fragmentation caused by signal-related occlusions. Local 01 remained challenging because of congestion, steep viewing angles, and vehicle overlap.

Mean and median speeds were generally consistent across locations, suggesting that the speed estimation and trajectory smoothing approach produced stable results. Local 01 was the main exception, where a steep camera angle, congestion and merging maneuvers reduced tracking performance and increased speed variability. Local 02 exhibited the most reliable trajectories, characterized by low outlier rates and high speed autocorrelation, although performance deteriorated under reduced lighting conditions. In contrast, Local 04 showed a higher proportion of low-speed outliers, likely due to vehicles being detected after initiating motion outside the intersection area.

TTC events were identified using a 1.5 s threshold selected empirically. Manual inspection showed that many triggered events resulted from tracking inaccuracies, particularly near image boundaries and during close vehicle interactions that did not involve evasive actions. Consequently, TTC indicators should be interpreted as relative measures of trajectory interaction rather than direct estimates of conflict occurrence.

At Local 01, merging maneuvers and significant speed differentials lead to frequent trajectory interactions, although typically at non-critical speeds. Local 03 is characterized by a roundabout geometry, where intersecting trajectories are inherent to the traffic flow. In contrast, the analyzed segment of Local 02 includes limited conflict points, with only a right-side merging lane; as traffic density decreases and speeds increase, merging behavior becomes smoother. Finally, Location 04 exhibits complete directional flow separation, and no violations of the prescribed traffic patterns were observed.

However, manual video inspection reveals hazardous behavior by motorcyclists across all locations (Porto and Andrade, 2023), which is not adequately captured by the current automated pipeline. In addition, Local 04 presents a significant pedestrian presence. Although vehicular traffic generally complies with signalization and flow separation, pedestrians frequently cross with greater flexibility, often disregarding traffic signals and relying on subjective gap acceptance. With improved tracking calibration and higher-quality footage, a conflict analysis focused on pedestrian interactions could provide insights into the adequacy of signal timing, as well as highlight the risks associated with pedestrian exposure to high-speed vehicular traffic.

4. Discussion

The results indicate that the performance of state-of-the-art algorithms remains highly dependent on input data quality. In this case study, the classification of motorized road users by type (car, motorcycle, bus, truck) represents an improvement over previous assessments conducted at the same locations. However, persistent tracking issues, particularly in the detection and identification of motorcycles, highlight important limitations. These shortcomings indicate that additional refinements are necessary before the proposed conflict detection framework can be reliably applied in operational settings.

For Locations 02 and 03, the proposed YOLO-based pipeline increased the proportion of correctly tracked trajectories from approximately 75% (Porto and Andrade, 2023) to over 80%. The improvement was particularly evident at Local 03, where static occlusions caused frequent trajectory fragmentation in the previous workflow. By providing more robust object detections under partial occlusion, YOLO reduced track interruptions and improved trajectory continuity.

In one of the study locations, Local 02, there was a high success rate in the identification and tracking of road users. For an even better assessment, some improvements can be considered: (i) performing detection and analysis at the lane level, which would enable a more detailed characterization of traffic behaviour; (ii) incorporating colour histogram features normalized by background pixels to enhance object re-identification - preliminary tests using raw bounding box histograms were inconclusive, but background removal may provide more distinctive object signatures without requiring computationally intensive segmentation; and (iii) further exploration of feature representations from intermediate layers of the detection network for object matching. In this study, features from the 22nd and 23rd layers of the YOLO model were evaluated using cosine similarity but did not provide sufficient discriminative power, leaving centroid-based distance as the primary matching criterion. Nonetheless, prior work suggests that intermediate features can support tracking, and a more systematic evaluation may identify suitable representations.

Finally, manual inspection of the video data revealed frequent erratic behaviour among motorcyclists and pedestrians. As vulnerable road users, they are particularly exposed to risk when disregarding traffic signals or designated movement patterns. Traffic conflict analysis can therefore serve not only as a monitoring tool but also as a basis for targeted educational interventions, by quantifying proximity to hazardous situations and raising awareness of risk exposure. Those road users are, however, precisely the most difficult ones to track using automated techniques, indicating an important challenge for future work.

5. Conclusion

This study evaluated the feasibility of using pre-existing HD-CCTV systems for traffic conflict assessment. The proposed pipeline employed YOLOv11m for object detection, focusing on four classes of motorized road users (bus, car, motorcycle, and truck). Trajectories were estimated using an optical flow approach implemented in OpenCV, while object matching between frames relied primarily on bounding box centroid proximity, despite exploratory use of colour histograms and intermediate network features to enhance tracking performance. Tracking results were heterogeneous across locations, with the lowest performance observed at Local 01 and the highest at Local 02. Improved performance was associated with more favourable camera configurations, particularly shorter distances and less oblique viewing angles, which reduce image distortion. Among the limitations discussed throughout the paper, the adoption of fixed calibration parameters and a uniform TTC threshold across locations may affect the comparability of conflict indicators. Future work should investigate adaptive threshold selection and location-specific calibration procedures to improve the robustness and interpretability of conflict-based safety metrics.

Overall, the proposed pipeline represents an advancement over previous automated monitoring efforts conducted at the same sites, although further improvements are required, especially in the detection and tracking of motorcyclists and pedestrians. Regardless, the pipeline was enough to correctly represent interaction among vehicles as a function of road geometry and traffic flow control. In addition to the implementation of the processing pipeline and the provision of tracking outputs, this study contributes new ground truth data for object detection at Location 04 (PS), supporting future benchmarking and comparative analyses.

Acknowledgements

This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101119590.

References

Abdel-Aty, M., Wang, Z., Zheng, O., Abdelraouf, A. (2023). Advances and applications of computer vision techniques in vehicle trajectory generation and surrogate traffic safety indicators. Accident Analysis and Prevention, 191, 107191. https://doi.org/10.1016/j.aap.2023.107191.

Allen, B. L., Shin, B. T., & Cooper, P. J. (1978). Analysis of traffic conflicts and collisions. Transportation Research Record, 667, 67–74.

Archer, J. (2005). Indicators for traffic safety assessment and prediction and their application in micro-simulation modelling. Royal Institute of Technology.

Coifman, B., Beymer, D., McLauchlan, P., & Malik, J. (1998). A real-time computer vision system for vehicle tracking and traffic surveillance. Transportation Research Part C, 6(4), 271–288.

Elvik, R. (2008). A survey of operational definitions of hazardous road locations in some European countries. Accident Analysis & Prevention, 40(6), 1830–1835.

Gettman, D., & Head, L. (2003). Surrogate safety measures from traffic simulation models. Transportation Research Record, 1840, 104–115.

Ismail, K., Sayed, T., Saunier, N., & Lim, C. (2009). Automated analysis of pedestrian–vehicle conflicts using video data. Transportation Research Record.

Laureshyn, A., Svensson, Å., & Hydén, C. (2017). Evaluation of traffic safety based on micro-level behavioural data: The traffic conflict technique. Accident Analysis & Prevention, 102, 1–10.

Lord, D., & Mannering, F. (2010). The statistical analysis of crash-frequency data. Transportation Research Part A, 44(5), 291–305.

Porto, J. A. & Andrade, M. (2023). Conflict Data Assessment: A Case Study in the Federal District of Brazil, In: XXII PANAM GUAYAQUIL: XXII Pan-American Conference on Transportation Engineering and Logistics 2023, Guayaquil/Ecuador.

Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You Only Look Once: Unified, real-time object detection. CVPR.

Saunier, N., & Sayed, T. (2010). Automated analysis of road safety using video data. Transportation Research Record.

Sayed, T., & Zein, S. (2018). Traffic conflict standards for intersections. Transportation Research Record.

Tarko, A., Davis, G., Saunier, N., Sayed, T., & Washington, S. (2009). Surrogate measures of safety. Transportation Research Board.

Ventura, R., Roussou, S., Ziakopoulos, A., Barabino, B., Yannis, G. (2025). Using computer vision and street-level videos for pedestrian-vehicle tracking and behaviour analysis. Transportation Research Interdisciplinary Perspectives, Volume 30, 101366. https://doi: 10.1016/j.trip.2025.101366.

Wojke, N., Bewley, A., & Paulus, D. (2017). Simple online and realtime tracking with a deep association metric. ICIP.

Yannis, G., Papadimitriou, E., & Folla, K. (2014). Effect of GDP changes on road traffic fatalities. Safety Science, 63, 42–49.

Zheng, L., Ismail, K., & Meng, X. (2019). Traffic conflict techniques for road safety analysis: Open questions and some insights. Accident Analysis & Prevention, 132.