



Road Safety and Simulation 2026 – RSS2026

Are telematics-based harsh occurrences associated with street-level visual features? A case study of motorway intersections.

Julia Alves Porto^{a*}, Apostolos Ziakopoulos^a, John Avgeros^b, Daniel Felipe Lopez^c,
George Yannis^a

^aNational Technical University of Athens, Department of Transportation Planning and Engineering,
5 Iroon Polytechniou Street, 15773, Athens, Greece

^bOSeven Single Member Private Company, 27B Chaimanta Street, 15234, Chalandri, Greece

^cFRED Engineering, 31 Via Treviso, RM-00161, Rome, Italy

Abstract

Road intersections are well-known high-risk locations within transportation networks due to their susceptibility to traffic conflicts. This study investigates whether elements visible in street-level imagery are associated with intersection hazard levels, using telematics-based surrogate safety measures: harsh braking and harsh acceleration events. In this exploratory study, telematics data from 322 trips conducted nearby the city of Gorizia, Italy, were analyzed. A total of 98 motorway junctions were selected based on two criteria: (i) at least one trip crossed the intersection, and (ii) at least one connected road was classified as a motorway or primary road. The number of harsh events within a 50-meter buffer around each junction was selected as the target variable. Street-level imagery was sourced from the crowdsourced platform Mapillary, and image features were segmented using a DeepLabV3+ model trained with Cityscapes data. The Spearman test showed weak correlation between number of harsh events and terrain pixels, and no correlation at all when corrected by the Benjamini-Horchberg false detection rate. However, the log-transformed street variables provided a good fit with a Poisson regression model. Street variables were aggregated into five types: construction, flat, poles and signs, road users and nature, whereas the log-ratio was calculated relative to road pixels, present in all images. Flat pixels showed the most statistical significance amongst the environmental attributes, consistent with the correlation results that showed terrain pixels having the smallest p-value. Adding mean speed and number of trips as explanatory variables as opposed to exposure increased the goodness of fit of the model, resulting to a pseudo-R² of 0.57. Due to the limited number of data points, these findings have limited generalization. However, the proposed framework offers a valuable and transferable approach that can be applied to larger datasets in future image-based research.

© 2026 Porto, J.A., Ziakopoulos, A., Avgeros, J., Lopez, D. F., Yannis, G.. Published by ELSEVIER B.V.

This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0>)

Peer-review under responsibility of the scientific committee of the Road Safety and Simulation 2026 – RSS2026

Keywords: intersection; telematics; street-level imagery; Spearman; Poisson; road safety

* Corresponding author. Tel.: +30-694-996-1497

E-mail address: julia_porto@mail.ntua.gr

1. Introduction

Road safety remains an important research topic, as road crashes are still the number one cause of death for young people aged from 5 to 29 years old and cause considerable human and material toll globally (WHO, 2023). Throughout road networks, intersections are known to be hazardous places due to the natural interactions among road users crossing the road in different directions. That intrinsic characteristic has led to several safety research papers targeting those areas specifically (Yue, 2024; Zheng & Sayed, 2020; Zhou et al., 2011).

Historically, safety level measurement has relied on crash data (Nikolaou et al., 2023), but on recent years, the use of surrogate safety measurements instead has gained increasing attention in the academic community (Tarko, 2018). Surrogate safety measures are non-crash occurrences that are correlated with crash occurrence, either by sharing the same causal events or by having proven statistical close relationship with crash occurrence. A major data source that can indicate the risk of a crash occurrence is road telematics (Boylan et al., 2024), which comprises data referring to geo-positioning and monitoring of road users over time, usually by assessing mobile phone or in-vehicle sensors.

In parallel, technological advancements have reshaped how safety research is done, relying more on open access, big datasets and computational tools. One type of data increasingly used for safety assessment is street-view imagery (SVI) data. Not only is there near-global coverage provided by Google, Baidu Maps and other open sources like Mapillary, but there are also some open datasets with labeled data for Computer Vision (CV) tasks such as semantic segmentation (Neuhold et al., 2017; Cordts et al., 2016; Zhou et al., 2017). Some behavior patterns based on SVI attributes have begun to emerge, but still further research is required to allow for generalized conclusions. For example, the simultaneous presence of high quantities of road and sky pixels have been associated with the increase of fatal crashes (Fan & Loo, 2025; Stiles et al., 2022) and with a reduction of crash risk for cyclists (Jeon & Woo, 2024), highlighting that the same environment can relate to traffic behavior in complex ways.

The aim of this research was to explore whether SVI-derived attributes can be used to estimate risk level based on telematics data, more specifically, the number of harsh braking and harsh acceleration events. To the best of the authors' knowledge, this combination itself consists of a novel approach, meriting investigation.

2. Data and methods

2.1. Data collection and processing

Telematics data was provided by OSeven (www.oseven.io), a Greek company that provides smartphone-based telematics solutions. OSeven has developed a seamless integration platform for collecting and transferring raw data from smartphones' sensors. The platform includes a classifier able to flag unusual behaviour, such as harsh braking/acceleration, based on data fusion and machine learning algorithms that interpret values of the accelerometer and other phone sensors (e.g., GPS, gyroscope), without relying on specific thresholds (Kontaxi et al., 2021). The collected data consists of 322 trips between March 2021 and March 2025 on the region surrounding the city of Gorizia, close to the border between Italy and Slovenia. Per each trip, the per-second behaviour is provided, including geolocation, speed, harsh event occurrence and intensity of harsh event. Abiding to General Data Protection Restriction law, data was completely anonymized before being shared for research purposes.

Network data in the form of nodes and edges was collected from OpenStreetMap (OSM), through the OSMnx library, in Python. It was observed that most harsh events were concentrated around intersections, thus consistent with the well documented higher risk in such locations (Kwon & Cho, 2020; Yue, 2024; Tay & Rifaat, 2007), and providing an additional incentive to limit the scope of the present research to intersections. Other than that, to keep the covered data points consistent with each other, we further restricted the dataset to “motorway junctions”, i.e., intersections where at least one of the roads involved is classified as a primary road or motorway, and to intersections with at least two trips, in order to avoid bias towards a specific person's behaviour.

To relate the telematics data to the junction, a 50-meter buffer was created around each junction and grouped by the trip id number. Although in dense networks there is a risk that, using a simple buffer to connect telematics data to an intersection may lead to a wrong classification without considering the edge origin of the trip (Paradiso et al., 2024), the rural context of the present data resulted in intersections being sufficiently sparse that no overlaps occurred within a 50-m buffer, allowing the use of a simple buffer. Figure 1 shows the selected 98 data points.




Fig. 1. Network and telematics data layout. Red dots represent trip data points, black dots represent the motorway junctions that obey the described selection criteria, grey continuous lines represent the roads, where thicker lines indicate motorways or primary roads, and the dashed lines represent the Italian and Slovenian borders. The North points up and the South border is the Mediterranean Sea.

To gather environmental attribute data, we relied on the Mapillary database, which is a crowd-sourced platform with street-level imagery all around the globe, with a CC-BY-SA (Creative Commons Attribution-ShareAlike) license; therefore, it is reproducible for research without restrictions under registration. To correctly fetch the available images, the junction dataset was divided into a grid compatible with the default configurations of `get_image_close_to` function from the Mapillary SDK package for Python, and all images' IDs within the correspondent grid to each of our junction points were selected. Next, a spatial join was performed to retrieve only the images within a 50-m radius from the junction coordinates, consistent with the geofenced area regarding the telematics data. Mapillary has uneven imagery distribution around the world, but it has a very high coverage percentage in Europe, and at least one image was successfully retrieved for each junction point. On total, 101 images were fetched (3 junctions have 2 images within the 50-m buffer).

The images were then processed with the DeepLabV3+ (Chen et al., 2018) semantic segmentation model, available on Segmentation Models Pytorch library (Iakubovskii, 2019). DeepLabV3+ is a Deep Neural Network model that relies on encoder-decoder structure. The ResNet-50 was used as an encoder. For the decoder, default 5-layer configuration was used. The DeepLabV3+ decoder relies on atrous (dilated) separable convolutions, which combine depthwise separable convolution with atrous convolution to improve computational efficiency while preserving spatial information. Training was performed locally on NVIDIA GeForce RTX-2080 GPU, Cuda version 12.6, with the Cityscapes data (Cordts et al., 2016). 19 different classes were monitored using three metrics (accuracy, Intersection of Union – IoU and F1-Score), both overall and per class, and training was performed during 58 epochs before early stop was triggered after 10 epochs elapsing with no IoU improvement on the validation dataset. The model reached an average IoU of 0.5881 on the best epoch. Every one of the images was then processed with the trained model. For the junction points with two images, the average pixel rate per class was used in the next steps.

2.2. Correlation and regression

The first round of analysis searched for the presence of a linear relationship between the environmental elements and telematics-based events. Since the harsh event occurrence follows a non-normal distribution, Spearman Correlation was found more suitable, as it is more robust to non-normality. The same correlation test was also employed for autocorrelation assessment between the environmental variables. The Spearman correlation test is a ranked correlation test, where the strength of the correlation, ranging from -1 to 1, is given by their rank positions along a dataset. The correlation tests were performed using the SciPy library, in Python (Virtanen et al., 2020).

In the second round of analysis, the environmental data was used to estimate the number of harsh event occurrences through a regression count model. Two models were tested: Poisson (Equation 1) and Negative Binomial (NB) (Equation 2). While Poisson assumes that the variance is equal to the mean, NB admits the existence of overdispersion, regulated by a dispersion parameter (α) that dictates the difference between the mean (μ) and the variance (Equation 3). When α is approximately 0, NB collapses to Poisson distribution.

$$P(Y_i = y_i | \mu_i) = \frac{\mu_i^{y_i} e^{-\mu_i}}{y_i!} \tag{1}$$

$$P(Y_i = y_i | \mu_i, \alpha) = \frac{\Gamma(y_i + 1/\alpha)}{\Gamma(1/\alpha) y_i!} \left(\frac{1}{1 + \alpha \mu_i} \right)^{1/\alpha} \left(\frac{\alpha \mu_i}{1 + \alpha \mu_i} \right)^{y_i} \tag{2}$$

$$Var(Y_i | X_i) = \mu_i + \alpha \mu_i^2 \tag{3}$$

The value of μ is estimated using the provided variables (X_i) through maximum likelihood. For that, the Statsmodels package for Python was used. The model output includes the coefficient β , confidence interval and p-value for each variable.

To perform variable selection, some considerations had to be made. Since the semantic segmentation pixel data is compositional, that is, they belong to the same dataset and sum up to 1 (or 100), the regression model crashes out when given the raw variable values. To work around this problem, the arithmetic log-ratio (ALR) rate with reference to the road pixels was provided as explanatory variables rather than the raw values itself, when working with all listed variables. Finally, some of the telematics variables were also incorporated into the models. First and foremost, the number of trips for each junction in our dataset was considered as an exposure for the number of harsh-event occurrences, and its log was incorporated as an offset to the model. Alternatively, we explored incorporating the number of trips not as exposure, but as a variable, as well as the average speed on the junction.

3. Results and Discussion

Figure 2 and Table 1 portray, respectively, the autocorrelation results amongst the environmental attributes and their correlation to number of harsh events, regressed by number of trips.

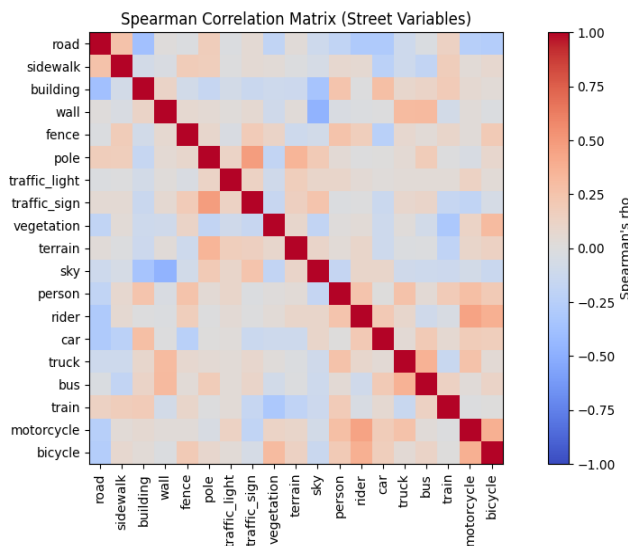


Fig. 2. Autocorrelation results

Table 1. Correlation results

Variable	Spearman ρ	p-value	p-value (BH)
Road	0.1770	0.0813	0.5147
Sidewalk	0.1241	0.2235	0.7121
Building	-0.0345	0.7356	0.8735
Wall	0.0395	0.6995	0.8735
Fence	-0.0014	0.9893	0.9893
Pole	-0.0382	0.7091	0.8735
Traffic light	0.0906	0.3748	0.7121
Traffic sign	-0.0043	0.9662	0.9893
Vegetation	0.1092	0.2844	0.7121
Terrain	-0.2339	0.0204	0.3885
Sky	0.0146	0.8868	0.9893
Person	0.0630	0.5379	0.8735
Rider	-0.0477	0.6413	0.8735
Car	-0.1790	0.0778	0.5147
Truck	0.1077	0.2913	0.7121
Bus	-0.1058	0.2999	0.7121
Train	-0.0532	0.6030	0.8735
Motorcycle	-0.0913	0.3711	0.7121
Bicycle	-0.1239	0.2240	0.7121

As expected, there is some trade-off between categories that is shown by the autocorrelation. For example, when the image presents a higher percentage of sky, there is a lower frequency of wall and fence, and vice versa. Also, when motorcycles or bicycles are present, the model tends to also identify riders. This level of correlation was expected given the nature of the variables and the kind of data. Nevertheless, this correlation does not show strict co-occurrence patterns, which might appear if a magnitude higher than 0.8 or 0.9 was identified. Therefore, no variables were dropped for the correlation tests.

As for the correlation test itself, it evaluates whether there is a monotonic association between two variables. In this case, the non-null hypothesis would be that changes in the proportion of a given street element are associated with changes (increase or decrease) in the number of harsh events. Since the number of trips is an exposure variable and presented a significant level of correlation with number of events ($\rho = 0.30$), it was considered suitable to perform a partial correlation test; that is, instead of correlating the rank values of the variable of interest and of the number of harsh events, we first removed the linear effect of the exposure variable from them. This was done by fitting a linear regression model of each variable on the number of trips rank and extracting the residuals (i.e., the differences between observed and fitted ranks). The correlation was then computed between these residuals.

The Spearman- ρ output did not show any strong correlation for our variables of interest. In fact, the only variable that showed a correlation at the 5% confidence level was terrain, which could signalize that open spaces have a positive effect on driving behavior, reducing the number of harsh-event occurrences. However, when the p-values are corrected through the Benjamini and Hochberg (1995) procedure to account for the higher probability of false positives when multiple tests are performed jointly, the null hypothesis that there is no monotonic association could not be rejected at all. The adjusted p-value equation is given by the column “p-value (BH)” on Table 1.

The results do not indicate that there is no relationship at all, but only that any relationships cannot be interpreted linearly. Therefore, no impediment for other levels of analysis, such as the regression tests, were determined.

For the present dataset, the mean value of number of harsh events is 0.35 and the variance is 1.96, indicative of considerable overdispersion. We compared goodness of fit of three models: Poisson, Negative Binomial with estimated alpha value, and Negative Binomial with alpha value fixed as 1, the default value from Statsmodels. However, when running the NB with estimated alpha values in some trials, the model did not converge, indicating that either the overdispersion is not big enough to justify using NB instead of Poisson, or that the log-transformation was not sufficient to eliminate the dependency caused by the compositional nature of the data. We also experimented with different composition of explanatory variables: using all ALR variables with reference to the road pixels and number of trips as an exposure offset; using ALR aggregated variables according to their types (Table 2) and number of trips as an exposure offset; using ALR aggregated variables and number of trips directly; and using ALR aggregated variables, number of trips and the mean speed.

Table 2. Data aggregation for regression modelling

Component variables	Aggregated variable
Building, wall, fence	Construction
Sidewalk, terrain	Flat
Pole, traffic light, traffic sign	Poles and signs
Vegetation, sky	Nature
Person, rider, car, bus, train, truck, motorcycle, bicycle	Road users

To evaluate model fit, we relied on three criteria: the log-likelihood (LL), the corrected Akaike Information Criterion (AICc), and the Bayesian Information Criterion (BIC). The log-likelihood measures how well the model explains the observed data; higher values (i.e., less negative in count models) indicate better fit. However, because the log-likelihood generally increases as more parameters are added, it does not penalize model complexity and therefore should not be used alone for model comparison. Table 3 indicates the results for each of our models.

Table 3. Model fit comparison

Model	Variables used	Alpha	AICc	BIC	Log-likelihood	Significant / total variables
Poisson	All street variables	0	128.3420	-318.9659	-41.2976	4/18
Poisson	Aggregated variables	0	112.7693	-360.7311	-50.1672	3/5
Poisson	With trips	0	111.4521	-359.6879	-48.3938	3/6
Poisson	With trips and speed	0	105.0355	-363.7935	-44.0511	4/7
NB	All street variables	1	136.4259	-327.9652	-45.3396	2/18
NB	Aggregated variables	5.855*10 ⁻⁶	114.7693	132.4293	-50.1672	3/5
NB	Aggregated variables	1	116.8745	-373.8093	-52.2198	2/5
NB	With trips	1.623*10 ⁻⁸	113.4521	133.4725	-48.3939	3/6
NB	With trips	1	115.5369	-372.7864	-50.4388	3/6
NB	With trips and speed	1	107.3726	-378.6397	-45.2200	2/7

AIC was our primary indicator of goodness of fit, since BIC tends to over-penalize model complexity. Besides, we took into account how well the selected variables explain the traffic behavior. A smaller absolute value of both AICc and BIC indicate a better fit.

First, it is clear that Poisson models provided a better fit than Negative Binomial ones. Although preliminary inspection of the target variable's mean and variance suggested substantial overdispersion, the model-based dispersion statistic (Pearson chi-square over residual degrees of freedom) indicated only moderate overdispersion, with values ranging from 0.85 to 2.34 depending on the selected variables. This is further supported by the low estimated values of the α parameter, as well as convergence issues encountered in some NB estimations. Therefore, these results strongly suggest that Poisson provides a better fit for the distribution of harsh-event occurrences in the present dataset.

Second, considering the number of trips directly an explanatory variable rather than an offset slightly improved model fit, and even more so when combined with average speed information. The main difference between these approaches is that, as an offset, the exposure variable's coefficient is fixed as 1. However, introducing those two variables reduced the statistical significance of the SVI-derived variables. This finding is consistent with Stiles et al. (2022), who reported that segmentation-derived variables tend to play a secondary role in predicting injury and fatal crashes when pixel rates are included directly to the model. Moreover, the inclusion of an offset implies its direct entanglement with the rest of the explanatory variables, which might not always be the case.

Table 4 shows the variable coefficient, p-value and confidence interval for the two Poisson models (highlighted in Table 3) with better fit considering number of trips either as an offset or as a variable. Additionally, it presents the pseudo-R² value and AICc for both models.

Table 4. Coefficient and pseudo-R² values for the best-fitting models

Model	Number trips as offset			Number of trips and speed as variables		
	β	p-value	CI	β	p-value	CI
Constant	-7.3278	0.000	[-9.093, -5.562]	-1.8787	0.199	[-4.748, 0.991]
Construction	0.3165	0.049	[0.001, 0.632]	0.2894	0.168	[-0.122, 0.701]
Flat	-0.2593	0.000	[-0.353, -0.165]	-0.1549	0.010	[-0.273, -0.036]
Poles and signs	-0.0598	0.475	[-0.224, 0.104]	-0.1098	0.217	[-0.284, 0.064]
Nature	0.6816	0.117	[-0.172, 1.535]	1.1323	0.027	[0.126, 2.139]
Road users	-0.4889	0.005	[-0.833, -0.145]	-0.2778	0.156	[-0.662, 0.106]
Number of Trips				0.0222	0.000	[0.013, 0.032]
Mean speed				-0.0294	0.003	[-0.048, -0.010]
Pseudo-R ²	0.3730			0.5431		
AICc	112.7693			105.0355		

As can be seen by Table 4, adding number of trips and mean speed increases the predictive power of the model. However, relying solely on information of the street environment derived from semantic segmentation already provides a reasonable fit for predicting harsh events (pseudo-R² of 0.373).

Lastly, some patterns were observed throughout the model comparison regarding the street variables: flat variables, particularly terrain, showed negative coefficients significant at the 95% confidence level in most model specifications, whereas poles and signs exhibited very high p-values. This suggests that higher proportions of open or unobstructed space are associated with a lower expected number of harsh events. For poles and signs, the limited variance in pixel proportions, especially when expressed relative to road pixels, likely reduced their statistical significance within the models.

As stated in the previous Section, the street variables were not included as raw percentages, but as log-ratios relative to road pixels. For instance, for flat pixels (terrain and sidewalk), if they double their relative percentage to road pixels, indicating narrower road presence, there is an expected decrease of 10% to 17% in harsh-event occurrence, while doubling the proportion of sky and vegetation pixels relative to road pixels is associated with an expected increase of 80% to 120% on the same events. Besides, although the magnitude and statistical significance of the coefficients vary across model specifications, their signs remain largely consistent, suggesting at least a weak but stable association between environmental attributes and harsh-event occurrence. Road users and poles and signs were statistically significant only in a limited number of models. Given the discrete nature of these elements, an object detection approach or a discrete representation of their occurrence may provide a more appropriate modelling strategy for these attributes.

Finally, the small coefficient but strong statistical confidence regarding the telematics variables indicates that they are indeed more suitable as explanatory variables than as direct exposure metrics, as used in previous research (Ziakopoulos, 2024).

4. Conclusion

This research explored the relationship between environmental attributes and dangerous behaviour as a case study focused on junction motorways nearby the city of Gorizia, Italy. Traffic behaviour was derived from telematics data provided by OSeven, encompassing over 300 trips taken in the area between 2021 and 2025. Besides the date, anonymized trip identification and geolocation, the dataset contains speed and harsh-event (harsh braking or harsh acceleration) information. The total number of harsh events per motorway junction was chosen as our variable of interest, where a motorway junction is defined as so by OSM data. A 50-meter buffer boundary was set to identify each individual junction's zone, summed up the number of harsh events identified in each of them and calculated the average speed, required that at least two trips had crossed the buffer area. Environmental attributes were derived from street-level imagery from the Mapillary database. The DeepLabV3+ model was trained to perform semantic segmentation on the standard Cityscapes categories, achieving 0.588 mIoU.

Correlation test results did not dismiss confidently the null hypothesis that there is no direct correlation between individual environmental attribute and harsh-event occurrence. However, they were able to achieve a reasonable fit for count regression models. Although the variance is meaningfully above the mean, Poisson distribution still presented a better fit than Negative Binomial. To achieve meaningful explanatory power, the variables were aggregated into five categories per type: construction, flat, nature, poles and signs, and road users. The ALR transformation relative to road pixels was performed to allow modelling with compositional data. Finally, a greater prevalence of flat (terrain and sidewalk) pixels over road consistently has a negative effect on number of harsh events, which reinforced the weak (and negative) correlation found between terrain pixels and harsh-event occurrence with the Spearman test. Adding number of trips and mean speed as explanatory variables also increased the model predictability strength, achieving a pseudo- R^2 of 0.54, while relying solely on street variables still achieves a reasonable fit of 0.37.

This research has some limitations. First, the dataset is considerably small: only 98 motorway junctions, and some of them have very few trip occurrences, which vary from 3 to 127 per junction. A bigger collection of datapoints would provide more reliable results capable of being directly used by stakeholders to understand what aspects of the environment have most influence on traffic behaviour. Secondly, semantic segmentation models have been reported with higher metrics than the one used in this work. Revising the parameters and the classes used for training and training the model to recognize the ego-vehicle, could improve the results and reliability on identified classes. Future work should try to address those two items as well as include either object detection model for a count-based identification of poles, signs and road users. The model could also be improved by incorporating other explanatory variables, namely intersection type, road classification, demographic characteristics and traffic volume.

Acknowledgements

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

References

- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, 57(1), 289-300.
- Boylan, J., Meyer, D., & Chen, W. S. (2024). A systematic review of the use of in-vehicle telematics in monitoring driving behaviours. *Accident Analysis & Prevention*, 199, 107519. <https://doi.org/10.1016/j.aap.2024.107519>
- Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., & Adam, H. (2018). Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation (arXiv:1802.02611). arXiv. <https://doi.org/10.48550/arXiv.1802.02611>
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B. (2016). The Cityscapes Dataset for Semantic Urban Scene Understanding. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
- Fan, Z., & Loo, B. P. Y. (2025). Urban visual clusters and road transport fatalities: A global city-level image analysis. *Communications in Transportation Research*, 5, 100193. <https://doi.org/10.1016/j.commtr.2025.100193>
- Iakubovskii, P. (2019). Segmentation Models Pytorch. In GitHub repository. GitHub. https://github.com/qubvel/segmentation_models.pytorch
- Jeon, J., & Woo, A. (2024). The effects of built environments on bicycle accidents around bike-sharing program stations using street view images and deep learning techniques: The moderating role of streetscape features. *Journal of Transport Geography*, 121. Scopus. <https://doi.org/10.1016/j.jtrangeo.2024.103992>
- Kontaxi, A., Ziakopoulos, A., & Yannis, G. (2021). Trip characteristics impact on the frequency of harsh events recorded via smartphone sensors. *IATSS Research*, 45(4), 574–583. <https://doi.org/10.1016/j.iatssr.2021.07.004>
- Kwon, J.-H., & Cho, G.-H. (2020). An examination of the intersection environment associated with perceived crash risk among school-aged children: Using street-level imagery and computer vision. *Accident Analysis and Prevention*, 146. Scopus. <https://doi.org/10.1016/j.aap.2020.105716>
- Neuhold, G., Ollmann, T., Bulò, S.R., Kotschieder, P. (2017). The Mapillary Vistas Dataset for Semantic Understanding of Street Scenes. In: Proceedings of the IEEE International Conference on Computer Vision, October 2017, 5000-5009. <https://doi.org/10.1109/ICCV.2017.534>
- Nikolaou, D., Ziakopoulos, A., & Yannis, G. (2023). A Review of Surrogate Safety Measures Uses in Historical Crash Investigations. *Sustainability*, 15(9), 7580. <https://doi.org/10.3390/su15097580>
- Paradiso, S., Ziakopoulos, A., Yannis, G. (2025). Aggregating telematics data for road safety analysis. Paper presented at the 12th International Congress on Transportation Research (ICTR), Thessaloniki, Greece. Available at: <https://www.nrso.ntua.gr/geyannis/pub/pc602-aggregating-telematics-data-for-road-safety-analysis/>.
- Stiles, J., Li, Y., & Miller, H. J. (2022). How does street space influence crash frequency? An analysis using segmented street view imagery. *Environment and Planning B: Urban Analytics and City Science*, 49(9), 2467–2483. Scopus. <https://doi.org/10.1177/23998083221090962>
- Tarko, A. P. (2018). Surrogate measures of safety. *Transport and Sustainability*, 11, 383–405. doi:10.1108/S2044-99412018000011019.
- Tay, R., & Rifaat, S. M. (2007). Factors contributing to the severity of intersection crashes. *Journal of Advanced Transportation*, 41(3), 245-265.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., ... SciPy 1.0 Contributors. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>.
- World Health Organization - WHO (2023). Global Status Report on Road Safety 2023 (1st ed).
- Yue, H. (2024). Investigating the influence of streetscape environmental characteristics on pedestrian crashes at intersections using street view images and explainable machine learning. *Accident Analysis and Prevention*, 205. Scopus. <https://doi.org/10.1016/j.aap.2024.107693>
- Zheng, L., & Sayed, T. (2020). A novel approach for real time crash prediction at signalized intersections. *Transportation Research Part C: Emerging Technologies*, 117. Scopus. <https://doi.org/10.1016/j.trc.2020.102683>
- Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralla, A. (2017). Scene Parsing through ADE20K Dataset. *Computer Vision and Pattern Recognition (CVPR)*.
- Zhou, S., Sun, J., An, X., & Li, K. (2011). The development of a conflict hazardous assessment model for evaluating urban intersection safety. *Transport*, 26(2), 216–223. Scopus. <https://doi.org/10.3846/16484142.2011.589494>
- Ziakopoulos, A. (2024). Analysis of harsh braking and harsh acceleration occurrence via explainable imbalanced machine learning using high-resolution smartphone telematics and traffic data. *Accident Analysis and Prevention*, 207. Scopus. <https://doi.org/10.1016/j.aap.2024.107743>