

A Crowd Sensing Approach to Video Classification of Traffic Accident Hotspots

Conference Paper**Author(s):**

Gahr, Bernhard; Ryder, Benjamin; Dahlinger, André; Wortmann, Felix

Publication date:

2018-07

Permanent link:

<https://doi.org/10.3929/ethz-b-000297283>

Rights / license:

In Copyright - Non-Commercial Use Permitted

Originally published in:

Lecture Notes in Computer Science 10935, https://doi.org/10.1007/978-3-319-96133-0_14

A Crowd Sensing Approach to Video Classification of Traffic Accident Hotspots

Bernhard Gahr¹, Benjamin Ryder², André Dahlinger¹, and Felix Wortmann¹

¹ University of St. Gallen, St. Gallen 9000, Switzerland

² ETH Zurich, Zurich 8000, Switzerland

Abstract. Despite various initiatives over the recent years, the number of traffic accidents has been steadily increasing and has reached over 1.2 million fatalities per year world wide. Recent research has highlighted the positive effects that come from educating drivers about accident hotspots, for example, through in-vehicle warnings of upcoming dangerous areas. Further, it has been shown that there exists a spatial correlation between to locations of heavy braking events and historical accidents. This indicates that emerging accident hotspots can be identified from a high rate of heavy braking, and countermeasures deployed in order to prevent accidents before they appear. In order to contextualize and classify historic accident hotspots and locations of current dangerous driving maneuvers, the research at hand introduces a crowd sensing system collecting vehicle and video data. This system was tested in a naturalistic driving study of 40 vehicles for two months, collecting over 140,000km of driving data and 36,000 videos of various traffic situations. The exploratory results show that through applying data mining approaches it is possible to describe these situations and determine information regarding the involved traffic participants, main causes and location features. This enables accurate insights into the road network, and can help inform both drivers and authorities.

1 Introduction

In recent years the number of road fatalities have been steadily increasing to over 1.2 million deaths per year [35]. This increase is not limited to developing countries with poor infrastructure, for example, the USA saw a 5.6% increase in the number of deaths from traffic accident from 2015 to 2016, and 8.4% from 2014 to 2015 [27]. Similarly, the number of road deaths in the European zone has increased in 2015 and 2016, widening the gap to the goal of reducing road fatalities from 2010 by 50% until 2020 [1]. Besides the well known reasons of driver distraction, i.e., driving under influence of alcohol or drugs, speeding, and other reckless driving behavior, many external factors can lead to traffic accidents. These external factors, such as poor road management, infrastructure challenges and weather conditions, can develop into locations with a high likelihood of traffic accidents, otherwise known as ‘accident hotspots’ [15, 30]. Further, vehicle collisions with wildlife might not have such a high impact on

human mortality, however damages related to these accidents are estimated to have an annual cost of over \$8.3 Billion in the USA alone [9].

Since incorrect driver behavior has been identified as one of the main reasons for traffic accidents [5], many of the road safety increasing programs focus on educating drivers and finding countermeasures for wrong behavior. Authorities within the EU, or the NHTSA in the USA, hence introduced guidelines on driver education to improve the traffic safety situation [13, 16]. Moreover, several research groups have shown the positive effects that in-vehicle warnings of upcoming accident hotspots have on driver behavior [3, 31, 33]. In these situations, people preferred contextualized warnings of danger rather than simple unspecific warnings [26]. A second approach to reduce traffic accidents involves programs focused on improvements to the road infrastructure, examples of this include EuroRAP [20], RANKERS [11], or EUROTAP [2]. For example, recent research on road network characteristics shows that the highest risk of fatal or severe crashes occur on road networks with a very low intersection density [22], and it was found that safety outcomes improve as the intersection density increases. Additionally, the reorganization of road structures implementing ‘shared spaces’ have shown potential in increasing safety for pedestrians and reducing the number of accidents [24]. Other projects, such as in the Banff National Park, found that the building of animal crossings led to a severe reduction of animal related accidents [8]. Independent of the countermeasure employed at a certain location, a concrete knowledge of the actual cause for the accidents can aid in improving the road infrastructure.

One generalized approach to develop this vital understanding of our road networks includes programs such as the DaCoTa project. This project led to the combination and analysis of accident investigations from 30 European countries [19], and resulted in the Road Safety Knowledge System [12]. However, since this approach relies on historic accidents, it falls short in preventing new upcoming accident hotspots and contextualize them in real-time. Therefore we investigate, whether recent content deprived accident hotspots and current dangerous situations can be detected and classified in a immediate manner. As such, in the paper at hand we introduce a crowd, or community, sensing approach which collects data from the Controller Area Network (CAN-Bus) of vehicle’s on board computer, and enriches this with video data from an on-board smartphone camera. Videos are recorded when these vehicles cross a historically known dangerous location, or performs a dangerous driving maneuver. This way, we tackle two hotspot classification issues at the same time: First, locations which are historically known to be dangerous, but lack contextual information about the main accident cause can be analyzed in a retrospective manner. This community sensing approach means that expensive and high maintenance equipment, like stationary surveillance cameras, do not have to be installed. Second, by recording video sequences from dangerous driving maneuvers, developing dangerous locations can be identified and classified before an accident occurs. The increasing number of dash cameras in recent years shows the potential of identifying accident causes out of video sequences. The approach is hereby based on the idea

that only a subset of all accidents (A) are reported (A_r). Further, we assume that all accidents are a subset of critical brake events (B_c), which are a subset of all brake events (B).

$$A_r \subset A \subset B_c \subset B. \quad (1)$$

In the following chapters, the paper at hand introduces a crowd sensing system, and the associated field study, which resulted in a dataset enabling the automated classification of traffic accident hotspots. The remainder of the paper is structured as follows. In Section 2 related work is given and in Section 3 the system is presented. Further, Section 4 provides a description of a field study setup and the collected data is presented in Section 5. Preliminary results are discussed in Section 6 and the paper concludes with an outlook to the remaining work in Section 7.

2 Related Work

Crowd sensing, sometimes referred to as crowd sourcing, is a relatively new paradigm which typically utilizes mobile devices to passively gather data and provide wide-ranging ubiquitous applications [21]. In previous research it was shown that sensors in smart phones, such as an accelerometer, gyroscope and GPS, are well suited for real-time community sensing tasks, e.g., predicting thermal hotspots by clustering sink- and climb rates of paragliders [34]. Crowd sensing of insights related to the road network has previously been used for detecting traffic anomalies, such as disasters or protests [28], along with identifying infrastructure, e.g., traffic lights and stop signs [18]. More recently, a crowd sourcing approach was analyzed for the detection road anomalies from driving video records [7].

The aim of the research at hand is to classify both accident hotspots and dangerous driving behavior on the road network following a crowd sensing approach. Prior research has demonstrated a similar contextualization task, focusing on traffic scenes using video data from static surveillance cameras [23]. The need for such classification is motivated by recent findings which demonstrate that users prefer contextualized in-vehicle warnings over an ‘all purpose’ general warning sign [26]. A prerequisite of the first task, accident hotspots classification, is identifying the locations of existing accident hotspots where driving videos should be recorded. To achieve this, we build upon the results of previous research in Switzerland, where accident hotspots were identified using police accident reports and the DBSCAN clustering technique [31]. This hotspot dataset provided GPS locations of over 1,600 accident hotspots, along with the three descriptive layers of the involved traffic participating objects (“What”), the predominant cause (“Why”), and the location (“Where”).

The second task, contextualizing video captured during dangerous driving situations, builds upon a large body of existing research on driving event detection. Since the system should collect the dangerous driving maneuvers out of a real-world study, we cannot rely on Time-to-collision (TTC) values found in previous research [17], but rather velocity, deceleration, and jerk (change of

acceleration, $j = \frac{da}{dt}$) measurements. Rather high deceleration values of $0.75\ g$ ($7.35\ m/s^2$) to cover 4% to 0.5% of all braking events in an urban area were found in [14]. More practical, but not rigorously tested, values to classify braking events were given in [4] with decelerations of below $2.0\ m/s^2$ as “low danger”, above $3.9\ m/s^2$ as “high danger”, and everything in between as “medium danger”. These same thresholds were taken for both acceleration and turning. Recent research also considered the jerk-rate as a good indicator for critical driving events [29] and values of $-2\ m/s^3$ were found to be good classifier thresholds when brake event locations were correlated with historical accident locations [32]. Moreover, a sophisticated categorization of dangerous driving has been given by the “jerk feature”, calculated via the standard deviation of the jerk rate within a specific time window over the average jerk rate of the current road type [25]. The thresholds from this study were suggested to be between 0.5 and 1.0.

Summing up, we find that smartphones and driving data from vehicles have previously been utilized, and are well suited, for community sensing applications. Driver behavior is improved when users are provided with in-vehicle warnings of upcoming dangerous locations, and contextualized information in these warnings is preferred. Finally, thresholds for identifying brake events vary significantly across research studies, however, medium and high danger events typically lie around $2\ m/s^2$ for deceleration and $-2\ m/s^3$ for the jerk rate.

3 System Description

In the following section we present the crowd sensing system which was developed to address the challenge of classifying dangerous locations on the road network. The system is mainly comprised of three parts: the car and its associated sensors, the smartphone which relays all data from the vehicle to the server and contributes recoded video sequences and its own sensor information, and the server which stores and processes the data and provides accident hotspot information back to the driver’s smartphone. An overview of the system is shown in Figure 1.




Fig. 1: Overview of the system setup.

The car's CAN-Bus was accessed via a dongle mounted at the OBD-II port and data collected was transferred via Bluetooth to the smartphone. The signals which were queried from the CAN-Bus by the dongle and transmitted to the server are shown in Table 1. These signals are not typically available with a constant frequency on the CAN-Bus, rather they are triggered on a change of the associated sensor. However, many change frequently enough to consider them a continuous signal, and for these we present an approximate frequency.

Table 1: Collected CAN-bus Signals

Signal name	Range/Measure	Approx. Frequency
Wheel Speed {fr, fl ,br ,bl} ³	$\frac{km}{h}$	10 Hz
Wheel Slip Status	p-/ n-/ no- slip ⁴	onChange
ABS Active	on/off	onChange
ABS Failed	on/off	onChange
ESP	on/off	onChange
Engine Speed	rpm	10 Hz
Odometer	km	onChange
Brake Pressure	Pa	30 Hz
Brake Pedal & Throttle Position	%	30 Hz
Brake Pressure Detected	on/off	onChange
Steering Wheel Angle	deg	10 Hz
Gear	p/r/n/d ⁵	onChange
Vehicle Acceleration	$\frac{m}{s^2}$	10 Hz
Long. & Lat. Acceleration	$\frac{m}{s^2}$	10 Hz
Lights	o/p/d/f ⁶	onChange
Head Lights	on/off	onChange
Fuel Level	%	onChange
Fuel Consumption	$\frac{l}{h}$	10 Hz
Fuel Consumption / 100km	$\frac{100km}{l}$	10 Hz
Outside Temperature	Celsius	onChange
Yaw Rate	deg	10 Hz

On the smartphone, the received CAN-Bus signals were further enriched with video sequences, location information and additional sensor data from phone itself. The additional signals generated from the smartphone are presented in Table 2. In order to get an optimum view of the street scene ahead of the car, users of the system were asked to mount the smartphone so that the rear camera was facing out of the windshield onto the street. To better aid in the calibration

³ f: front, b:back, l:left, r:right

⁴ p: positive, n: negative

⁵ p: park, r: reverse, n: neutral, d: drive

⁶ o: off, p: park, d: dimmed, f: full

⁷ Meters above sea level

⁸ Only for calibration, to get the phones orientation

⁹ 2G, 3G, LTE

Table 2: Collected Smart Phone Signals

Signal name	Range/Measure	Approx. Frequency
Bearing	deg	1 Hz
GPS Altitude	Masl ⁷	1 Hz
GPS Latitude	deg	1 Hz
GPS Longitude	deg	1 Hz
GPS Speed	$\frac{m}{s}$	1 Hz
GPS Accuracy	m	1 Hz
Accelerometer	$\{x,y,z\}-\frac{m}{s^2}$	1 Hz
Magnetometer	$\{x,y,z\}-\mu T$	onCalib ⁸
Cellular Network Signal Strength	Type ⁹ , dBm	1 Hz
Video Recordings	Video Sequences	onTrigger

of the smartphone’s field of view, a live camera feed was shown to the user for six seconds on the smartphone’s screen whenever a start of the engine was detected from the CAN-Bus. After the calibration screen the system provided eco-driving feedback, as introduced in [10], together with accident hotspot warnings [31] to the users. The collected data from the CAN-Bus and the smartphone were transmitted via the cellular network to a server, where they were stored in a database. With the latency of the cellular network as the main delay contribution, the system can provide insights in a real-time manner to interested road authorities and potentially other traffic participants.

3.1 Video Recordings

As introduced in Section 1, the aim of the system is to enrich historic accident hotspots with context from captured video sequences, and in addition classify current dangerous driving maneuvers. In order to accurately capture relevant video footage, live data from the smartphone camera was constantly written to and stored in a ring buffer of four seconds length. Whenever a user crossed a historic hotspot location, the buffer together with a further video sequence of three seconds was extracted and sent to the server. With this implementation, we received a video sequence that typically showed the entry and exit each time a historic hotspot was encountered.

To enable to video classification of dangerous driving maneuvers, the video extract was triggered on the detection of these events. Due to the many different definitions of a “dangerous driving behavior” we took the simple, yet practical definition of such events as a magnitude acceleration of greater than 2 m/s² in the horizontal plane [4], i.e.

$$\sqrt{\left(\frac{d^2x(t)}{dt^2}\right)^2 + \left(\frac{d^2y(t)}{dt^2}\right)^2} > 2 \left[\frac{m}{s^2}\right], \quad (2)$$

in other words the L2-norm, whereby $x(t)$ and $y(t)$ denoted the coordinates in the horizontal plane dependent on the time t . Hence, the videos are triggered

on heavy deceleration, acceleration, cornering, or a combination of these. Once the dangerous event is detected, the ring buffer and an additional sequence of three seconds is extracted and transmitted to the server. Similar to the historical hotspots, this results in video sequences which show how the user enters and exits each dangerous situation.

4 Study Description

In order to evaluate the proposed crowd sensing system, a field study was conducted with the Swiss road assistance service “Touring Club Suisse” (TCS). In



Fig. 2: The base locations of the drivers distributed in the German speaking area of Switzerland.

total, 41 TCS drivers, or patrollers, were recruited to take part in a two month naturalistic driving study. All drivers were male, between 23 and 64 years old, and 37 years on average. We installed our system into each of the patrollers’ cars, all same make and model of Chevrolet Captiva 2.2, without assigning them with any specific driving tasks. As such, they were instructed to simply follow their daily work and so did not follow any specific routes nor did they perform any predefined driving maneuvers. These drivers were situated at nine different base locations in the German speaking area of Switzerland with at least three drivers per base location, as shown in Figure 2. The locations covered mountainous (southern area, base locations 3 to 5) as well as hilly to almost flat (northern area) regions. Part of the general employment conditions of the participants are taking driver safety training courses, as well as fuel efficient driving lessons. As such, a safe and “calm” driving behavior was typically expected of the drivers.

The study ran for over two months, including the staging and dismantling phases at the beginning and the end of the study respectively. The number of active drivers per day is given in Figure 3. Disregarding the first and last days, we can see that at least 10 drivers were actively using our system on most days,




Fig. 3: Number of active drivers per day.

with peak days reaching over 25 active drivers. The high fluctuations are a result of week days and the demand of call-outs per day¹⁰.

5 Collected Dataset

Within the two months of the field study over 140,000 kilometers of driving data were collected in total. The accumulated kilometers over the whole fleet is shown in Figure 4. A constant increase in the accumulated kilometers can be seen. In Figure 5 the blue line shows the total amount of driven kilometers per




Fig. 4: Total driven kilometers of the whole fleet.

driver. The driver IDs are ordered by the total distance driven during the field

¹⁰ Beside the normal road rescue tasks, patrollers had stand-by duty.

study. Out of the 40 drivers that signed up for the field study, three never used the system. Another three drivers only marginally used the system, driving less than 1,000 kilometers each. The three most active drivers reached almost 8,000 kilometers within the two and a half months.




Fig. 5: Total kilometers driven and videos recorded per driver.

Furthermore, Figure 5 additionally shows the number of videos recorded per driver, denoted by the orange dashed graph. Within the whole field study over 36,000 video sequences were collected, resulting in over 36 GB of video data. As described in Section 3.1 the videos were triggered during an accident hotspot crossing and on the detection of dangerous driving maneuvers. The number of videos increased more or less proportional to the total kilometers recorded, reaching a maximum of over 2,300 videos for one driver. Importantly, since the patrollers ultimately chose freely the final mounting of the smartphone within the car, two sets of videos were deemed to be unusable. These videos were generated by Driver 19, who placed the phone in such a way, that only the motor toll sticker was recorded, and Driver 31, who left the phone somewhere in the side door without any sight onto the road. Therefore of the 36,341 videos, approximately 34,000 videos are labeled as usable.

6 Preliminary Results

In the following section we present the results of two exploratory classification analyses conducted for the collected historic accident hotspot and dangerous event recorded videos respectively. Hereby, for each task we attempt to generate warning feedback similar to the guidelines introduced by the NHTSA, so that drivers are provided with information about an accident hotspot in form of a warning sign and non-critical supporting text [6]. The hotspot information will be described — whenever possible — by the three categories of location description,

predominant cause, and involved traffic objects similar to the approach in [31]. The available information is hereby ranked by the following descending order of importance: objects, cause, and location. If no specific attribute in any of the category can be found a general warning sign with the text “Attention: Dangerous Location” is assigned to the hotspot.




Fig. 6: Frames of five videos at a historic tunnel accident hotspot.

6.1 Historic Hotspot Contextualization

The videos recorded from existing hotspots can serve as a good indicator for the location contextualization layer describing the scenery. Out of the video sequences, large static objects can be easily identified, for example, a tunnel as shown in Figure 6. Further, as it can be seen in Figure 7, road junctions, crosswalks, bike-lanes and other infrastructure can also be identified.

In addition, the traffic participants, such as pedestrians (c.f. Figure 7 (e)) or cyclists (Figure 7 (f)), can be determined through investigating multiple videos of the same hazardous location, and a dominant object class assigned to the hotspot. In the same way, if the majority of video sequences of a hotspot show a particular traffic situation which could cause accidents, then this can be identified. For example, videos in Figure 7 might lead to a “rear-end collision” categorization due to the repeatedly observed close proximity of vehicles.

6.2 Dangerous Driving Contextualization

As described in Section 3.1, we used a simple yet practical threshold of acceleration magnitude of 2 m/s^2 for dangerous event detection. However, since the system collected the vehicle accelerometer measurements in addition




Fig. 7: Frames of six videos at a historic rear-end collision accident hotspot.

to the video sequences, we can assign the maximum magnitude acceleration and maximum jerk rates a-posteriori to each video. The following Figures will show situations, which are characterized either by a high acceleration value following the definition of [4], or by a high jerk rate, following the definition of [29, 32] as a dangerous situation. Similar to the classification of historic hotspots, the classification of a locations where dangerous driving maneuvers occurred is demonstrable from the recorded video sequences.




Fig. 8: Two frames of a video labeled as a “dangerous” driving maneuver by a high acceleration rate.

The CAN bus recordings for the videos shown in Figure 8 show medium jerk values and high acceleration values. For the sequences in Figure 8, the location attribute can be described as a sharp curve. In Figure 9 a sequence characterized by high jerk rates is shown. Similar to Figure 8, the location is characterized by a sharp curve. However, the driver enters this second curve with an excessive speed, so much so that we observe objects sliding through the driver’s cabin.




Fig. 9: Three frames of a videos labeled as a “dangerous” driving maneuver by a high jerk rate.

Where in the first sequence, a “sporty”, but still controlled type of driving was performed, the cause attribute of the later one can be characterized as heedless driving. As such, the classification of the predominant cause in this situation would lead to a “swerving” or “control speed” warning. Further, weather and light conditions, i.e. entering from the bright street into the dark underpass can lead to uncontrolled maneuvers.

The location shown in Figure 10 is classified as a crossing, where the patroller had the right of way. In this sequence we find that in the object layer another vehicle is identified, and later observe that this car disregards the patroller’s right of way causing him to brake sharply. Hence, on the predominant cause layer, this situation is described as “disregarding the right of way”.




Fig. 10: Frames of a videos characterized by disregarding the right of way.

Finally, in the sequence in Figure 11 the driver had a delayed reaction to the vehicle ahead braking, and had to come to an emergency stop in order avoid a collision. In addition to the classification of the cause, in this case “heedless diving”, the vehicle ahead can be identified in the object layer and leads to a “rear-end collision” contextualization.




Fig. 11: Two frames of a videos characterized by a rear-end collision situation.

7 Conclusion

In this work we introduced a crowd sensing system to tackle the problem of classifying traffic accidents and hotspots, and presented initial exploratory results demonstrating the potential of this approach. In previous work we saw that concrete countermeasures improve the road safety if the correct, i.e. the cause of accident matching counteraction, are taken. In order to identify the correct cause of a hotspot the system was taken into the field and collected a large dataset of CAN-Bus and video recordings with over 140,000 kilometers and over 34,000 video sequences. Our first results indicate that with this data set we are able to identify dangerous hotspots within a large part of Switzerland and classify the main accident causes in order to warn drivers and road authorities with concrete information of these areas.

The preliminary results of this work should be assessed in light of its limitations. Despite the size and uniqueness of the collected dataset, we have to acknowledge that the first results are based on a homogeneous group of experienced drivers. As such, a generalization of the field study data to the normal population of traffic participants should be undertaken with caution, and future field studies should consist of more regular drivers. However, due to the experience and risk averse driving style of the patrollers, we assume that dangerous driving maneuvers occurred less frequently than a more representative group of study participants.

Further, a promising next step to this research, omitted here for brevity, are the first insights from using publicly available image classification models for object identification. This approach shows very promising results, and indicates that automated models can already be used to identify objects and reduce the laborious task of labeling all videos by hand. In future work, the object identification of such videos should be extended towards automated classification of traffic accident hotspots and dangerous driving maneuvers. Implications of video sequences from the three layers used in this work to contextualize a hotspot should be drawn automatically to build an actual real-time system.

References

1. Adminaite, D., Jost, G., Stipdonk, H., Ward, H.: Ranking EU progress on road safety. Tech. rep. (2016), <http://etsc.eu/10th-annual-road-safety-performance-index-pin-report>
2. Allgemeiner Deutscher Automobil Club: EUROTAP (2007), https://ec.europa.eu/transport/road_safety/sites/roadsafety/files/pdf/projects/eurotap.pdf
3. An, P.E., Harris, C.J.: An intelligent driver warning system for vehicle collision avoidance. *IEEE Transactions on Systems, Man, and Cybernetics Part A: Systems and Humans*. 26(2), 254–261 (1996)
4. Bergasa, L.M., Almeria, D., Almazan, J., Yebes, J.J., Arroyo, R.: DriveSafe: An app for alerting inattentive drivers and scoring driving behaviors. In: *IEEE Intelligent Vehicles Symposium, Proceedings*. pp. 240–245 (2014)
5. Bohnenblust, D., Pool, M.: Verkehrsunfälle in der schweiz 2016. Bundesamt für Statistik, BFS (2017), <https://www.bfs.admin.ch/bfs/de/home/statistiken/mobilitaet-verkehr/unfaelle-umweltauswirkungen/verkehrsunfaelle.assetdetail.3103126.html>
6. Campbell, J.L., Richard, C.M., Brown, J.L., McCallum, M.: Crash Warning System Interfaces: Human Factors insights and lessons learned. Tech. rep. (2007)
7. Chen, H.T., Lai, C.Y., Shih, C.A.: Toward community sensing of road anomalies using monocular vision. *IEEE Sensors Journal* 16(8), 2380–2388 (April 2016)
8. Clevenger, A.P., Ford, A.T., Sawaya, M.A.: Banff wildlife crossings project: Integrating science and education in restoring population connectivity across transportation corridors. Final report to Parks Canada Agency, Radium Hot Springs, British Columbia, Canada (June), 165 (2009)
9. Clevenger, T., Cypher, B.L., Ford, A., Huijser, M., Leeson, B.F., Walder, B., Walters, C.: Wildlife-Vehicle Collision Reduction Study. Tech. rep. (2008)
10. Dahlinger, A., Wortmann, F., Tiefenbeck, V., Ryder, B., Gahr, B.: Feld-experiment zur wirksamkeit von konkretem vs. abstraktem eco-driving feedback. In: *Wirtschaftsinformatik Konferenz (WI) 2017* (March 2017), <https://www.alexandria.unisg.ch/250432/>
11. European Commision: Ranking for European Road Safety (2008), https://ec.europa.eu/transport/road_safety/sites/roadsafety/files/pdf/projects/rankers.pdf
12. European Commision: Road safety knowledge system (2008), <http://safetyknowsys.swov.nl>
13. European Commission: A strategic approach to implementing countermeasures (2018), https://ec.europa.eu/transport/road_safety/specialist/knowledge/young/implementation_process/a_strategic_approach_to_implementing_countermeasures.en
14. Glassco, R.A., Cohen, D.S.: Collision avoidance warnings approaching stopped or stopping vehicles. In: *ITS America (ed.) 8th World Congress on Intelligent Transport Systems*. Sydney, Australia (2001)
15. Goniewicz, K., Goniewicz, M., Pawłowski, W., Fiedor, P., Lasota, D.: Road safety in poland: magnitude, causes and injuries. *Wiadomosci lekarskie* (Warsaw, Poland: 1960) 70(2 pt 2), 352–356 (2017)
16. Goodwin, A., Thomas, L., Kirley, B., Hall, W., O'Brien, N., & Hill, K.: Countermeasures That Work: A highway safety countermeasure guide for State highway safety offices, Eighth edition. (Report No. DOT HS 812 202). Washington, DC: National Highway Traffic Safety Administration. (January) (2015)

17. Hogema, J., Janssen, W.: Effects of intelligent cruise control on driving behaviour: a simulator study (jan 1996)
18. Hu, S., Su, L., Liu, H., Wang, H., Abdelzaher, T.F.: Smartroad: Smartphone-based crowd sensing for traffic regulator detection and identification. *ACM Trans. Sen. Netw.* 11(4), 55:1–55:27 (Jul 2015)
19. Jähi, H., Muhlrads, N., Buttler, I., Gitelman, V., Bax, C., Dupont, E., Giustiniani, G., Machata, K., Martensen, H., Papadimitriou, E., Persia, L., Talbot, R., Vallet, G., Yannis, G.: Investigating Road Safety Management Processes in Europe. *Procedia - Social and Behavioral Sciences* 48, 2130–2139 (2012)
20. Lynam, D., Castle, J., Martin, J., Lawson, S.D., Hill, J., Charman, S.: EuroRAP 2005-06 technical update. *Traffic Engineering and Control* 48(11), 477–484 (2007)
21. Ma, H., Zhao, D., Yuan, P.: Opportunities in mobile crowd sensing. *IEEE Communications Magazine* 52(8), 29–35 (Aug 2014)
22. Marshall, W.E., Garrick, N.W.: Street network types and road safety: A study of 24 California cities. *Urban Design International* 15(3), 133–147 (2010)
23. Misra, I., Shrivastava, A., Hebert, M.: Watch and learn: Semi-supervised learning of object detectors from videos. *CoRR abs/1505.05769* (2015)
24. Moody, S., Melia, S.: Shared space research, policy and problems. *Proceedings of the Institution of Civil Engineers - Transport* 167(6), 384–392 (2014)
25. Murphey, Y.L., Milton, R., Kiliaris, L.: Driver’s style classification using jerk analysis. In: 2009 IEEE Workshop on Computational Intelligence in Vehicles and Vehicular Systems, CIVVS 2009 - Proceedings. pp. 23–28 (2009)
26. Naujoks, F., Neukum, A.: Specificity and timing of advisory warnings based on cooperative perception. *Mensch und Computer Workshopband* pp. 229–238 (2014)
27. NHTSA, N.H.T.S.: 2016 Motor Vehicle Crashes: Overview. Traffic safety facts research note 2017, 1–9 (2017), <https://crashstats.nhtsa.dot.gov/Api/Public/Publication/812456>
28. Pan, B., Zheng, Y., Wilkie, D., Shahabi, C.: Crowd sensing of traffic anomalies based on human mobility and social media. In: *Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. pp. 344–353. SIGSPATIAL’13, ACM, New York, NY, USA (2013)
29. Pande, A., Chand, S., Saxena, N., Dixit, V., Loy, J., Wolshon, B., Kent, J.D.: A preliminary investigation of the relationships between historical crash and naturalistic driving. *Accident Analysis & Prevention* 101, 107–116 (2017)
30. Ruikar, M.: National statistics of road traffic accidents in India. *Journal of Orthopedics, Traumatology and Rehabilitation* 6(1), 1–6 (2013)
31. Ryder, B., Gahr, B., Dahlinger, A.: An In-Vehicle Information System Providing Accident Hotspot Warnings. In: *ECIS 2016 Proceedings. Prototypes, AIS Electronic Library (AISeL)* (2016)
32. Ryder, B., Gahr, B., Dahlinger, A., Zundritsch, P., Wortmann, F., Fleisch, E.: Spatial Prediction of Traffic Accidents with Critical Driving Events — Insights from a Nationwide Field Study. *Transportation Research: Part A* (2017), submitted
33. Werneke, J., Vollrath, M.: How to present collision warnings at intersections? - A comparison of different approaches. *Accident Analysis and Prevention* 52 (2013)
34. Wirz, M., Strohrmann, C., Patscheider, R., Hilti, F., Gahr, B., Hess, F., Roggen, D., Tröster, G.: Real-time detection and recommendation of thermal spots by sensing collective behaviors in paragliding. In: *Proceedings of 1st International Symposium on From Digital Footprints to Social and Community Intelligence*. pp. 7–12. SCI ’11, ACM, New York, NY, USA (2011)
35. World Health Organization: Global Status Report on Road Safety 2015. Tech. rep. (2015)