

Autonomously Detecting and Classifying Traffic Accident Hotspots

Benjamin Ryder
ETH Zurich
Zurich, Switzerland
bryder@ethz.ch

Felix Wortmann
University of St. Gallen
St. Gallen, Switzerland
felix.wortmann@unisg.ch

ACM Classification Keywords

I.4.8. Image Processing and Computer Vision: Scene Analysis;
I.4.8. Image Processing and Computer Vision: Sensor Fusion;
I.4.9. Image Processing and Computer Vision: Applications

Author Keywords

Ubiquitous and mobile computing; Computer vision;
Machine learning; Traffic Accident Hotspots

PROBLEM DOMAIN AND OVERVIEW OF RELATED WORK

The number of road traffic fatalities has been steadily increasing since 2001 and is currently the eighth leading cause of death globally, with the loss of life of 1.2 million people each year according to the World Health Organization (WHO) [11]. In addition, the National Highway Traffic Safety Administration (NHTSA) reported that the number of deaths from traffic accidents in the USA increased by 7% from 2014 to 2015, rising to 35,092 fatalities [4]. Amid growing humanitarian concerns of so many injuries and fatalities worldwide, the Department of Transport issued a call to action encouraging the continuous research into different approaches that can improve the situation. As such, there are various research studies which are geared towards how in-vehicle systems can encourage drivers to adapt their driving behaviour and help to reduce the amount of both fatal and non-fatal traffic accidents. Typically, these systems aim to prevent a collision with an upcoming vehicle or pedestrian by providing warnings to drivers, and latest studies demonstrate promising evidence that these systems can indeed have significant positive effects [2, 9, 10]. However, the vast majority of studies have focused on simulation experiments [3, 8] and controlled lab experiments [6, 13]. We have recently contributed to this field by bringing an in-vehicle warning system into a field studying setting, utilising real world location analytics on traffic accident hotspots to generate in-vehicle warnings [7]. Going one step further, the benefit of ubiquitously detecting dangerous locations from data gathered by connected vehicles, and using these locations as a source for in-vehicle warnings,

has widely not been addressed in this growing domain and is the focus of our research.

The remainder of the paper is structured as follows, we first describe a preliminary investigation into accident hotspots and the impact specific locations had on the number of traffic accidents occurring nearby. We then outline the key idea of the project, along with two approaches which we are exploring to enable the vision of autonomously detecting and classifying accident hotspots with driving data. Finally, we conclude with our expected contributions to the field of ubiquitous computing.

RESEARCH AND EVALUATION CARRIED OUT SO FAR

Preliminary Analysis of Traffic Flow and Accidents relationship

Methodology

The first step of the research at hand was an exercise investigating the link between road traffic frequency and the number of traffic accidents, and whether the affect of individual locations, i.e. accident hotspots, also impacts the number of accidents. To investigate this, traffic accident data for five years was obtained from the Swiss Road Authority (FEDRO), this dataset contained GPS locations, as well as contextual information on the causes, on over 266,000 traffic accidents which occurred in Switzerland between 2011 and 2015. Additionally, the rate of traffic flow, i.e. the number of cars travelling on a particular stretch of road, has long been associated with traffic accidents. In order to assess the impact specific locations have on the number of traffic accidents, it is therefore important to account for traffic frequency. As such, traffic data was obtained from FEDRO, comprised of the average number of vehicles per day passing approximately 400 counting stations across Switzerland. This data was filtered to cover the same 5 year time period as the traffic accident data. Moreover, the locations for analysis were further limited to those where there was traffic frequency data collected over the full five-year period. Therefore, the final set of observations were constructed from 190 locations, each of which had five traffic frequency measurements, one per year. Accident data from each of these locations, for each year, was then collected by applying a naive grid-count approach. Here the number of accidents which occurred within a 1km² grid were counted for each year, where the centre of the grid was the traffic frequency counting station, similar to a previous traffic accident study [12]. Over the five year period, this came to a total of 9,102 traffic accidents, with an arithmetic mean

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UbiComp/ISWC'17 Adjunct, September 11–15, 2017, Maui, HI, USA
© 2017 Association for Computing Machinery.
ACM ISBN 978-1-4503-5190-4/17/09...\$15.00
<https://doi.org/10.1145/3123024.3123199>

	FE1		FE2		RE1		RE2	
ln (Traffic Frequency)	0.530	(1.56)	0.995**	(2.71)	0.500***	(10.14)	0.509***	(10.28)
Year=2011			0	(.)			0	(.)
Year=2012			-0.00209	(-0.05)			-0.00239	(-0.05)
Year=2013			-0.0362	(-0.82)			-0.0290	(-0.66)
Year=2014			-0.0643	(-1.43)			-0.0507	(-1.16)
Year=2015			-0.136**	(-2.93)			-0.116**	(-2.64)
Constant	1.763***	(126.65)	1.810***	(57.35)	1.763***	(29.52)	1.802***	(27.35)
Sigma_U	0.822		1.014		0.801		0.801	
Sigma_E	0.429		0.427		0.429		0.427	
ICC - Intraclass Correlation	0.786		0.849		0.777		0.779	

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1. Panel Data Analysis - The Impact of Traffic Frequency on Traffic Accidents (natural log)

of 9.58 and a geometric mean of 5.65 accidents per location per year. As such, the dependent variable in our analysis was the number of traffic accidents occurring, with traffic frequency as an independent variable. Since the data collected is longitudinal for each location, i.e. across time, Panel Data Analysis with a grouping on each location is suitable for analysis, and additionally allows us to add the time fixed effects as an independent variable. Using this approach we set to answer our two preliminary research questions:

RQ1 Does traffic frequency through a location impact the number of traffic accidents occurring near (within 1km²) that location?

RQ2 Does the location itself impact the number of traffic accidents occurring near (within 1km²) that location?

The first question can be answered by considering the significance of the traffic flow independent variable in our model. The second question can be answered by testing whether using fixed effects or random effects of the location lead to a significantly improved model fit to the data. Finally, fundamental assumptions for regression analysis are violated as we are utilising count data following a poisson distribution. As such, both the accident count data and traffic frequency count data were transformed via the commonly used natural log function [5]. Additionally, the natural log transformed traffic frequency was mean centred to generate the independent variable.

Evaluation

A summary of the results from the analysis can be found in Table 1. The evaluation starts with fixed effects to analyse the impact of variables that vary over time. Fixed effects remove the effect of time-invariant characteristics so we can assess the net effect of the predictors on the outcome variable. In this first model (FE1), this is simply the impact of traffic flow on the number of traffic accidents. Here we do not see a significant effect ($p=0.119$) of the number of vehicles travelling through a location and the number of traffic accidents near that location. We can extend model FE1 by including the year as a categorical variable in the analysis, in order to test whether these time fixed effects are needed in the model. Model FE2 adds these variables, and we see that the effect of traffic flow on the number of traffic accidents

becomes significant ($p=0.007$). With the exception of 2015, none of the year categories are significant when compared to 2011. However, when we test the null hypothesis that these year categories are simultaneously equal to zero, we are able to reject it ($p=0.026$). Hence, we can conclude that we should include time fixed effects in our fixed effect model, and thus that FE2 is the better model when compared to FE1. Since we are utilising natural log transformed dependent and independent variables, the interpretation of the coefficient of 0.995 is the following: for a 10% increase in traffic flow, we would expect an almost identical increase of 9.9% ($1.1^{0.995} = 1.099$) in the geometric mean of yearly accidents occurring near that location.

We next test whether incorporating random effects of the locations improve the quality of our model. Unlike the fixed effects model, the rationale behind the random effects model is that the variation across entities is assumed to be random and thus uncorrelated with the independent variables included in the model. In the first random effects model (RE1), we take the same approach as FE1 and test first for the impact of traffic flow on the number of nearby accidents to the location. Here we see a strong significant effect of the traffic flow ($p=0.000$), where the number of yearly traffic accidents occurring within 1km² of the location would be expected to increase by 4.9% for a 10% increase in traffic flow at that location ($1.1^{0.500} = 1.049$). In the same way as FE1 and FE2, we can continue to build this model by additionally testing whether time fixed effects are needed. Model RE2 adds these variables, and we see that the effect of traffic flow on the number of traffic accidents remains highly significant, and none of the year categories are significant when compared to 2011, with the exception of 2015. However, unlike the fixed effects model FE2, when we test the null hypothesis that these year categories are simultaneously equal to zero, we fall short of satisfyingly rejecting the null hypothesis ($p=0.053$). Hence, we can conclude that we should not include time fixed effects in our random effect model, and thus that RE1 is the better model when compared to RE2.

To determine whether the random effect models are a better fit for the data than the fixed effects models, we apply Hausman tests between FE1 and RE1, and FE2 and RE2. This tests whether the unique errors in the fixed effects model are

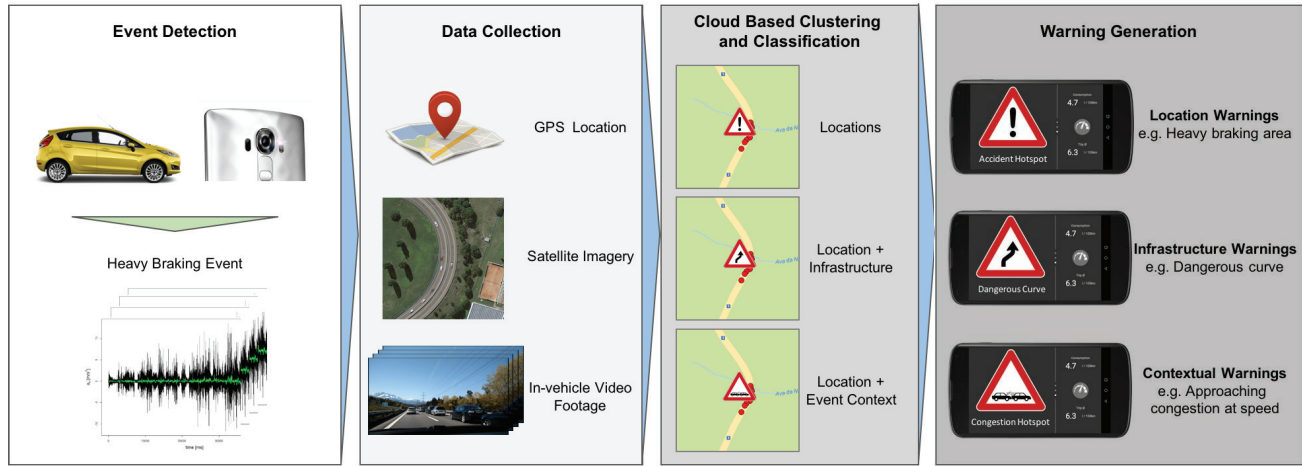


Figure 1. Pipeline of event detection, data collection and classification

correlated with the regressors, and as such, the null hypothesis is they are not. Thus, the null hypothesis is that the preferred model is random effects vs. the alternative fixed effects model. In this case, we fail to reject both the null hypotheses that 1. RE1 is better than FE1 ($p=0.878$), and 2. RE2 is better than FE2 ($p=0.929$). As such, both the models incorporating the random effects of locations are preferred over the fixed effects models. This affirms our second research question, that the location itself impacts the number of traffic accidents occurring near (within 1km^2) that location. Additionally, since the effect of the traffic flow on the number of traffic accidents remained significant through both the random effects models, we are able to answer our initial research question, that traffic frequency through a location impacts the number of traffic accidents occurring near (within 1km^2) that location.

Finally, the intraclass correlation (ICC) of the models gives an insight into the fraction of the variance which is associated with locations. It is defined as the fraction of the total variance that is due to variation between groups:

$$ICC = \frac{\text{Sigma}_U^2}{(\text{Sigma}_U^2 + \text{Sigma}_E^2)}$$

In our best fitting random effect model, 77.7% of the variance is due to differences across locations. While our analysis finds high significance in the effect of traffic flow on the number of accidents, the majority of variance in our data, and associated models, comes from the locations themselves. Finally, this is shown more clearly by applying a Breusch and Pagan Lagrangian multiplier test for random effects on the original data. The null hypothesis of this test is that the variance of the random effect is zero for the grouping of location. Since this test is highly significant in our dataset ($p=0.000$), we reject the null hypothesis and can conclude that the area surrounding the location itself plays a significant role in the number of accidents which occur there.

Hence, with the knowledge that certain areas on the road network are more dangerous than others, we continue our research with the motivation of improving driving behaviour through these particularly hazardous locations, otherwise known as accident hotspots.

ORIGINAL KEY IDEA AND CORRESPONDING HYPOTHESIS OF THE THESIS

In a previous paper, we demonstrated that providing drivers with warnings of upcoming accident hotspots improved their driving behaviour through these hazardous locations over time [7]. Additionally, prior simulation studies have shown that in less critical situations contextual warnings are more suitable and preferred by users than a standard stop sign warning [3]. The most reliable method of generating these in-vehicle warnings is applying location analytics to national traffic accident data. However, utilising this data comes with a variety of challenges. First, such data is not typically available at a national level, i.e. Germany and many other countries, and might not be openly or commercially available. Second, accident data is often only provided in an aggregated format, and does not include the locations and causes of individual accidents. Finally, any useful data is commonly only made available on a yearly basis. As such, rapidly developing accident hotspots, e.g. due to temporary road maintenance, are detected too late to prevent accidents and might have been resolved by the time the data is analysed. Thus, we see a growing need to analyse real-time data collected from vehicles to detect and classify hazardous areas. To tackle this problem, we are considering the combination of two techniques. The first is to use satellite imagery of the GPS locations of driving events to identify the road challenges, i.e. the 'Infrastructure', which might influence the accident likelihood. The second aims to discover the causes for these dangerous events, i.e. the 'Event Context', and considers video captured from systems inside the vehicle itself to identify the diverse reasons why an event could have occurred. An overview of these approaches is shown in Figure 1.

RESEARCH AND EVALUATION PLANNED AHEAD

Satellite Imagery Classification of Traffic Accident Locations

As part of our previous field study, we were able to identify the GPS locations of dangerous driving behaviour, i.e. heavy braking or swerving, through a prototype system collecting



Figure 2. Example of video frames captured from a heavy braking event, triggered by a vehicle disregarding the right of way



Figure 3. Examples of video quality issues experienced in a naturalistic field study setting

Controller Area Network (CAN) Bus data from the vehicles. Thus, we are able to provide new warnings to drivers of ‘Heavy Braking’ or ‘Swerving’ hotspots on the basis of the driving behaviour. Augmenting these warnings with road infrastructure information, such as ‘Roundabout’ or ‘Intersection’, could help the driver to adapt their behaviour more effectively. The first approach we are exploring is using image recognition techniques on satellite imagery of hazardous locations to identify the possible infrastructure causes of previous traffic accidents. This model can then be used to predict the infrastructure reasons for newly detected dangerous event locations.

For the classification of traffic accident locations, we make use of the previously mentioned data from over 266,000 traffic accidents covering five years and collected from the Swiss Road Authority (FEDRO). This dataset contained the GPS location and contextual information on the road type, infrastructure involved and causes of traffic accidents occurring in Switzerland between 2011 and 2015. We split this accident data into 20 different combinations of two categories: road type, e.g. ‘motorway’ and ‘main road’, and infrastructure, e.g. ‘roundabout’ and ‘intersection’. The accidents making up these categories were then randomly sampled in a balanced way, and satellite imagery of the location of each accident was captured from Google Maps. This resulted in a dataset over 7,000 labelled satellite images of the road type and infrastructure surrounding the area of an accident. High level feature vectors for the labelled images were then generated using TensorFlow Inception-v3. These

feature vectors can be used for machine learning in order to detect characteristics of road infrastructure from the images, and then identify infrastructure in new images which have not been previously seen by the model. Initial classification results of the retrained Inception neural network indicate an accuracy rate of approximately 30%, significantly higher than randomly guessing between the classes, and thus indicate that the approach has potential. However, we see many improvements which can be made on our initial tests and are continuing the analysis in order to raise the prediction accuracy of the model to a usable level.

Smart-phone Video Classification of Driving Event Context

Whilst knowing the infrastructure challenges which may be faced at specific hazardous locations can enable more detailed warnings to drivers, it does not capture the full picture of an accident hotspot. In order to develop deeper insights into why near-miss accidents and other dangerous events occur, the next stage of our research will make use of video data captured from vehicles while they are driving. Videos are recorded from a smart-phone mounted in the vehicle which acts as a ‘dashcam’, continuously capturing video data and streaming this to a server when a dangerous event is detected by the vehicle. Figure 2 shows an example of the video data collected during a heavy braking event. As such, in our upcoming field study, we will collect both vehicle CAN and smart-phone video data from a fleet of approximately 40 professional drivers over the course of two months. These videos will then be labelled with contextual information describing the situation leading to the

detected event, e.g. ‘another driver disregarding the right of way at an intersection’. Insights from this rich dataset could help us provide even greater situational context in warnings of upcoming hazardous locations to drivers. To achieve this, the second step of the project is to train machine learning classifiers on the video data, making use of computer vision techniques, such as object detection and recognition, in order to autonomously generate more detailed warnings for drivers.

It is important to also highlight the challenges which can be encountered through such a naturalistic field study setup, where the drivers even install the system in the vehicle themselves. The first of these is each of the driver’s interpretation on the correct smartphone setup and position in the vehicle, since it can lead to scene obstruction or poor field of view if the phone is incorrectly positioned. Where these issues exist it becomes almost impossible to use the videos for classification. Other issues come from adverse weather conditions, i.e. heavy rain, and videos captured during the night. Figure 3 shows examples of these image quality challenges already experienced in our naturalistic field study setting. However, these issues raises an interesting challenge for the project in the future, where we will try to automatically identify videos which are unsuitable for classification. The second core issue includes the additional work which will be required to accurately and reliably label the videos. Options for tackling this include making use of distributed human intelligence task platforms, such as Amazon Mechanical Turk [1]. Another approach would be to consider semi-supervised learning techniques [14], which are able to make use of a small amount of labelled data and a large amount of unlabelled data for training and classification tasks. However, while there are many difficulties, we believe that studying true behaviour of users in a naturalistic setting gives a much more accurate portrayal of how such systems could work in the real world. This will include obtaining user feedback from the drivers on issues such as privacy concerns from automatically capturing videos from the smartphone while driving.

CONTRIBUTION EXPECTED IN THE FIELD OF UBIQUITOUS COMPUTING

We expect to make further contributions to the field of ubiquitous computing through exploring several other topics as part of our field studies, although these are not the focus of this paper. These additional elements consider how users react to different types of eco-driving feedback provided by the smartphone, including realtime vs. delayed information, and contextual vs. numeric feedback. Moreover, we also consider the detection of individual contextual attributes, for example, the identification of the driver in the vehicle from their driving style, and how this driving style is linked to the driver’s personality.

Finally, the paper at hand envisions real-time in-vehicle warnings on upcoming road hazards, ubiquitously detected and identified through data collected from the vehicle and transmitted to cloud-based classification services. While capturing simple insights from vehicles, e.g. the locations of heavy braking events, can provide generalised warnings

of potentially dangerous locations, this does not enable contextual insights as to why a location is hazardous. Thus, our research focuses on investigating the effectiveness of two techniques to autonomously identify the possible causes for evasive behaviour and heavy braking events detected while driving. In tackling these challenging problems, we endeavour to apply ubiquitous computing principles to contribute to the ‘Target Zero’ aim of preventing traffic fatalities and serious injuries. We hope to contribute to the field of computer vision by applying and assessing event detection and classification techniques in a new setting. In addition, through field study experiments surrounding the impact of generated warnings on driver behaviour, we expect to further contribute to the topics of in-vehicle human computer interaction in the field of ubiquitous computing.

ACKNOWLEDGMENTS

This work is part of the Bosch IoT Lab, headed by Prof. Dr. Elgar Fleisch at the University of St. Gallen, Switzerland. We would also like to thank our colleagues on our Connected Car project, Andre Dahlinger, Bernhard Gahr, Timo Gessmann and Prof. Dr. Markus Weinberger.

REFERENCES

1. Michael Buhrmester, Tracy Kwang, and Samuel D Gosling. 2011. Amazon’s Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on psychological science* 6, 1 (2011), 3–5.
2. Juela Kazazi, Susann Winkler, and Mark Vollrath. 2015. Accident Prevention through Visual Warnings: How to Design Warnings in Head-up Display for Older and Younger Drivers. In *2015 IEEE 18th International Conference on Intelligent Transportation Systems*. IEEE, 1028–1034.
3. Frederik Naujoks and Alexandra Neukum. 2014. Specificity and timing of advisory warnings based on cooperative perception. In *Mensch & Computer Workshopband*. 229–238.
4. NHTSA. 2015. Motor Vehicle Crashes: Overview. (2015). <https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812318> Accessed 16.09.2016.
5. Mohammed A Quddus. 2008. Modelling area-wide count outcomes with spatial correlation and heterogeneity: an analysis of London crash data. *Accident Analysis & Prevention* 40, 4 (2008), 1486–1497.
6. Daniele Ruscio, Maria Rita Ciceri, and Federica Biassoni. 2015. How does a collision warning system shape driver’s brake response time? the influence of expectancy and automation complacency on real-life emergency braking. *Accident Analysis and Prevention* 77 (2015), 72–81.
7. Benjamin Ryder, Bernhard Gahr, Philipp Egolf, Andre Dahlinger, and Felix Wortmann. 2017. Preventing traffic accidents with in-vehicle decision support systems - The impact of accident hotspot warnings on driver behaviour. *Decision Support Systems* 99 (2017), 64 – 74.

8. Florian Seeliger, Galia Weidl, D Petrich, Frederik Naujoks, Gabi Breuel, Alexandra Neukum, and Klaus Dietmayer. 2014. Advisory warnings based on cooperative perception. In *Intelligent Vehicles Symposium Proceedings*. IEEE, 246–252.
9. Li Sian Tey, Sicong Zhu, Luis Ferreira, and Guy Wallis. 2014. Microsimulation modelling of driver behaviour towards alternative warning devices at railway level crossings. *Accident Analysis and Prevention* 71 (2014), 177–182.
10. Julia Werneke and Mark Vollrath. 2013. How to present collision warnings at intersections? - A comparison of different approaches. *Accident Analysis and Prevention* 52 (2013), 91–99.
11. WHO. 2015. Global status report on road safety. (2015). http://www.who.int/violence_injury_prevention/road_safety_status/2015/en Accessed 24.02.2016.
12. Bong-Min Yang and Jinhyun Kim. 2003. Road traffic accidents and policy interventions in Korea. *Injury control and safety promotion* 10, 1-2 (2003), 89–94.
13. Junyi Zhang, Koji Suto, and Akimasa Fujiwara. 2009. Effects of in-vehicle warning information on drivers' decelerating and accelerating behaviors near an arch-shaped intersection. *Accident Analysis and Prevention* 41, 5 (2009), 948–958.
14. Xiaojin Zhu. 2006. Semi-supervised learning literature survey. *Computer Science, University of Wisconsin-Madison* 2, 3 (2006), 4.