Rethinking transport
27–30 April 2020

TRA 2020 HELSINKI

Proceedings of 8th Transport Research Arena TRA 2020, April 27-30, 2020, Helsinki, Finland

Technical requirements for real-time traffic detection and dynamic infrastructure measures for safer behaviour

Moritz Berghaus[a]*, Adrian Fazekas[a], Markus Oeser[a]

[a]*Institute of Highway Engineering, RWTH Aachen University, Mies-van-der-Rohe-Strasse 1, 52074 Aachen, Germany*

Abstract

One of the main reason for road accidents is unsafe driving behaviour due to wrong perception of the road. Infrastructure-based road safety measures are most effective if they only target those drivers that drive unsafely. In order to influence individual drivers towards safer behaviour, their behaviour must be captured and evaluated in real-time. This requires the collection of vehicle trajectory data. We present a system that detects vehicle positions and speeds using thermal cameras and computer vision algorithms. The system uses the concept of nudging to reduce the vehicles' speeds and guide them along a safe trajectory. In order to nudge unsafe drivers individually and in real-time, the detection system needs to fulfil several requirements which we discuss in this work. Furthermore, we present methods of data acquisition able to fulfil these requirements.

Keywords: vehicle tracking; traffic safety; nudging; trajectory analysis; road geometry

* Corresponding author. Tel.: +49-241-80-25-240;
E-mail address: berghaus@isac.rwth-aachen.de

1. Introduction

Although most road accidents occur due to human errors, road design is often responsible for these errors. Roads should be designed in such a way that they match the expectations of the drivers. As a result, drivers automatically drive safely. This design principle is called "self-explaining road" (Theeuwes and Godthelp, 1995). Many aspects of this principle have been implemented in national road design guidelines. One example is the rule that the radii of two subsequent curves should not differ too much. However, this rule cannot be applied on motorway exits. On the motorway itself the radii are large, while in the exit smaller radii are required. Although drivers are aware that they have to decelerate in front of an exit, they might overestimate the curve radius and underestimate their speed. The system presented in this paper aims at reducing speeds in front of the exit in order to increase safety in the exit.

There are several approaches to make driver behaviour safer. Road safety campaigns aim at enhancing traffic safety usually by means of mass media (Hoekstra and Wegman, 2011). The goal is to change driver behaviour in general, without focusing on those road sections where unsafe behaviour frequently occurs. Speed limit enforcement cameras can effectively reduce speeds at specific road sections, but they often lack acceptance as they are often perceived as "rip-off" rather than a means of increasing safety. Soft measures such as speed indicator devices contain radar detectors to measure vehicle speeds and give a visual feedback to each individual driver whether their speed is safe, without enforcing speeding. Walter and Knowles (2008) found that these devices can reduce mean vehicle speeds by 1.4 mph (2.2 km/h) in free flow conditions. Taylor *et al.* (2000) argue that 1 mph (1.6 km/h) of speed reduction reduces accident frequencies by 3 % on higher speed urban roads. However, Quddus (2013) states that the average speed has no significant effect on accident rates, but a smaller variation in speed can reduce the accident rates. This indicates that it is most effective to target only the drivers with the highest speeds and slow them down.

For this purpose, the vehicle positions and speeds of all vehicles taking the motorway exit must be known at any given time. Radar sensors used in speed indicator devices are not suitable for this as they are only able to measure the speed of one vehicle at a time without measuring its position. Video systems are however able to provide trajectories of all vehicles within their field of view (Bommes *et al.*, 2016). The trajectories are evaluated in real-time to decide which drivers are driving unsafely.

Informing drivers about their current speed is only effective when drivers make a conscious decision to drive more safely. In complex driving situations such as motorway exits, the driver might not have the cognitive capacity to make this decision. Instead, we use a more subtle approach called nudging. A nudging intervention works on an intuitive level and can therefore evoke safer behaviour without requiring the driver's awareness. (Karlsson *et al.*, 2017)

The remainder of this paper is structured as follows: Section 2 gives an overview of the infrastructure measure that we propose to make driver behaviour safer, and presents the field test location. In section 3 and 4 we focus on the methods to detect and track vehicles in real-time, which is essential for the presented infrastructure measure. Section 3 describes the necessary steps to transform image coordinates into road coordinates. Section 4 explains how the vehicle positions and speeds are estimated, and discusses the challenges associated with this infrastructure measure and the field test location. Unfortunately, we cannot present results on the effectiveness of the infrastructure measure since the field test is still ongoing, but we give an outlook in section 5.

2. Overview of the infrastructure measure

The case study location for this infrastructure measure is a motorway exit in Eindhoven, the Netherlands. The motorway leads from the motorway ring around Eindhoven into the city centre. The speed limit decreases from 80 km/h in front of the exit to 70 km/h at the beginning of the exit down to 50 km/h in the curve itself. Before the implementation of the infrastructure measure, the speed distribution and the traffic volume have been measured by simple radar sensors. On the straight part of the exit lane, the average speed was 82 km/h, at the beginning of the curve (where the speed limit of 50 km/h begins) it was 57 km/h and in the middle of the curve it was 56 km/h. On the main lanes, the average daily traffic volume adds up to 30,000 vehicles, while 4,700 vehicles per day use the exit. The ratio of heavy goods vehicles (HGV) is approximately 2 %. This ratio is characteristic for a road with predominantly urban traffic.

The main goal of the infrastructure measure is to influence those vehicles driving unsafely and nudge them towards safer behaviour. From this main goal, the technical and functional requirements for the whole system are derived backwards. In order to influence unsafe drivers exclusively, the driving behaviour of each vehicle has to be evaluated in terms of safety and the unsafe drivers have to identified. To achieve this, the positions and speeds of all vehicles must be captured. This leads to three main steps that are conducted in real-time:

1. Detecting vehicle positions and speeds
2. Identifying vehicles driving unsafely
3. Influencing those drivers to make their behaviour safer

In this section we describe the three steps in some detail. We start with the chronologically last step (3) as this is logically the first step in the design of the infrastructure measure.

The infrastructure measure aims at reducing vehicle speeds and make their trajectories along a curve safer. Drivers are influenced by a dynamic light pattern using the principle of nudging. The lights are state-of-the-art LED road studs embedded in the road surface. In order to avoid a haptic effect when driving over the road studs, the road studs are placed along the edge of the exit lane on top of the road marking. Each road stud must be controlled individually to create the desired pattern. The pattern must be shown only in front of the driver who is to be influenced to avoid influencing other drivers ahead or behind who drive safely. That means the light pattern is dynamic in space. Since the light pattern has a certain length, it cannot be shown if two vehicles follow each other too closely. This means the infrastructure measure cannot influence all unsafe drivers. However, it can be argued that small headways occur more frequently at low speeds.

Before influencing the drivers, their current behaviour must be assessed in terms of safety (step 2). Only those vehicles driving unsafely must be nudged. For this purpose, we define a speed threshold that is depending on the position along the exit lane. In front of the curve, the speed threshold decreases, while it is constant in the curve. The speed threshold is above the speed limit and above the average speed in order to target only the fastest drivers. If a vehicle is above the threshold, the light pattern will turn on and the driver will be nudged. If drivers change their speed along the road and pass the threshold between safe and unsafe, the light pattern has to react to this by turning accordingly on or off. At a later stage, more elaborate criteria for unsafe behaviour can be used by taking into account a combination of speed, longitudinal and lateral acceleration, and lateral deviation from the centre of the lane.

In order to show the light pattern at the right position and to the right drivers, the positions and speeds of all vehicles on the whole road section must be known (step 1). Cross-section based sensors such as radar detectors or inductive loops are not suitable for this because they cannot adapt to changes in speed early enough. Instead, a set of cameras are installed along the road, each covering a section of up to 100 m. In curves the range of a camera may be lower due to vegetation next to the road. With multiple cameras, vehicles can be tracked along the road, so their positions and speeds are known at any given time. For the case study, we use thermal cameras as they can detect vehicles also at night. An additional very important aspect is that number plates are not visible in thermal imaging, so there are no privacy concerns with this technique.

The remainder of this paper focuses on the detection of the vehicles (step 1). A more detailed description of the whole infrastructure measure can be found in (Köhler *et al.*, 2019).

3. Road Geometry and Coordinate Systems

In the process of detecting the positions of vehicles on the road using cameras, there are three different coordinate systems involved. The cameras use an image coordinate system, the road alignment is described in a global coordinate system, and the vehicle positions relative to the road are described in a local coordinate system. The correct transformation between these coordinate systems is crucial for the accuracy of the detection.

3.1. Global Road Coordinate System

The horizontal alignment of the road is described in a global coordinate system, e.g. UTM. The road can be described either as a sequence of straight lines, circles and clothoids or as a polygonal chain. If plans of the road are not available, a polygonal chain can be extracted from a map of orthophotos from satellites. The points of the polygonal chain can then be fitted to a continuous line, e.g. a polynomial of k-th degree. The polynomial is two-dimensional and gives the North- and East-coordinate as a function N(x) and E(x), where x is the distance to some reference point, measured along the road.

$$\begin{bmatrix} E(x) \\ N(x) \end{bmatrix} = \begin{bmatrix} \sum_{i=0}^{k} p_{E,i} x^i \\ \sum_{i=0}^{k} p_{N,i} x^i \end{bmatrix}$$

For the case study road, we manually extracted the coordinates of 104 points along the road edge from satellite images. We then divide the road in four sections, each approximately 90 m long, and fit a 3[rd] degree 2D polynomial to the road coordinates of each section. The sections overlap each other by 10 m to ensure that the transition between the polynomials is smoother. The goodness of fit in terms of the Root-Mean-Square Error (RMSE) is listed in Table 1 and the residuals are shown in Fig. 1.

Table 1. Goodness of fit for fitting the road coordinates to a set of polynomials

	Section 1	Section 2	Section 3	Section 4
RMSE East Coordinates [m]	0.060	0.076	0.300	0.039
RMSE North Coordinates [m]	0.015	0.021	0.126	0.067

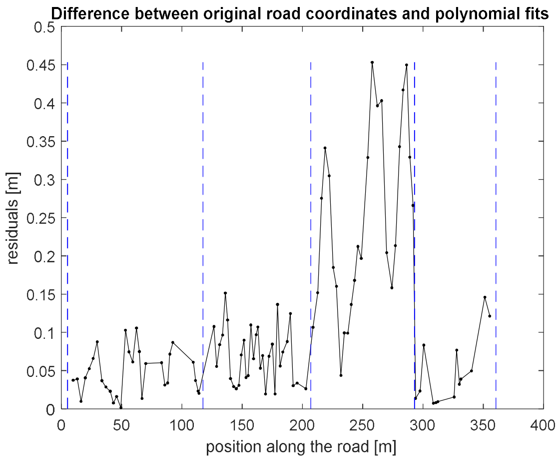

Fig. 1 Residuals of the polynomial fits computed as the Euclidian distance between the road points and the corresponding points of the polynomials

At the transition between the four sections, the gaps between the polynomials are 0.033 m, 0.188 m and 0.142 m respectively. Hence, this method gives a good approximation of the road alignment, even if the accuracy of the input data is low.

3.2. Local Road Coordinate System

The positions of vehicles are described in a local coordinate system, i.e. in coordinates relative to the road. The x-coordinate is the position along the road. The y-coordinate is the orthogonal distance to the (right) edge of the road (see Fig. 2). This enables to describe the trajectory of a vehicle relative to the road. The transformation from local to global coordinates is simple and efficient if the road is represented by a polynomial. The normalised normal vector of the polynomial is required to account for the distance to the road edge.

$$\begin{bmatrix} x \\ y \end{bmatrix} \to \begin{bmatrix} E \\ N \end{bmatrix}: \begin{bmatrix} E(x,y) \\ N(x,y) \end{bmatrix} = \begin{bmatrix} E(x) \\ N(x) \end{bmatrix} + y * \begin{bmatrix} -N'(x) \\ E'(x) \end{bmatrix} / \left\| \begin{matrix} -N'(x) \\ E'(x) \end{matrix} \right\|$$

The transformation from global to local coordinates requires computing the distance between a point and a polynomial curve. Whether this is algebraically possible depends on the degree of the polynomial (Lennerz and Schömer, 2002). However, a numerical solution, e.g. with the Newton-Raphson method, is much easier to implement.

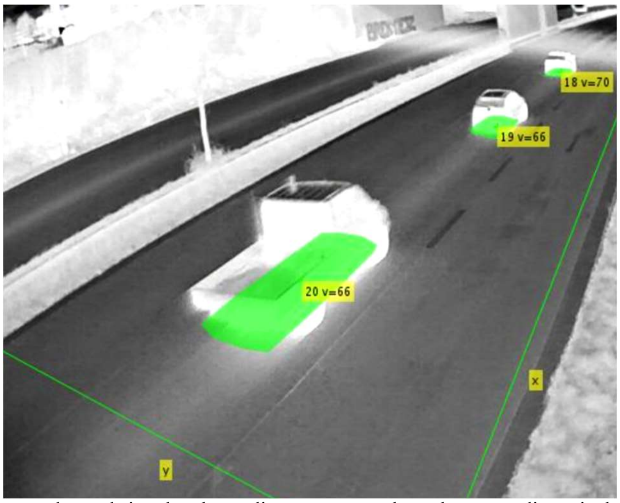

Fig. 2 Detection of vehicle positions and speeds in a local coordinate system, where the x-coordinate is the (longitudinal) position along the road and the y-coordinate is the lateral position, i.e. the orthogonal distance of the vehicle to the right edge of the road

3.3. Image Coordinate System and Camera Calibration

The vehicle positions are detected in image coordinates. The transformation between image coordinates and the global road coordinates requires an internal and external calibration of the cameras. For the internal calibration, the Matlab Camera Calibration Toolbox (Bouguet, 2015) is used. To estimate the camera parameters, usually a grid of black and white squares or circles is filmed from different orientations (Zhang, 2000). Since thermal cameras cannot distinguish colours, we use a laminated aluminium plate with circular holes that is heated up to make it visible for the thermal camera. For the external calibration we use a point cloud with characteristic points of the road and its surrounding, e.g. road markings, lampposts, traffic signs or trees. The points can be extracted from an aerial image. Six characteristic points serve to compute a first approximation of the rotation matrix and the translation vector. Afterwards, a manual fine adjustment is conducted by projecting the whole point cloud into the camera image.

4. Vehicle Detection and Tracking

In order to determine the position and velocity of a vehicle, two steps are necessary. Firstly, the detection of the vehicle in the area, in which the vehicles enter the field of vision. Each vehicle is represented by a point cloud that describes the corners and edges of the vehicle. The comparison of this projection to the distinctive features of the image enables the calculation of the probability of a car's presence on different positions in space (see Fig. 3) (Fazekas and Oeser, 2019).

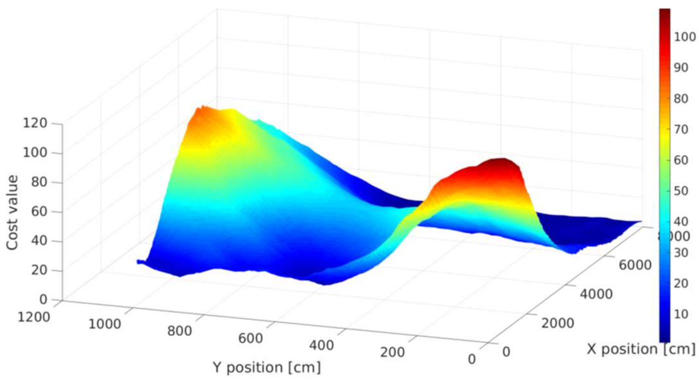

Fig. 3 Cost function representing the probability of the presence of a vehicle (Fazekas and Oeser, 2019)

As soon as the presence of a vehicle is detected, its position is tracked in the following images. For this, the sequential Monte Carlo method, also called particle filter (Forsyth and Ponce, 2012) (Fazekas *et al.*, 2013), is used. A certain number of states that represent the movement of the vehicle in spaces is created. For each state the corresponding position of the vehicle in the next time step is translated to image coordinates using the methods presented in section 3. Thus for each state a hypothesis how the vehicles boundaries should be visible in the image is proposed. By comparing the postulated boundaries to the actual image features these hypotheses are assessed. This method is referred to as Chamfer-Matching (Borgefors, 1988). For the state with the highest compliance, the corresponding position in the street system of coordinates is marked as the position of the vehicle. The speed is then computed as the derivative of the position with respect to time. Since inaccuracies in the position increase in the derivative, a moving average filter over the period of 1 second is used to smooth the speed values. This procedure is repeated until the vehicle leaves the field of vision of the camera.

In order to track vehicles over a longer period, multiple cameras are installed in a row. In the gaps between those cameras, the movements of the cars have to be calculated by extrapolation. Once a vehicle leaves the field of vision of one camera, it is assumed that the vehicle maintains the speed that was last detected by the camera. If the gaps between the cameras are sufficiently small, the exact time when the car enters the field of vision of the next camera is predicted reliably (see Fig. 4). Hence, each vehicle is assigned a distinct ID that relates to all cameras. This allows us to analyse driver behaviour over a longer distance.

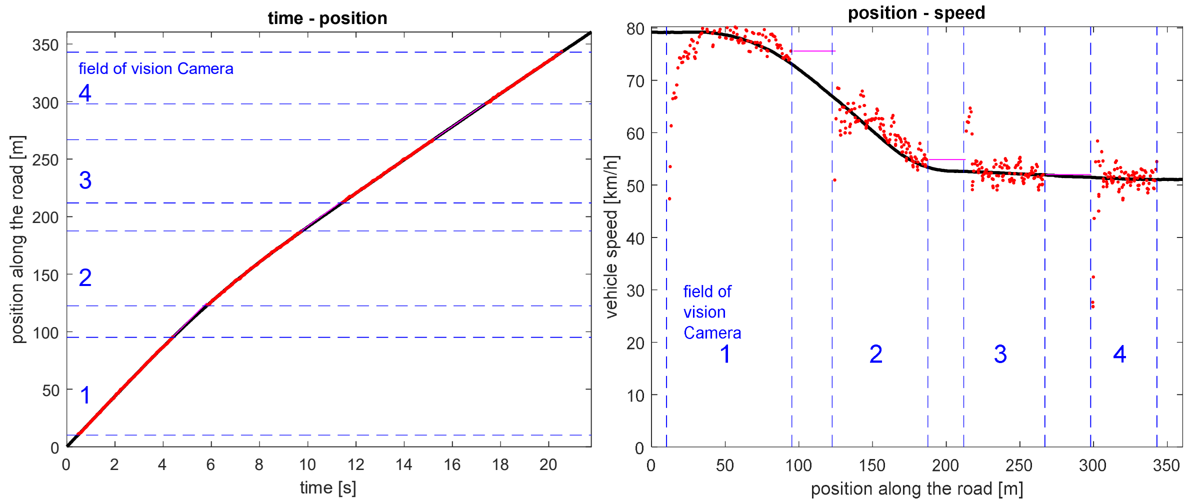

Fig. 4 Example trajectory with true positions and speeds (black line), detected positions (assuming a normally distributed position error of 0.5 m) and computed speeds (red dots), and extrapolated positions and speeds (magenta line)

As a result of the above-mentioned method, vehicle trajectories over the course of the exit are obtained. This method is efficient enough to be implemented in real-time and yet robust against measurement errors. The trajectories are used for two purposes. On the one hand, the Processing Unit receives the current position and velocity of each vehicle in real time and decides whether to nudge that vehicle. Due to the measurement errors, a hysteresis is implemented in order to avoid that the nudging turns on and off too frequently if the vehicle speed is close to the threshold. On the other hand, in order to evaluate the efficiency of the nudging measure, the trajectories of the field test are analysed.

5. Conclusions and Outlook

In this paper, we have presented a system that is able to detect vehicle speeds and positions along a motorway exit in order to individually nudge those vehicles driving unsafely. Computer vision algorithms are used to track vehicles in thermal images and translate image coordinates into road coordinates. Since the system operates in real-time, both the tracking and the coordinate transformation must be computationally efficient. We therefore use a polynomial model of the road rather than a polygonal chain or the correct alignment. This allows a fast generation and evaluation of states during the tracking process, which is crucial for an accurate position estimation. The system is robust enough to identify the vehicles driving above the speed threshold reliably.

The system has been developed and implemented on a motorway exit in Eindhoven, the Netherlands. A long term field test with different nudging scenarios is currently conducted. Afterwards the effectiveness of the infrastructure measure will be evaluated by analysing the vehicle trajectories. In the future, the system can be adapted for different locations where other safety interventions have not worked or are not applicable. It would especially be interesting to validate the system in several countries, where obedience rates to traffic signs can be very different. The prospect of connecting the system to a centralized Advanced Traffic Management System (ATMS) would enable a higher level of adaption to external factors like weather conditions and traffic diversion. This would result in a more flexible Cooperative-ITS solution.

Acknowledgements

This work is part of the project MeBeSafe which has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 723430.

References

Bommes, M., Fazekas, A., Volkenhoff, T., Oeser, M., 2016. Video Based Intelligent Transportation Systems – State of the Art and Future Development. Transportation Research Procedia 14, 4495–4504.
Borgefors, G., 1988. Hierarchical chamfer matching: a parametric edge matching algorithm. IEEE Transactions on Pattern Analysis and Machine Intelligence 10 No. 6, 849–865.
Bouguet, J.Y., 2015. Camera Calibration Toolbox for Matlab: Computational Vision at the California Institute of Technology.
Fazekas, A., Bommes, M., Oeser, M., 2013. Vehicle Tracking using 3D Particle Filter in Tunnel Surveillance and Incident Detection. Conference: Models and Technologies for Intelligent Transportation Systems 2013, At Dresden, Volume: 3.
Fazekas, A., Oeser, M., 2019. Performance Metrics and Validation Methods for Vehicle Position Estimators. IEEE Transactions on Intelligent Transportation Systems, 1–11.
Forsyth, D., Ponce, J., 2012. Computer vision: A modern approach, 2nd ed., Pearson, Boston.
Hoekstra, T., Wegman, F., 2011. Improving the effectiveness of road safety campaigns: Current and new practices. IATSS Research 34 No. 2, 80–86.
Karlsson, M., Brouwers, R., Fahrenkrog, F., Hof, T., Köhler, A.L., Ljung Aust, M., Selvefors, A., van Schagen, I., Sprajcer, A., Toffetti, A., Twisk, D., af Wåhlgren, A., Ziegler, J., 2017. Delivery report for MeBeSafe. Measures for behaving safely in traffic: Integrated Framework. Deliverable D1.1, available at: https://www.mebesafe.eu//wp-content/uploads/2018/04/MeBeSafe_D1.1_Integrated-framework.pdf (accessed 30 June 2019).
Köhler, A.L., Op den Camp, O., van Mierlo, M., Ladwig, S., Schwalm, M., 2019. Nudging Drivers Towards Higher Safety Margins - Applications of the H2020-Project MeBeSafe. 13th ITS European Congress, Brainport, The Netherlands, 3rd-6th June 2019.
Lennerz, C., Schömer, E., 2002. Efficient distance computation for quadratic curves and surfaces, in Suzuki, H. (Ed.), *Geometric modeling and processing:* Proceedings ; theory and applications ; 10 - 12 July 2002, Wako, Saitama, Japan.
Quddus, M., 2013. Exploring the Relationship Between Average Speed, Speed Variation, and Accident Rates Using Spatial Statistical Models and GIS. Journal of Transportation Safety & Security 5 No. 1, 27–45.
Taylor, M.C., Lynam, D.A., Baruya, A., 2000. The effects of drivers' speed on the frequency of road accidents: Prepared for Road Safety Division, Department of the Environment, Transport and the Regions, *TRL Report*, Crowthorne, Berkshire, UK.
Theeuwes, J., Godthelp, H., 1995. Self-explaining roads. Safety Science 19, 217–225.
Walter, L., Knowles, J., 2008. Effectiveness of Speed Indicator Devices on reducing vehicle speeds in London, *Published Project Report*.
Zhang, Z., 2000. A flexible new technique for camera calibration. IEEE Transactions on Pattern Analysis and Machine Intelligence 22 No. 11, 1330–1334.