

Collision Mitigation for Crossing Traffic in Urban Scenarios

Philip Heck, Jan Bellin, Martin Matoušek, Stefan Wonneberger,
Ondřej Sychrovský, Radim Šára, Markus Maurer

Abstract—Current collision mitigation systems focus on rear end collisions. To address the full spectrum of real world accidents, these systems could be enhanced to cover more traffic situations. Vehicle to vehicle accidents in crossing traffic situations make up around 25% of accidents in Germany. This paper discusses the requirements and differences compared to rear-end collisions. Presented here is an action concept that takes into account how the impact configuration is changed by breaking the host (impacting) vehicle. Based on this concept the requirements for the detection of crossing traffic were derived. These requirements were met by developing a video system based on a monocular wide field of view camera. It is further shown how this action concept and sensor were integrated into a demonstrator vehicle and evaluated in full scale testing.

I. INTRODUCTION

Recent advances in in the field of driver assistant systems contribute to improving road safety. Halving the number of road deaths by 2020 is the target set forth by the European Commission [1]. Assisting drivers by supporting the driving task and removing human error with intelligent vehicle technologies is the driving force behind the development of advanced driver assistance systems with active intervention. Current systems for accident avoidance and collision mitigation focus on rear-end collisions and vulnerable road-users.

Intersection safety is a topic that has been intensively researched during the past decade. This paper focuses on the mitigation of unavoidable collisions at intersections using an on-board sensor. Previous projects have focused on avoiding collisions by assisting the drivers (stop-sign and red-light assistance). Collision mitigation was performed analogously to longitudinal traffic by emergency braking in case of unavoidable collisions. Decision making was often dependent on C2X communication technology or detailed maps, technologies that are not available in series production cars today.

By engaging the brakes in crossing traffic situations the time to collision is lengthened. Due to the speed reduction of the host vehicle the required time to travel the distance to

the point of collision increases. During this additional time (Δt) the crossing vehicle produces a lateral offset thereby changing the location of impact. While the braking maneuver does reduce the collision speed (v_{host}) the influences on the collision severity due to changed impact configuration leading to a possible increase in Δv^1 have so far not been considered in decision making.

This paper introduces an action concept for collision mitigation in crossing traffic that considers the effects on accident severity due to a changed impact configuration. On the basis of a detailed accident analysis, we describe the mitigation possibilities for a set of possible traffic situations in Section III. Subsequently we propose requirements for sensors and actuators to address these situations. Taken into account are results from a vehicle in the loop survey focusing on driver reaction. Section IV details the levels of escalation and shows how uncertainties are handled in automated decision making. A new approach to meet the formulated requirements in perception is introduced in Section V. By combining a low cost camera with a real time wheel detector we show how the passenger compartment can be detected in crossing vehicles. While the overall system architecture is presented in Section VI, results of test track testing in the VeHIL facility conclude this paper in Section VII.

II. RELATED WORK

A. Collision Mitigation at intersections

The PREVENT Project found misinterpretation, inattention and lack of consideration given to sight obstructions to be the main reasons for collisions at intersections [3]. To support the driver, numerous projects have contributed to intersection assistance systems. First Collision mitigation functions shown in public demonstrator cars were presented within the PREVENT subproject Intersafe [4].

The Intersafe-2 project was a successor to Intersafe. The focus to develop Cooperative Intersection Safety Systems led to a series of demonstrated functions among which Fürstenberg and Rössler showed emergency braking in crossing traffic [5]. The system either increased a driver initiated braking in high risk situations or applied full force automatically when a collision became unavoidable. Consideration to the change of impact configuration was not given; it was assumed, that a lowered impact velocity would automatically lead to reduced crash severity. Brännström et al. suggest a split action concept between the two colliding cars depending

P. Heck and S. Wonneberger are with Volkswagen Group Research, Dept. of Driver Assistance Systems and Integrated Safety, 38436 Wolfsburg, Germany {philip.heck, stefan.wonneberger}@volkswagen.de

J. Bellin is with Volkswagen AG, Dept. of Powertrain development, 38436 Wolfsburg, Germany jan.bellin@volkswagen.de

M. Maurer is with the Institute of Control Engineering, Dept. Electronic Vehicle Systems, Technical University Braunschweig, Hans-Sommer-Str. 66, 38106 Braunschweig, Germany maurer@ifr.ing.tu-bs.de

M. Matoušek, O. Sychrovský and R. Šára are with Center for Machine Perception, Department of Cybernetics, Faculty of Electrical Engineering, Czech Technical University in Prague, 166 27 Prague 6, Technická 2, Czech Republic, {xmatousm, sychrol, sara}@cmp.felk.cvut.cz

¹ Δv denotes the speed reduction of the vehicle as a result of the impact and is a measure for accident severity. For a detailed discussion see [2].

on the relationship of the velocities [6]. Pointing out the high potential for collision avoidance, a cooperative concept with only one car intervening could reduce the intervention time. The possibility of an intervention by the other car is considered by Kim et al. [7]. Based on a hypothesis of acceleration and deceleration of the crossing car the possible change in impact location is calculated.

Calculation of collision risk is an enabling factor in scene interpretation, but it is challenging in crossing traffic. Herrmann and Schroven combine a driver intention prediction with probabilistic reachable sets to determine the criticality of a scene [8]. Lefèvre et al. take the context of the traffic situation into account [9]. This paper focuses on intervention strategies for unavoidable collisions, therefore a simpler approach is used that can be combined with the mentioned approaches for further improvement.

B. Image-Based Object Detection for Crossing Traffic

Image-based crossing traffic detection, either feature-based or motion-based, has been discussed throughout the last decade. Having the advantage of utilizing a single moving camera with a wide field of view defined by the application, both approaches have to ensure required detection rates with correct classification and position estimation.

Motion based approaches have been shown e.g. in [10], [11]. The key is to separate ego and object motion. Klappstein et al. outline the limitations of such approaches depending on ego and obstacle trajectory when only relying on motion vectors [10].

Instead, feature-based object detection relies on a single frame interpretation, with the help of an image model and has been intensively studied in the past. In recent years, the need for real-time applications has risen, and hardware systems have been proposed to accelerate the computing-intensive object detection part. Many employ popular AdaBoost-based classifiers inspired by Viola and Jones [12]. Haar-like features with cascaded classification implemented on parallel hardware reach real-time for a large number of test samples, such as in [13], [14]. Beside Boosting, Support Vector Machines are a popular machine learning and detection tool, demonstrated on parallel hardware as well, e.g. [15]. Due to an enormous search space, when considering all possible object positions in an image, another approach is to calculate ROIs in advance using information such as edges at the price of increased false negatives [16].

Our approach combines the AdaBoost framework with more complicated features that use the magnitude of the dot product between the image sample and a complex-valued kernel to achieve rotational invariance. These rotationally invariant responses help to avoid overfitting that would result if every possible rotation would be represented in the training set. To achieve a low latency we avoid using a classification cascade and use a massively parallel implementation instead.

III. ACCIDENT ANALYSIS AND DERIVED REQUIREMENTS

Collision mitigation by braking in longitudinal traffic assures that the accident severity is reduced. Accident data

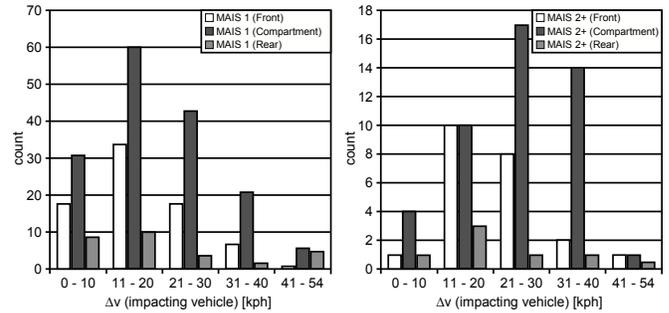


Fig. 1. GIDAS injury distribution for crossing traffic accidents grouped in minor (left) and major accidents (right).

of the German in depth accident study (GIDAS)² can help to understand the influences caused by changing the impact configuration in crossing traffic and derive new requirements.

A. Accident Analysis

To describe the accident situation we grouped all accidents into one of three impact zones. Front and rear impact mean that no part of the host vehicle collided with the side of the passenger compartment. As no precise definition for the passenger compartment exists³, a definition for this paper has been made by assuming that the passenger compartment is the part of the car between the front and rear axle.⁴ The injuries are described by the commonly used abbreviated injury scale (AIS) in which the maximum score for each involved passenger is given (MAIS), and only belted occupants are taken into account.

Fig. 1 shows the distribution and injury severity found for the analyzed 338 accidents. The data is classified into three impact zones on the crossing vehicle (front, compartment, rear). Since the accident severity is determined mainly by Δv of the host vehicle all graphs are referenced by this value. An increased risk for minor and major injuries resulting from collisions with the passenger compartment is apparent down to low speeds.⁵ In addition, Δv is influenced by the impact configuration. Impacting at or near the center of gravity (which for most cars is located in front of the passenger compartment) leads to the highest Δv . Thus, an initial reduction of v_{host} might still lead to a higher Δv .

We conclude that besides reducing Δv , avoiding collisions with the compartment is the main contributing factor to reducing injury severity in crossing traffic. Therefore, we suggest implementing a system that avoids a collision with

²The data analyzed takes into account all accidents in the GIDAS database as of October 2012 that took place after 1999. The vehicle types are limited to normal passenger vehicles including MPVs and SUVs as well as light commercial vehicles. In addition accidents where a loss of control before the primary collision had taken place or multiple collisions/cars where involved are not taken into account. Every accident report was revisited to assure that the impact location was precisely coded.

³Neither SAE J1100, DIN 70200 or standard German literature give a common definition for the passenger compartment.

⁴Vehicles where the passenger compartment extends behind the rear axle e.g. vehicles with a second seat row are addressed by the action concept in Sec. III-B.

⁵See [2] for a more detailed analysis. O'Brien concludes that a median velocity at which MAIS2+ injuries occur is for $\Delta v \geq 34$ kph.

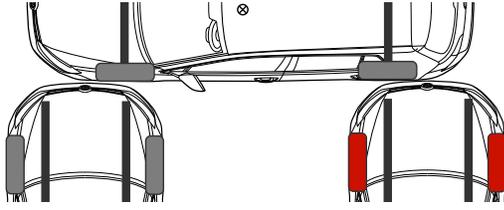


Fig. 2. Optimal frontal (left) and rearward (right) impact location each respectively with the longitudinal support that is closer to the compartment aligned with either front or rear axle.

the passenger compartment while reducing the host vehicle speed as much as possible.

B. Action Concept

According to our definition of the passenger compartment, limited by the front and the rear axles, we propose two optimal impact configurations illustrated in Fig. 2. For each configuration one longitudinal support structure of the host vehicle is aimed at the wheel or rim of the crossing vehicle.

For initial impacts in the front zone, the aim is to hit the front wheel of the crossing car with the right support structure of the host. This decreases impact speed while avoiding the risk of hitting the passenger compartment. When initial impact is inside the passenger compartment zone the aim is to hit the rear wheel of the crossing car with the left support structure of the host vehicle. This leads to an impact outside the passenger compartment while reducing the risks associated with hitting further back. These risks can be an increase in angular momentum and intrusion in the rear of the crossing car. As the rear axle is structurally rigid, even to side forces, it withstands a strong impact and limits intrusion. This is important when looking at the exception of our passenger compartment definition where a possible second or third seat rows exist. If possible, the aim should be to avoid a collision and pass behind the crossing vehicle.

As a result from this, we have to detect the front and rear wheel of the crossing vehicles and aim the host vehicle to collide precisely in this area leading to a required system accuracy of approximately the crossing car's wheel radius. If we assume a standard rim size of 16" this gives ± 20.3 cm of acceptable deviation.

C. Survey of driver reaction to intervention in crossing traffic

For an automatic intervention the driver reaction to the traffic situation has to be taken into account. This is important to find the trigger time for activation and to account for any driver reactions to the system intervention. As crossing traffic poses a different challenge for driver perception and scene interpretation, we conducted a study to evaluate driver reactions. The survey was conducted using a vehicle in the loop system where the driver was driving a real car with a helmet to display a virtual environment.

1) *Scenario description:* An urban environment was chosen to simulate a situation where crossing traffic collisions mainly occur. The host vehicle was driven with a speed of 50 kph. After a learning phase to familiarize the test subject



Fig. 3. Driver view of crossing in vehicle in the loop simulation during approach (left) and view with crossing car shortly before impact (right).

with the simulator a crossing traffic collision was provoked by the simulator. A car from either left or right was started to cross at a speed of 30 kph. The crossing car was hidden by sight obstructions until 1.3 seconds prior to impact. Fig. 3 shows the driver's field of view during approach and before the collision. The crossing car was timed so that a resulting collision would occur with the middle of the host vehicle impacting at the center of the crossings car front door.

2) *System intervention:* The system intervention was two staged. The first stage triggered an emergency brake with maximum deceleration and an audio warning 700 ms prior to impact. The second stage added additional steering in the direction of the back of the crossing car with a torque of 8 Nm at the steering wheel 600 ms prior to impact. No manual abortion was possible to study driver behavior and apprehension of the system reaction.

3) *Driver reaction:* During the experiment, all driver inputs were recorded and a questionnaire was answered by the test subjects afterwards. A total of 50 persons evenly distributed by gender, age, driving experience, and familiarity to driver assistance systems were tested. In the questionnaire, all drives rated the situation as surprising and dangerous. The total of 48% noticed a system intervention; of those 82% noticed a braking maneuver and 21% a steering motion. Only 8% wanted to change the system behavior and overrule the system while 7% felt a loss of control. Controllability was rated very high with 96% of the drivers feeling confident to overrule the system if required.

A total of 48 drivers claimed to have braked. Fig. 4 shows when the driver pressed the brake pedal and when full brake pressure was reached in relation to the system intervention time. Only three drivers started to brake before the automated intervention while 62% released the gas pedal until this point. The total of 49 subjects braked prior or during the intervention. Of those 33 reached full brake pedal force before impact.

The steering reaction was counteracted by all drivers. Fig. 4 shows that 80% stopped the turning steering wheel within 250 ms. The average offset produced by steering during the test was thus limited and stayed below 0.2 m. While an uncountered intervention would have produced around 0.66 m. In the questionnaire, only 29% of the drivers either noticed a steering reaction or claimed that they tried to steer. A reaction to the scenario by steering to avoid the accident was shown by none of the drivers.

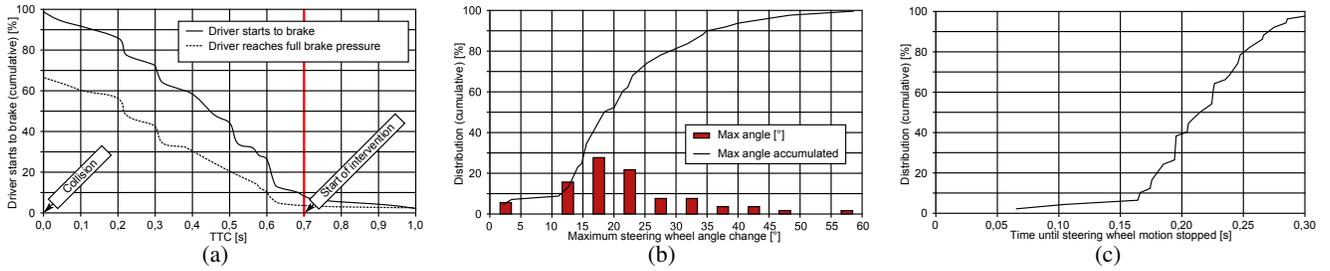


Fig. 4. (a) Time when the driver started to brake and time when full brake pressure was reached by the driver (with system support). (b) Maximum steering wheel angle change achieved by the system before driver stopped movement and cumulative value. (c) Time until steering wheel motion was fully stopped and 8 Nm of steering wheel momentum were compensated by the driver.

4) *Conclusion:* The automatic braking intervention was rated beneficial by all subjects. The unintentional reaction to the steering wheel motion by all drivers seems to show that a steering wheel momentum with 8 Nm will be overruled instinctively during unapprehended emergency situations. Further studies have to prove this for a drive-by-wire system. With more than 62% of the drivers already reacting to the situation before intervention, we conclude from the survey that an intervention at 700-900 ms TTC is best to support the driver.

IV. FUNCTIONAL DEFINITION OF DECISION MAKING ALGORITHM

The goal of the implemented decision making algorithm is to find the action that minimizes the risk of any road user involved in an accident to sustain an MAIS2+ injury.

A. Vehicle environment interaction

Predicting the outcome of a possible collision is an essential part in making the decision whether a particular maneuver should be conducted. Therefore, the first step in making a decision prior to an upcoming collision is finding a detailed analytical description of all involved road users and their interaction with the environment.

Since the action concept to influence an upcoming collision relies on the vehicle's ability to steer and brake, the most influential parameter is the friction coefficient μ . It is a direct measure of the forces that can be transferred from the rolling vehicle onto the road surface and depends on the condition of the road surface as well as the vehicle's tires and suspension. The transfer of forces from the demonstrator vehicle's actuators to its suspension and inertia of the vehicle body are represented in a simplified, statically parameterized vehicle model.

B. Handling of uncertainties by generating an evaluation parameter

The decision making algorithm has to consider the possibility that the parameters on which a decision is based may be inaccurate due to measurement errors or made assumptions. The possible influence of such an error on the impact configuration is illustrated in Fig. 5. In the example, the probability of injury to the passengers of the crossing vehicle could increase if the road friction coefficient were 0.6 while the system assumed it to be 0.9. The algorithm uses

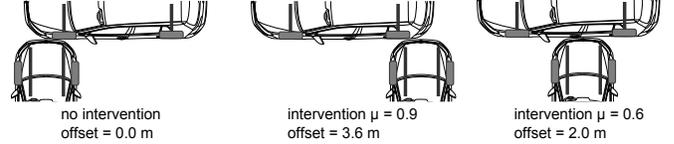


Fig. 5. Impact of an erroneous evaluation of μ

a stochastic approach to solve this problem. Based on the measurement accuracy of the friction coefficient estimation, it generates an evaluation parameter EP as a measure of criticality for every possible alternative action.

It can be assumed that every alternative requires a certain minimum road friction coefficient in order to be executed successfully and that it will result in one specific, determined accident configuration. This means that the probability of a successful execution $P_{success}$ equals the probability that the minimum required road friction coefficient is available $P_{\mu_{req}}$, because a μ lower than μ_{req} must result in a different accident outcome. The criticality of one specific action therefore has to be an accumulation factoring in all accident outcomes for $\mu \leq \mu_{req}$. The algorithm uses an array of elements $C_1 \dots C_n$, each one representing an accident outcome with C_1 containing the injury severity of the accident configuration caused by the desired action and a (theoretically indefinite) number of elements C_n containing those for $\mu \leq \mu_{req}$. The number of elements is limited to 10 per alternative action in the implemented algorithm which proved to be a good compromise between accuracy and computing time.

In Fig. 5 three possible outcomes of the attempt to minimize crash severity by changing the point of impact to the rear of the crossing vehicle in order to avoid hitting the passenger compartment are shown:

- No action can be executed (e.g. due to system error).
- The desired accident configuration can be reached.
- The road friction coefficient leads to an impact into the passenger compartment (erroneous assumption of μ).

Assuming, as an example, these outcomes are the only ones taken into account by the system, for each of them the algorithm determines the probabilities $P_{\mu \geq 0}$, $P_{\mu \geq 0.6}$ and $P_{\mu \geq 0.9}$ as well as the corresponding risks of an MAIS2+ injury ($P_{MAIS2+\mu \geq 0}$, $P_{MAIS2+\mu \geq 0.6}$, $P_{MAIS2+\mu \geq 0.9}$) for any road user participating in the collision. As explained earlier,

we can assume:

$$P_{MAIS2+\mu \geq 0.6} > P_{MAIS2+\mu \geq 0} > P_{MAIS2+\mu \geq 0.9}.$$

It is now possible to generate the elements that the evaluation parameter consists of:

- Element 1, which contains the criticality of the collision, assuming successful execution of the desired action:
 $C_1 = P_{\mu \geq 0.9} \cdot (1 - P_{MAIS2+\mu \geq 0.9}).$ ⁶
- Element 2, which contains the criticality of the collision, assuming $\mu \geq 1$ is not possible, but $\mu \geq 0.6$ is:
 $C_2 = (P_{\mu \geq 0.6} - P_{\mu \geq 0.9}) \cdot (1 - P_{MAIS2+\mu \geq 0.6}).$
- Element 3, which contains the criticality of the collision, assuming $\mu \geq 0.6$ is not possible, but $\mu \geq 0$ is:
 $C_3 = (P_{\mu \geq 0} - P_{\mu \geq 0.6}) \cdot (1 - P_{MAIS2+\mu \geq 0}).$

The evaluation parameter EP consisting of all three elements is:

$$EP = \sum_{n=1}^3 C_n = C_1 + C_2 + C_3.$$

To significantly reduce computing time when a large number of alternatives are evaluated, the EP s are generated sequentially, starting with the accident outcome requiring the lowest friction coefficient and ending with the highest one physically possible. The decrease in computing time is due to the fact that the result of previous calculations of:

$$EP_{n-2} = (P_{\mu \geq \mu_{n-2}} - P_{\mu \geq \mu_{n-1}}) \cdot (1 - P_{MAIS2+\mu \geq \mu_{n-2}}).$$

can be used by any evaluation parameter EP_x with $x > n-1$.

C. Choosing the appropriate action

Assuming there is a definite number of alternative actions, the algorithm can assign an evaluation parameter to each of them. The action with the greatest evaluation parameter EP_{max} represents the alternative that will currently yield the stochastically lowest crash severity.

D. Intervention criteria and levels of escalation

There are two types of a possible system intervention that have to be distinguished between. Passive intervention by warning the driver and active intervention through braking (or possibly steering). To trigger a system intervention two criteria have to be met:

- 1) A minimum situation criticality has to be reached. The measure for the criticality is EP_{max} .
- 2) The TTC has to be low enough to not cause an intuitive countermeasure by the driver. This measure is based on the driver reaction time.

For each criterion there is a lower and an upper threshold intervention TTC. The upper threshold time representing a value that would allow for an intervention, the lower threshold time representing a value that demands an intervention.⁷

⁶ $1 - P_{MAIS2+\mu \geq 0.9}$ represents the probability that none of the participants sustained an injury of grade MAIS2+.

⁷For the driver reaction time criterion, no lower threshold TTC exists for active intervention, because it only has to be assured that no maneuver will be executed by the driver.

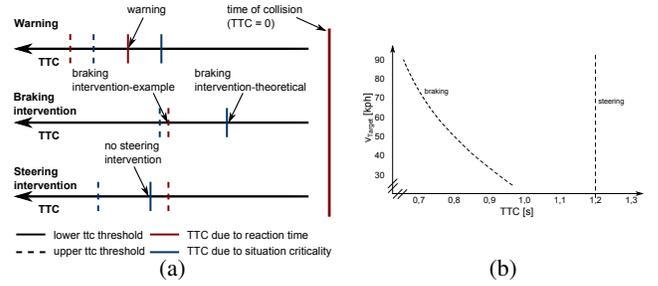


Fig. 6. (a) Levels of escalation and means of intervention (b) minimum TTCs depending on intervention type and target vehicle velocity.

This means that if the first lower threshold TTC is reached for any means of intervention while all upper threshold TTCs are passed, a corresponding action can be triggered. If there are various valid alternative actions, the last threshold TTC will trigger the intervention thereby prolonging the time left for a driver intervention. This process is described in Fig. 6 (a) for an example in crossing traffic. Based on the different upper and lower TTC threshold values, the system only allows for intervention through braking.

Fig. 6 (b) illustrates that braking will always yield superior results to steering in urban crossing traffic situations. It shows the TTCs at which a system activation has to be triggered by either means of intervention (braking or steering) to result in an offset large enough to avoid hitting the passenger compartment based on the crossing vehicle's speed v_{target} . For small TTCs, the transversal offset a car produces after a certain time only depends on the forces that are transferred transversally between its tires and the road surface. Therefore in this application the lateral offset of a steering intervention is independent of v_{target} . For crossing traffic, it can be concluded that a steering intervention will always require a larger TTC than accepted by the driver reaction time criterion (see Sec. IV-E).

E. Termination of intervention

An active system intervention will not be triggered or will even be aborted during execution, if:

- 1) The driver intervenes either by a strong steering motion or additional pressure on the gas pedal.
- 2) It is unlikely that the accident configuration can be influenced in a way that yields a reduction of crash severity without risking its increase. This is the case if the situation criticality demands a time of intervention that is higher than the maximum accepted TTC.

These criteria have to be met in order to satisfy the following demands:

- 1) The driver is always in control of the vehicle as demanded by the *Vienna Convention* [17].
- 2) A system intervention will most likely result in a crash severity lower than in a driver's intervention or no intervention whatsoever.

V. CAR DETECTION IN CROSSING TRAFFIC

With the requirement of using a sensor suitable for series production providing a wide field of view $> 100^\circ$ with a

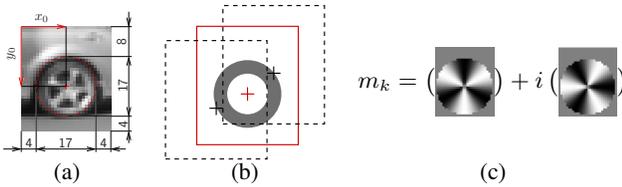


Fig. 7. (a) Scale-normalized image of a wheel with pixel dimensions. (b) Negative training samples (dashed) taken from the vicinity of the positive sample (solid). (c) Example of one rotationally invariant convolution kernel.



Fig. 8. Part of the input image (left), wheel confidence map (middle), and wheel confidence map aggregated by local structure model (right).

range of > 20 m and resolution sufficient for identifying the passenger compartment, a camera sensor is the first choice to meet these requirements. The presented image-based car detector for crossing traffic is build from several independent processes, running in a feed-forward pipeline.

A. Image Setup and Preprocessing

Gray-scale images are captured from a wide-field-of-view (190° horizontally) camera covering the front area of the vehicle. Radial distortion of the image region of interest is removed, and a conformal warp is applied, optimized to preserve local pixel resolution in the region. This ensures that wheels of a car moving perpendicularly to the optical axis appears circular (they are elliptical before the un-distortion in general), and are neither shrunk nor unnecessarily enlarged w.r.t. the original image.

B. Single-Scale Patch Based Wheel Classifier

The core element of the car detector is a car wheel classifier, designed for a single scale – an image of a wheel is assumed to fit into a 25×29 image patch, Fig. 7a. The AdaBoost learning uses features constructed by image patch projection to a set of kernels. Some of the kernels are rotationally symmetric and complex-number valued, the L_2 norm of their response is invariant to wheel rotation, Fig. 7c. A weak learner compares the individual kernel response with a learned threshold. For computational efficiency, we use quantization of kernel values to $\{-1, 0, 1\}$, and L_1 norm. We did not observe any significant performance loss over the case when real-valued kernels and L_2 norm are used.

The training set contained 24k annotated wheel samples and 100k of negative samples collected automatically from their vicinity, Fig. 7b. The classifier consists of 150 features, some kernels are re-used, so we only have 26 complex and 29 real convolution kernels.

Wheel confidence in every pixel is computed by classifying its neighbouring patch and then spatially aggregated using a wheel neighbourhood structure model, Fig. 8.

C. Wheel Candidate Detector

To extend the wheel classifier to scale-space, single-scale classifier is run on a Gaussian image pyramid with the scale



Fig. 9. A typical car detector result.

factor of $\sqrt{2}$. The maximal confidence across the scales is selected in every pixel of the original image. Finally, the positions of local maxima of the maximal confidences exceeding a learned threshold become wheel candidate detections, each with a scale and confidence value.

The method was tested for 8 scales, functioning well in 1:10 depth range. However, in the current real-time implementation, only 3 scales were used for 1:2 depth range required by the crossing traffic scenario.

To overcome computational complexity, kernels are convoluted on a dedicated hardware using FPGAs which computes 28 kernels in parallel resulting in a cycle time of 39ms for all scales.

D. Car Detector

Image is interpreted as a configuration of hypothesized cars (wheel pairs) and unassigned wheels. A probability density is defined over all such parameterized configurations. The probability term for a car hypothesized over two wheel candidates uses wheel confidence of both candidates, observed wheel-base to scale ratio, and a ratio of average horizontal to vertical image value variance in the middle section. The probability for an unassigned wheel uses wheel confidence only. All distribution parameters are learnt from annotations. The probability distribution includes a combinatorial prior term balancing different number of pairs in a configuration. This helps the method work well for an unknown number of cars. The optimization task is then maximum a posteriori problem over all configurations of non-overlapping pairs. We approximate the maximization by a greedy algorithm.

E. Target Car Depth and Velocity Estimation

Every observed car is now represented as a pair of wheel center locations in the image. The relative car-to-car distance is estimated from a single view using known camera calibration, the assumption of perpendicular motion and a prior on the true wheel-base in 3D. A large class of passenger cars (excluding SUVs and small cars) has a wheel-base length of 2.6 m with about 5% deviation, estimated from the 100 most frequent car types provided by the Czech Transportation Authority. In order to estimate the relative velocity vector of the target car, the 3D positions of cars are matched across subsequent frames using a nearest-neighbor algorithm.

A typical output is shown in Fig. 9. Using the parallel feature extraction and code optimized pre- and postprocessing the detector achieves a latency of 90 ms while running

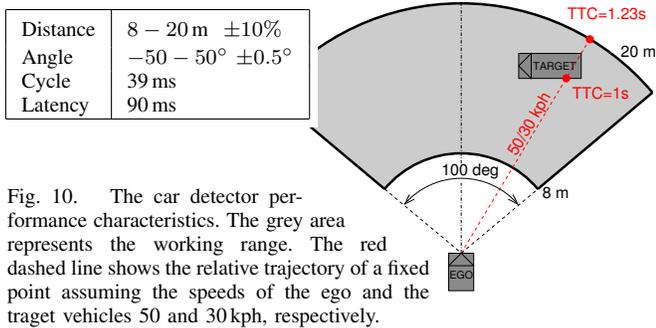


Fig. 10. The car detector performance characteristics. The grey area represents the working range. The red dashed line shows the relative trajectory of a fixed point assuming the speeds of the ego and the target vehicles 50 and 30kph, respectively.

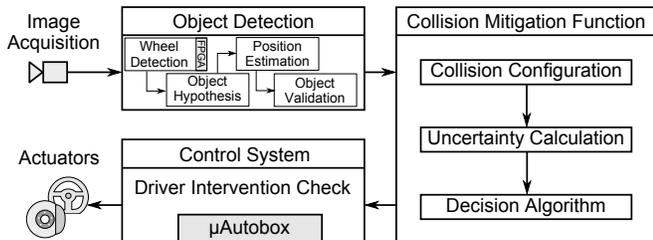


Fig. 11. Components of the collision mitigation system and data workflow.

with a cycle time 39 ms using the advantage of a pipelined workflow.

VI. SYSTEM ARCHITECTURE

The structural components and data workflow of the overall system are described in Fig. 11. An automotive 1.2 Megapixel camera has been integrated in the front grille. An image processing computer system performs preprocessing and vehicle hypothesis plausibilisation on a local CPU, while the wheel detection is performed on the attached FPGA co-processor card to achieve real-time performance. Crossing objects, validated over multiple frames, are transferred to a second computing system using a real-time middleware. The function processing system first creates the evaluation parameter map of resulting collision impacts for different alternative actions combined with the host and crossing vehicle’s state for each potential collision object (see IV). Intervention events are sent to the low level vehicle control system. In order to react immediately to driver interventions it has been realized using a μ Autobox which provides minimal latency. For an active intervention calculated braking and/or steering parameters are delivered to the control system taking actuator latency into account. A control algorithm has been developed to assure that the collision position is reached at the calculated point in time ($TTC + \Delta t$) and speed.

VII. EVALUATION

Evaluation of the system and sensor performance was performed at the TNO VeHIL facility in Helmond, Netherlands.

A. Test Setup

The demonstrator vehicle approached the target in a straight line with 90° angle. Fig. 12 shows the impact location. The target was pulled on a sled with 20kph. The start was synchronized to the demonstrator start in order to



Fig. 12. Location of impact without intervention in VeHIL testing (left) and detailed view including laser measurement of target (right).

TABLE I
VEHIL TEST RUN RESULTS

	Testrun				
	1	2	3	4	5
x distance measurement [m]	8.23	8.90	8.66	8.75	8.78
x distance measurement error [m]	0.16	0.22	-0.05	0.3	0.24
x brake position error [m]	0.1	-0.17	-0.05	0.14	-0.09
Error in calculated TTC [ms]	-40	56	14	-42	30
Latency compensation error [ms]	14	19	-4	27	21
Calculated impact location y [m]	4.02	4.06	4.09	4.03	4.07
Real Impact location y [m]	3.71	4.26	4.19	3.64	4.11
Impact position error y [m]	0.3	-0.2	-0.1	0.39	-0.04

impact at the same position. The demonstrator was traveling at constant 41 kph prior to impact. The floor in VeHIL has a value of $\mu = 1.2$. To achieve comparable results to normal roads the deceleration of the demonstrator was limited to a maximum of -9.5 m/s^2 . Reference measurements of the target were taken by laser and of the demonstrator by Correvit optical measurement. Once the collision was unavoidable the system activated to prevent a collision with the passenger compartment and aimed to crash into the rear wheel of the target with the host’s left longitudinal support structure.

B. Test results

Fig. 13 shows the recordings of five planned interventions in consecutive test runs. Tab. I compares the measured values to the reference measurements taken. It can be seen that the x distance measurement of the target by the camera has an error of about 2%. The x brake position error gives the difference between the calculated position given to the control system and the actual position. It is in the 2% range for the distance traveled. The error in calculated TTC is a combined result from the measurement and positioning error that is influenced by the assumed latency. The assumed latency is calculated during runtime on a non real-time windows system and is used for prediction of the incoming object data. The range is within the expected performance of the used hardware and software. Summed up, this results in the impact position error y which is the difference between the calculated and the real impact location ranging from +39 cm to -20 cm.

C. Conclusion

Within the confined conditions of VeHIL, the system performed within the desired specifications in situations where an unavoidable collision had to be mitigated by active

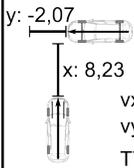
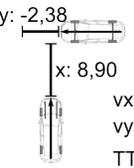
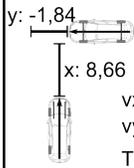
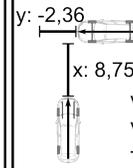
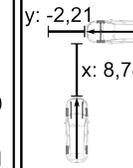
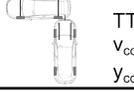
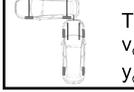
Scenario	Testrun 1	Scenario	Testrun 2	Scenario	Testrun 3	Scenario	Testrun 4	Scenario	Testrun 5
	$y: -2,07$ $x: 8,23$ $V_{x_{host}}: 11,18$ $V_{y_{target}}: 5,56$ $TTC: 0,736$		$y: -2,38$ $x: 8,90$ $V_{x_{host}}: 11,23$ $V_{y_{target}}: 5,61$ $TTC: 0,792$		$y: -1,84$ $x: 8,66$ $V_{x_{host}}: 11,17$ $V_{y_{target}}: 5,61$ $TTC: 0,775$		$y: -2,36$ $x: 8,75$ $V_{x_{host}}: 11,20$ $V_{y_{target}}: 5,61$ $TTC: 0,781$		$y: -2,21$ $x: 8,78$ $V_{x_{host}}: 11,20$ $V_{y_{target}}: 5,61$ $TTC: 0,784$
Collision without intervention	 $TTC: 0,736$ $V_{col}: 11,18$ $y_{col}: 2,02$	Collision without intervention	 $TTC: 0,792$ $V_{col}: 11,23$ $y_{col}: 2,07$	Collision without intervention	 $TTC: 0,775$ $V_{col}: 11,17$ $y_{col}: 2,51$	Collision without intervention	 $TTC: 0,781$ $V_{col}: 11,20$ $y_{col}: 2,03$	Collision without intervention	 $TTC: 0,784$ $V_{col}: 11,20$ $y_{col}: 2,18$
Collision with intervention	 $TTC+\Delta t: 1,095$ $V_{col}: 2,49$ $y_{col}: 4,02$	Collision with intervention	 $TTC+\Delta t: 1,148$ $V_{col}: 3,14$ $y_{col}: 4,06$	Collision with intervention	 $TTC+\Delta t: 1,057$ $V_{col}: 4,16$ $y_{col}: 4,09$	Collision with intervention	 $TTC+\Delta t: 1,137$ $V_{col}: 3,01$ $y_{col}: 4,03$	Collision with intervention	 $TTC+\Delta t: 1,120$ $V_{col}: 3,34$ $y_{col}: 4,07$

Fig. 13. Shown are five consecutive test runs. The upper box displays the situation before system intervention. Shown is the x distance [m] to the target in host vehicle coordinates as well as relevant host ($v_{x_{host}}$ [m/s]) and target ($v_{y_{target}}$ [m/s]) speeds. y [m] gives the distance traveled by the target in the remaining TTC. To compensate for the y offset during approach by the demonstrator the target was linked via radio to the VeHIL controller to adjust its acceleration at start. TTC [s] denotes the calculated time to collision without intervention. The middle box displays the scenario at collision without intervention as calculated by the system. The offset from the middle of the host car to the front of the target is given as y_{col} [m]. As the start velocities differ slightly in every run this value changes accordingly. The impact velocity of the host is v_{col} [m/s]. The lowest box displays the scenario the system tries to achieve by intervention. $TTC+\Delta t$ [s] is the new time until impact, v_{col} [m/s] the reduced impact speed and y_{col} [m] the resulting impact location.

intervention. The overall accuracy in impact location was within the required 20.3 cm in three of five consecutive tests.

VIII. CONCLUSIONS

An active intervention by braking has to take into account the resulting change of impact configuration. This paper presents a novel approach to assist the driver in crossing traffic. Intervention by braking is adjusted to either hit the front or rear axle without any overlap to the passenger compartment thereby avoiding the risk of increased injuries when impacting to the side of the passenger compartment.

A low cost system for wheel detection was implemented using a wide field of view camera for automotive systems. The system was validated by a full scale test track experiment. The results demonstrated the ability to prevent collisions with the passenger compartment. A driver survey showed that driver intervention below 700 ms prior to impact is only conducted by braking and such a system is considered helpful in mitigating the crash severity.

In its current form, the system only addresses crossing cars as collision targets. Classification of crossing objects could lead to enhanced intervention strategies that address e.g. collisions with truck trailers. Further work will enhance the video based system with new approaches like optical flow which will help achieve accurate car outline segmentation.

ACKNOWLEDGMENT

This work was supported by the European Commission under interactIVe, a large scale integrated project part of the FP7-ICT for Safety and Energy Efficiency in Mobility. The authors would like to thank all partners within interactIVe for their cooperation and valuable contribution. Special thanks go to D. Wilhelmssen and J. Vissers at TNO for their support with the VeHIL tests and to E. Wykowski and T. Wohllebe at Volkswagen Group Research.

REFERENCES

- [1] European Commission, "Towards a European road safety area: policy orientations on road safety 2011-2020," SEC(2010) 903, 2010.
- [2] S. O'Brien, "Measurement and Assessment of Passenger Vehicle Compatibility in Front and Side Collisions," Ph.D. dissertation, RMIT University, Melbourne, 2010.
- [3] PReVENT project, "Deliverable D40.4 Requirements for intersection safety applications," 2005.
- [4] J. Chen, S. Deuschle, and K. Fürstenberg, "Evaluation methods and results of the INTERSAFE intersection assistants," in *Proc. IEEE Intelligent Vehicles Symposium*, 2007, pp. 142–147.
- [5] K. Fürstenberg and B. Rössler, "Advanced intersection safety - the EC project INTERSAFE," in *Proc. IEEE Intelligent Vehicles Symposium*, 2006, pp. 89–93.
- [6] M. Brännström, E. Coelingh, and J. Sjöberg, "Decision making on when to brake and when to steer to avoid a collision," in *Proc. International Symposium on Future Active Safety Technology toward zero-traffic-accident*, 2011.
- [7] J. Kim, J. Lee, and E. Kim, "Technical feasibility study on avoidance or mitigation of side collisions at intersections," SAE Technical Paper 01-0090, 2012.
- [8] S. Herrmann and F. Schroven, "Situation analysis for driver assistance systems at urban intersections," in *Proc. IEEE International Conference on Vehicular Electronics and Safety*, 2012, pp. 151–156.
- [9] S. Lefevre, C. Laugier, and J. Ibanez-Guzmán, "Risk assessment at road intersections: Comparing intention and expectation," in *Proc. IEEE Intelligent Vehicles Symposium*, 2012, pp. 165–171.
- [10] J. Klappstein, F. Stein, and U. Franke, "Detectability of moving objects using correspondences over two and three frames," in *Proc. DAGM conference on Pattern recognition*, 2007, pp. 112–121.
- [11] I. Sato, C. Yamano, and H. Yanagawa, "Crossing obstacle detection with a vehicle-mounted camera," in *Proc. IEEE Intelligent Vehicles Symposium*, 2011, pp. 60–65.
- [12] P. Viola and M. J. Jones, "Robust real-time face detection," *International Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154, 2004.
- [13] J. Cho, S. Mirzaei, J. Oberg, and R. Kastner, "FPGA-based face detection system using Haar classifiers," in *Proc. ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, 2009, pp. 103–112.
- [14] C. Gao and S.-L. Lu, "Novel FPGA based Haar classifier face detection algorithm acceleration," in *Proc. International Conference on Field-Programmable Logic and Applications*, 2008, pp. 373–378.
- [15] M. Papadonikolakis and C. Bouganis, "Novel cascade FPGA accelerator for support vector machines classification," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 23, no. 7, pp. 1040–1052, 2012.
- [16] C. Kyrkou, C. Ttofis, and T. Theocharides, "FPGA-accelerated object detection using edge information," in *Proc. International Conference on Field-Programmable Logic and Applications*, 2011, pp. 167–170.
- [17] Economic Commission for Europe, "Vienna Convention on Road Traffic," pp. Article 8, Paragraph 5, 1968.