Traffic Awareness Driver Assistance based on Stereovision, Eye-tracking, and Head-Up Display

Tobias Langner[1], Daniel Seifert[1], Bennet Fischer[1], Daniel Goehring[1], Tinosch Ganjineh[2], and Raúl Rojas[1]

***Abstract*— This paper presents a system which constantly monitors the level of attention of a driver in traffic. The vehicle is instrumented and can identify the state of traffic-lights, as well as obstacles on the road. If the driver is inattentive and fails to recognize a threat, the assistance system produces a warning. Therefore, the system helps the driver to focus on crucial traffic situations. Our system consists of three components: computer vision detection of traffic-lights and other traffic participants, an eye tracking device used also for head localization, and finally, a human machine interface consisting of a head-up display and an acoustic module used to provide warnings to the driver. The orientation of the driver's head is detected using fiducial markers visible in video frames. We describe how the system was integrated using an autonomous car as experimental ADAS platform.**

Fig. 1. Our testbed, the autonomous car "MadeInGermany".

I. Introduction

A significant percentage of traffic accidents is a result of driver drowsiness or driver unawareness [2]. Many new cars are entering the road every year: in 2010 the number of vehicles worldwide crossed the one billion mark [17]. According to the National Highway Traffic Safety Administration (NHTSA), falling asleep while driving causes approximately one hundred thousand car crashes in the United States, about 40,000 non-fatal, and more than 1500 fatal injuries [15] yearly. This paper addresses these issues and describes a driver assistance system which:

1) Recognizes traffic-lights and vehicles at intersections using a self-developed stereo-camera and FPGA for fast 3D preprocessing.
2) Estimates the gaze direction of the driver and the position of his head relative to the car, checking if the driver fails to notice critical traffic conditions (for example a red traffic-light).
3) Warns the driver using a self-developed head-up display (HUD) and acoustic signals (e.g. speech) when he or she fails to notice a red traffic-light or vehicles arriving or waiting at intersections.

The system was developed using our autonomous vehicle as a testbed (Fig. 1).

A. Information Data-Flow

The data-flow between the system modules is shown in Fig. 2. The stereo camera detects the traffic-lights and vehicles on the road relative to the car. For improved performance we have been using road maps with annotated intersections in order to avoid false positive detections of traffic-lights. The gaze direction estimation module is based on wearable glasses which detect the gaze direction of the user using small video cameras pointed to the iris. The main module uses the information from the vision and gaze modules to determine if the driver has seen or not other vehicles and traffic-lights. If necessary, a warning message or symbol is shown in the HUD, and an acoustic warning is played.

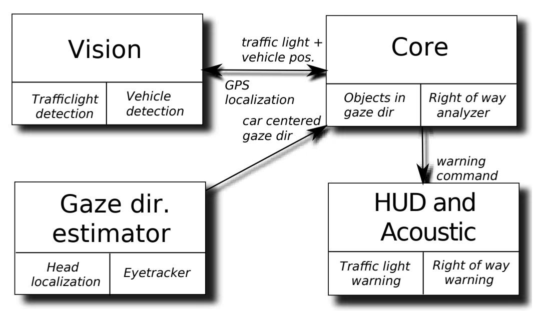

Fig. 2. Module Diagram which shows the data flow between traffic-light and vehicle detector, gaze direction estimator, HUD and acoustic output, and core module.

B. Related Work

Strong and reliable sensor fusion methods are crucial for driver assistance systems as well as for self driving cars. An important subtask is to detect other vehicles, traffic lights and road signs. An early version of a shape based road-sign detection was presented in [13] or, using two-layered neural networks in [4]. Monocular vision for autonomous driving has been successfully applied by Dickmanns et al.

*This work was funded by the German Ministry of Education and Research (BMBF)

[1]Intelligent Systems and Robotics, Computer Science Institute, Freie Universität Berlin, Germany {tobias.langner, daniel.seifert, b.fischer, daniel.goehring, raul.rojas}@fu-berlin.de

[2]Autonomos GmbH, Berlin, Germany {tg@autonomos-systems.de}

in [3]. Another interesting approach combining lidar and 2D-camera data using HOG-based object recognition has been presented in [11]. A method for recognizing traffic-lights was described in [18]. A traffic-light detection method able to recognize left and right turn traffic-lights was presented by Google in [19]. A boosting framework for image-based classifiers was developed by Viola and Jones in [21], an approach which we have been using for vehicle detection. Much research effort has been spent in the detection of driver drowsiness and fatigue, e.g., in [10]. There have been many applications of head-up displays [1] for avionics. Such HUD systems are now conquering the market for assisted driving solutions. Much research has been performed on gaze direction detection with different sensors and for a variety of applications. Land and Lee did research on where a driver looks while entering a curve [12], finding a driver looks on the tangent of the road ahead. A system, which checked if a driver had missed or perceived speed limit signs was presented in [5]. In [20] a vision based eyes of the road detection system was introduced. The work from Ziraknejad et al. [24] tries to detect head positions using capacitive sensors for a better adaptation of head restraint systems. The work from Skodras et al. [16] uses a low-cost webcam to detect eye-movements. In [14] a solution which applies heat-maps and optical flow methods for a highly precise attention estimation was presented. The vigilance of a driver has been widely been within the focus of research, as in [9] showing how driver drowsiness could be measured. In [2] an IR-camera was used to measure eye closing duration, as well as eye blinking and nodding frequencies, or in [23] using time-of-flight cameras.

The paper is structured as follows: Section II presents the stereo camera system used for the detection of traffic-lights and other vehicles. The algorithm for the detection of traffic-lights is described in Section III. The method of vehicle detection is explained in Section IV. Section V presents the eye-tracking hardware used. We explain our head localization approach using fiducial markers and/or infrared LEDs installed in the cockpit. Section VI describes our self-developed head-up display. We present experimental results for the different modules, specially concerning recognition rates in Section VII, before providing a summary of the paper.

II. Our Stereo-Camera

Our robotic testbed has been equipped with a smart stereo camera used to perceive the 3D structure of the environment and to detect traffic-lights (Figure 3). Our camera, developed for our previous AutoNOMOS project, processes the images on-board using an FPGA which delivers 3D point clouds to an on-board application processor. Disparities are processed at a resolution of 752×480 pixels at 30 frames per second. An example of the disparity map produced by our camera is shown in Figure 3. The camera is well suited to mobile robotics due to its low power consumption of $1.5\,\mathrm{W}$. Its small size and its HDR global shutter sensors are advantageous for our application. The camera has been utilized in several vehicles ranging in size from model cars to full-size trucks.

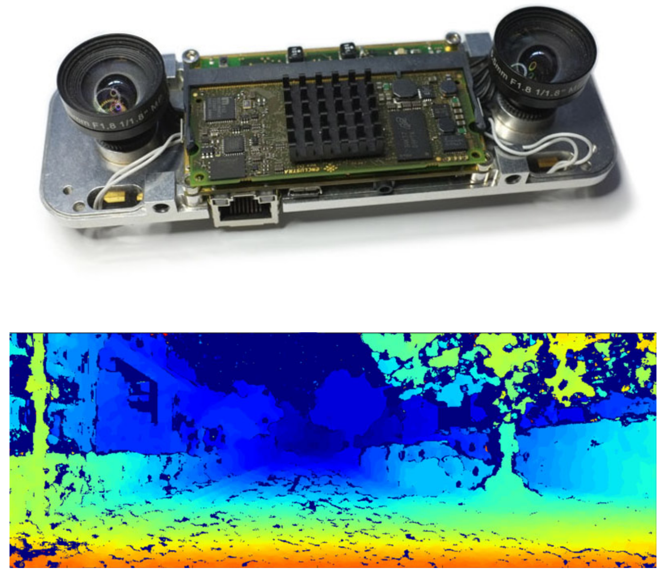

Fig. 3. Smart stereo camera and disparity image. Top: Smart stereo camera. Bottom: Example image processed by the camera's FPGA (image #8 of KITTI benchmark [7])

III. Traffic-Light Detection

Fig. 4. Two tracked traffic-lights and a detected vehicle.

A. Pixel Classification

We decided to use the HSV color space for pixel classification because it can separate luminance and chroma. An automatic analysis of images in a training set generated the thresholds for each color class. Fig. 5 shows the hue distribution of a few traffic-light samples. Obviously, the red and yellow distributions overlap, but have different modes. A probabilistic approach is feasible, assigning each hue a probability of being red or yellow. For our purposes (awareness detection) we decided to create one common

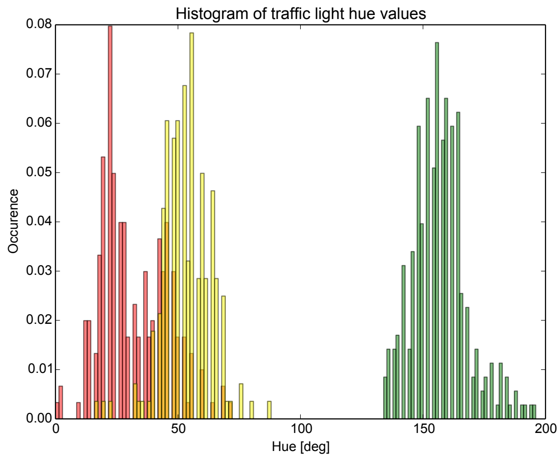

Fig. 5. Distribution of hue angles for red, yellow and green traffic-lights.

	Hue [deg]	Saturation [%]	Value [%]
Red-Yellow	0-70	0.4-1.0	0.6-1.0
Green	140-180	0.4-1.0	0.6-1.0
Dark	-	0.0-0.3	0.0-0.3

TABLE I

PIXEL COLOR CATEGORIES ACCORDING TO RANGES OF HUE, SATURATION, AND VALUE.

color class for red and yellow and deal with the exact color at a later processing stage (see below). Such hard color labeling simplifies the subsequent connected component analysis. Together with color, we also consider the 3D position of each pixel. In fact, most color-matching pixels can be rejected very early because of their implausible height or lateral distance to the street. In particular, we exploit the fact that in Germany the exact positions of traffic-lights are defined by the RiLSA standard (Richtlinien für Lichtsignalanlagen). Under this directive, the base of a hanging traffic-light frame lays at $4.5\,\mathrm{m}$. Signals mounted to a pole are positioned at $2.1\,\mathrm{m}$ above the ground. We assign a class label (*Red-Yellow*, *Green*, *Dark* or *Background*) to each pixel and pass the classified image through a connected components calculation.

B. Colored Connected Components

Only the color classes *Red-Yellow* and *Green* are considered for the connected component analysis. The label *Dark* is useful for the unlit parts of a traffic-light, but those pixels may not be distinguishable from the background. Especially at night, the segmentation of the traffic-light frame can fail. Therefore, the *Dark* pixels are used instead to verify if regions adjacent to light sources appear indeed dark. The connected components are built applying region growing to the color labels, whereby two pixels of the same color class in an 8-neighborhood are considered connected. Additional heuristics are applied to discard components of irregular shape, low density, or inconsistent depth.

C. Construction of Traffic-Light Entities

The detected circular regions of colored lights are used to construct complete traffic-light entities consisting of three (lit or unlit) circular light sources on a rectangular dark-colored frame. The length ratios between the lights and the frame are known. Thus, we use the measured size, color and 3D distance of the detected light source to determine the area where the two other unlit light sources should be found. We count the number of pixels classified as *Dark* within those regions, and compute their density. If the density is larger than a threshold (i.e. 75%), and the average 3D distance is comparable to that of the light source, the region will be accepted. At this point the ambiguity between red and yellow traffic signals can be resolved, because they have a different position.

D. Tracking and State Output

The positions of all detected traffic-lights are tracked over time. We use a recursive filter which performs alternate prediction and measurement updates. The prediction step displaces a tracked traffic-light under the assumption of constant pixel velocity to a new image position. Then new measurements and predicted traffic-lights are matched with respect to their distances in the image. Two typical trajectories can be seen in Fig. 4. The final traffic-light state is determined by a majority vote between all tracked lights. In case of a split vote, the closest semaphore wins. Finally, the state forwarded to the system is smoothed by a hysteresis filter requiring three consecutive identical states for triggering a state change. The possible states are *Stop*, *Go*, and *Unknown* (i.e. no traffic-light detected).

IV. VEHICLE DETECTION

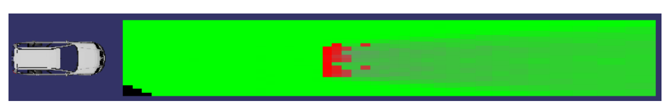

Fig. 6. The occupancy grid computed for the scene shown in Fig. 4. Green cells mark free space and red cells mark occupied space. The gray shaded cells are (partially) occluded by the obstacle.

A. Occupancy Grid Computation

Our stereo camera provides 3D point clouds, that we use for obstacle detection. The segmentation of objects from point clouds can be a challenging task. Thus, we decided to simplify the problem by computing a 2D histogram of the point cloud. Each cell covers a small area in front of the car and accumulates the 3D points being in a volume above that cell. In our case, only points within a height range of $0.5\,\mathrm{m}$ to $1.5\,\mathrm{m}$ are considered. Depending on the number of points every cell is labeled either occupied or free, hence the name occupancy grid.

B. State Decision for Grid Cells

Due to perspective projection close grid cells usually contain more points than far away cells. Thus, we cannot mark a cell as occupied by inspecting the absolute numbers, but instead have to calculate the degree of occupancy. For this purpose, we compute a priori the projection of each cell volume onto the image. This yields the upper bound for the number of points a cell can hold, allowing us to compute the degree of occupancy. Any cell with a ratio larger then 10% is marked as *Occupied*, while the rest is labeled as *Free*.

C. Obstacle Segmentation

Subsequently, we apply region growing to merge adjacent occupied cells into objects. Fig. 6 shows the occupancy grid for the scene depicted in Fig. 4. The measured distance to the car is the one to the closest cell. Tracking the distance over time yields information about the relative speed and the time to collision.

V. Estimation of Gaze Direction

We use commercially available eye tracking glasses built by SMI in order to detect the driver's focus of attention (Fig. 7). Each eye is captured by an infrared camera and the direction of the user's gaze relative to the glass mount is calculated by the vendor's software. To determine the driver's focus of attention, we need to know the position of the driver's head in order to transform the local gaze direction to the car's coordinate system.

Fig. 7. The eye tracking glasses produced by SMI. A mini camera is pointed away from the frame (as in a Google Glass frame). Two mini cameras are inside the frame and are pointed towards the eyes of the user.

One approach for tracking the position of the head is to use an external fixed camera for observing the driver, who needs to wear special markers. Both the observation of the driver by a camera and the need to wear the markers are impacting the acceptance of the system. We selected a different approach instead, where we are relying on the glasses' embedded *scene camera*. This mini camera is located above the nose and is pointing forward and slightly down. Image processing is used to determine the head orientation based on observed features inside the cockpit.

A. Infrared LED Marker Detection

Our first method for tracking the head orientation was to use infrared LEDs. Some 880 nm LEDs are placed at predefined positions along the windshield border. A bandpass filter (85% transmission, 10 nm FWHM) blacks out most irrelevant parts of the image. The LEDs were sufficiently bright to be visible in the image even at the lowest exposure setting, and a simple threshold was applied to further reduce noise. The remaining white blobs in the resulting black-white image were then extracted. A heuristic allowed us to filter out sunlight and most reflections in the remaining infrared band around 880 nm. Fig. 8 depicts our prototype board, the scene view with and without the bandpass filter, as well as the final detection result. Based on the position of the LEDs comprising the predefined pattern, we used OpenCV's solvePNP[1] function to determine the camera's position. Initial experiments using a web cam with a fixed, low exposure produced promising results. Unfortunately the embedded scene camera in the glasses only supported auto-exposure at the time of the experiments. Due to the bandpass filter, the camera images were usually quite dark, resulting in a high exposure and corresponding motion blur which prevented us to achieve a precise detection of the LEDs. An update to the glasses' software was not available at time which is why we switched to another approach to infer the head orientation of the driver.

B. Fiducial Marker Detection

To overcome the described challenges, we pursued an alternate approach using fiducial markers (Fig. 9). Numerous software solutions exist to create appropriate markers and locate them within images. Although they are commonly used for augmented reality applications, they can be also used to determine the position of a camera given the known location of the markers.

We selected the ArUco[2] library based on [6] for this task. Three markers were placed in the car: One slightly left of the driver, one in the center of the car's dashboard, and one on the right. The field of view of the camera and the position of the markers ensures that at least one marker is visible most of the time. This is sufficient to determine the position and rotation of the scene camera within the vehicle's coordinate frame.

C. Calibration

Several calibration steps are required to achieve sufficient precision for gaze direction estimation.

The eye tracking glasses require a manual calibration to adjust for differences in the driver's eyes anatomy. For this, the vendor's software provides one- and three-point calibration which we integrated in our software. The driver focuses attention towards a specific point which then is selected in the calibration software. Differences between the focus model and the actual focus will be used to calibrate the model to the driver's eyes.

The glasses come pre-calibrated to transform the infrared eye tracking camera coordinate system to the scene camera.

[1]`http://www.opencv.org`
[2]`http://www.uco.es/investiga/grupos/ava/node/26`

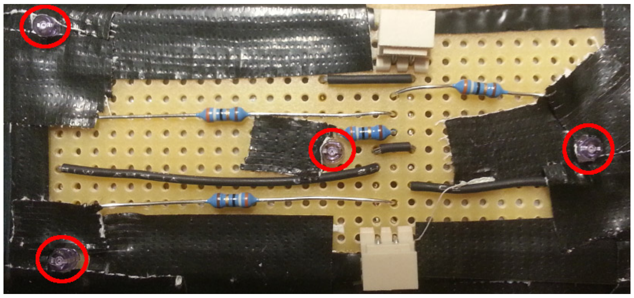

(a)

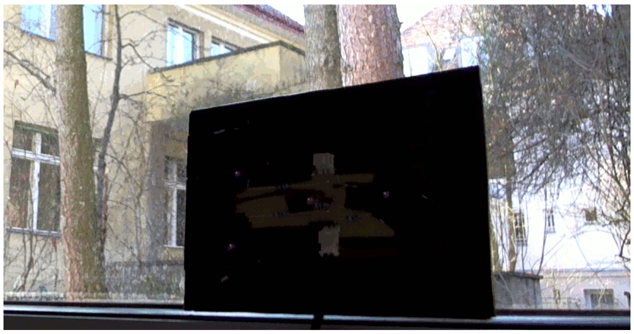

(b)

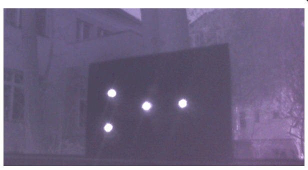

(c)

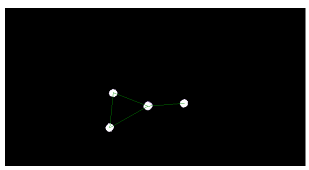

(d)

Fig. 8. Detection of the infrared marker, consisting of four 880nm infrared LEDs. (a) Close-up of the prototype marker board with the infrared LEDs highlighted. (b) The marker board in an experimental setup with a sunlit background. (c) The board recorded with the infrared bandpass filter. The LEDs are significantly brighter than the sun-lit background. (d) The output of the detection filter.

For our marker detection, we use OpenCV's camera calibration function to determine the intrinsic parameters of the scene camera.

Finally, it is necessary to determine the position of the markers in the car coordinate system. For this, we calibrate the location of the central marker in the car by means of an external calibration pattern which position is known by other means. This marker is fixed, allowing to calibrate its position and rotation once. For our test setup, the other two markers are positioned in a way that allows the scene camera to capture them together with the central marker. This provides the opportunity to calibrate the other markers' position during runtime. For increased accuracy, all markers will be calibrated individually.

VI. Head-Up Display

The goal of our system is to increase the driver's awareness of hazardous situations using technical aids. Therefore, it is of high importance for the user interface to be as unobstrusive as possible.

Experiments have shown that glances at car displays can take up to two seconds [22]. At city driving speeds, a car travels up to 30 meters during this period without driver attention to the road. Even when the driver looks back to the road the response time to a possible hazardous situation has to be added to the glance duration. The loss of road attention caused by car displays is caused by:

- Location: In-car displays are mounted on the dashboard. Glancing usually requires to turn the head.
- Distance: In-car displays are close to the driver. Therefore, the focus of the driver's eyes has to adapt to the distance.
- Brightness: Illumination is usually brighter outside than inside the car. A glance to the car display requires the iris to adapt.

The eyes have to adapt twice: When viewing changes from the road to the display, and vice versa.

Head-up displays are a well-known approach for reducing the attention problems of car displays. Usually, HUDs are implemented with a light projector. Its image is reflected by the windshield so that the information appears in the normal field of view of the driver. Since in this case the location, distance, and brightness do not change for the eyes, the glance duration is lower compared to a car display in the instruments panel [8].

Most devices available today are active systems, based on LED technology. Active systems have to be adapted to the environment, for example, during the night the image needs to be dimmed. To overcome the brightness problem and to provide a natural viewing experience, we propose a passive system that uses reflective technology. Figure 10 shows our head-up display. It was built using a 10 inch e-paper display and a custom e-paper controller.

The device is mounted on the dashboard of the car as seen in Figure 11. Images shown on the e-paper display are reflected by the windshield. This projects the image onto the road in front of the car which is in the natural field of view of the driver. To avoid a direct line of sight to the display a small barrier with a fiducial marker has been added to the device.

To overcome the slow update rate of e-paper, we developed a custom e-paper controller based on an STM32F4 ARM microcontroller. This allowed us to preload urgent messages like a hazard warning into the e-paper's memory without displaying them. Upon detection of a hazardous situation the

Fig. 9. Left: Fiducial marker example. Right: two markers inside the cockpit.

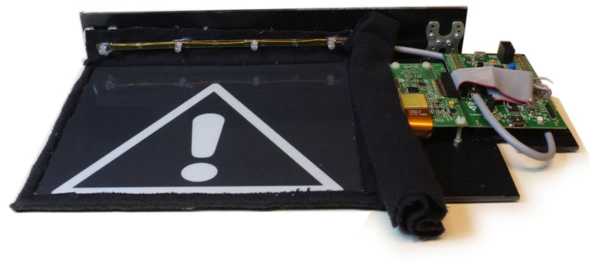

Fig. 10. Head-up display device. Left: Screen using e-paper technology. Right: Custom e-paper controller

Fig. 11. Head-up display seen from the driver's point of view. The display uses LED illumination during nighttime

e-paper can be enabled and displays the preloaded content with low latency.

VII. Experiments

A. Traffic-Light Detection Results

We obtained detection rates of above 95% in our experiments, which we want to improve, but is already significant and useful for a warning system. For our application it is also important to detect traffic signals as early as possible. As can be seen in Fig. 12 the typical detection range is at 40 m. Given the inner-city speed limit of 50 km/h, this gives the driver a reaction time of 2.8 s in case of a warning.

B. Gaze Direction Estimation

Lacking reasonable ground truth, we conducted a simple experiment to evaluate the accuracy of the gaze direction estimation. During a test drive in city traffic we asked the driver to focus certain points in the scene and compared it to the direction computed by our system. We observed deviations of up to 3 degrees, which is still sufficient for our purposes. To compensate for the errors, we enlarged the computed line of sight to a cone of 10 degrees aperture. Any object within this cone is assumed to have drawn the driver's attention.

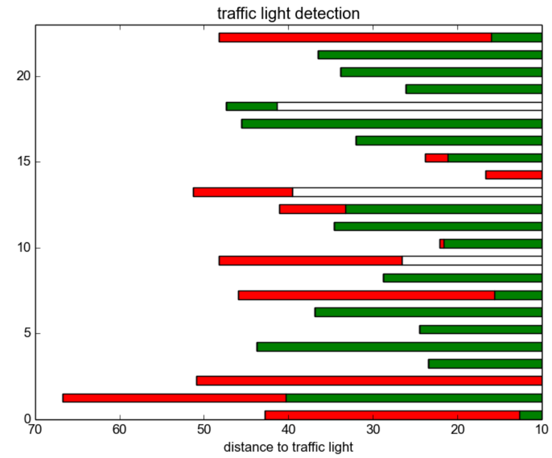

Fig. 12. Detection ranges for a sequence of 23 traffic-lights (represented by 23 bars). The distance to the traffic-light in meters is shown on the horizontal axis. The traffic-light is situated at the right end. The ranges start at 10 meters because any closer traffic light is beyond the camera's field of view. The color of the bar represents the state of the traffic-light. Some of the lights suddenly disappear because the car has turned left before reaching them. Traffic-light #3 (from the bottom) left the camera's field of view before it turned green.

In another simple test, we measured the standard deviation of the detected marker rotation for a non-moving camera outside the car. The rotation angles (roll, pitch, yaw) can differ by up to 0.2 degree. Hence, we assume the overall inaccuracy is caused by other factors.

C. Overall Results

Fig. 13 shows a driver using our system. The direction of his gaze is detected by the eye-tracker, the fiducial markers allow us to map the eye-direction to the frame of the car. Fig. 14 shows the 3D point cloud produced by our stereoscopic camera and the cone of attention of the driver, determined by his gaze. It is then possible to detect if the driver has not yet seen a red light or a car coming from the right. In that case a warning is given. The vehicle does not react to these events, all responsibility remains with the driver. We are now in the process of testing the system in the streets of Berlin. A group of selected drivers will be asked to drive in the city and the quality of the warnings will be assessed by a driving instructor.

VIII. Conclusions

We have described a system which allows the detection of a driver's awareness or unawareness of critical traffic situations. In case of distraction, the system helps the driver to guide his attention to the traffic-lights or to oncoming vehicles. More research needs to be done to improve the accuracy, reliability and practicality of the proposed approach. Variable lightning is a challenge for camera based ADAS – further research needs to be done in order to guarantee optimal performance during day or night. Regarding the vision sensor used, it is usually necessary to collect depth-data for objects 50 or 100 meters away. This requires higher resolution and/or

Fig. 13. Whole systems experiment. The driver is wearing the eye-tracker glasses, a camera in the glasses is perceiving the fiducial markers and localizes the head relative to the car.

Fig. 14. The depth and image information of the stereo camera is projected as a colored point cloud into the 3D-space with respect to the car. The gaze direction is depicted as the semi-transparent cone, starting from the driver's seat.

greater disparity between the stereo cameras. For practical purposes, an eye-tracking device which does not require wearing special glasses would provide optimal comfort and the best use-case.

Acknowledgements

The authors would like to thank the German Ministry of Education and Research (BMBF) for funding this project. We would also like to thank Sensomotoric Instruments (SMI) in Teltow, Germany, for their support in this project.

References

[1] K. Asai. *The Role of Head-Up Display in Computer- Assisted Instruction, Human Computer Interaction: New Developments*. 2008.

[2] L. Bergasa, J. Nuevo, M. Sotalo, and M. Vazquez. Real-time system for monitoring driver vigilance. In *Proceedings of the IEEE Intelligent Vehicle Symposium*, pages 78–83, 2004.

[3] E. Dickmanns and V. Graefe. Applications of dynamic monocular machine vision. *Machine Vision*, 241–261, 1988.

[4] C. Fang, C. Fuh, S. Chen, and P. Yen. A road sign recognition system based on dynamic visual model. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recog- nition*, volume 1, 2003.

[5] L. Fletcher, G. Loy, N. Barnes, and A. Zelinsky. Correlating driver gaze with the road scene for driver assistance systems. *Robotics and Autonomous Systems*, 52:71–84, 2005.

[6] S. Garrido-Jurado, R. Muñoz-Salinas, F. J. Madrid-Cuevas, and M. J. Marín-Jiménez. Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition*, 47(6):2280–2292, 2014.

[7] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.

[8] H. Hada. Drivers'visual attention to in-vehicle displays: Effects of display location and road types. Technical report, 1994.

[9] N. Haworth, T. Triggs, and E. Grey. Driver fatigue: concepts, measurement and crash countermeasures. Technical report, Federal Office of Road Safety Contract Report 72 by Hu- man Factors Group, Department of Psychology, Monash Uni- versity, 1988.

[10] W. B. Horng, C. Y. Chen, Y. Chang, and C. H. Fan. Driver fatigue detection based on eye tracking and dynamic template matching. In *Proceedings of the IEEE International Conference on Networking, Sensing and Control*, pages 7–12, 2004.

[11] Hyunggi Cho, Young-Woo Seo, B. Vijaya Kumar, and Ragunathan (Raj) Rajkumar. A Multi-Sensor Fusion System for Moving Object Detection and Tracking in Urban Driving Environments. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2014.

[12] M. Land and D. Lee. Where we look when we steer. *Nature*, 369:742–744, June 1994.

[13] G. Loy and N. Barnes. Fast shape-based road sign detection for a driver assistance system. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2004.

[14] R. Perko, M. Schwarz, and L. Paletta. Aggregated mapping of driver attention from matched optical flow. *2014 IEEE International Conference on Image Processing*, pages 214–218, January 2014.

[15] D. Royal. *Volume I - Findings report; national survey on distracted and driving attitudes and behaviours, 2002*. DOT HS 809 566. The Gallup Organization, Washington, D.C., 2003.

[16] E. Skodras, V. Kanas, and N. Fakotakis. On visual gaze tracking based on a single low cost camera. *Signal Processing: Image Communication*, 36:29–42, August 2015.

[17] J. Sousanis. World vehicle population tops 1 billion units, August 2011.

[18] H. Tae-Hyun, J. In-Hak, and C. Seong-Ik. Detection of traffic lights for vision-based car navigation system. *Advances in Image and Video Technology*, Lecture Notes in Computer Science(4319):682–691, 2006.

[19] N. F. C. Urmson. Traffic light mapping and detection. Technical report, Google Research, 2011.

[20] F. Vicente, Z. Huang, X. Xiong, F. D. la Torre, W. Zhang, and D. Levi. Driver gaze tracking and eyes off the road detection system. *IEEE Transactions on Intelligent Transportation Systems*, 16(4):2014–2027, August 2015.

[21] P. Viola and M. Jones. Rapid object detection using a boosted cascade of simple features. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages I511–I518. Cambridge, Mass, USA, 2001.

[22] A.-S. Wikman, T. Nieminen, and H. Summala. Driving experience and time-sharing during in-car tasks on roads of different width. *Ergonomics*, 41(3):358–372, 1998.

[23] N. Ziraknejad, P. Lawrence., and D. Romilly. The effect of time-of-flight camera integration time on vehicle driver head pose tracking accuracy. In *IEEE International Conference on Vehicular Electronics and Safety*, pages 247–254, 2012.

[24] N. Ziraknejad, P. Lawrence, and D. Romilly. Vehicle occupant head position quantification using an array of capacitive proximity sensors. *IEEE Transactions on Vehicular Technology*, 64(6):2274– 2287, June 2015.