Proceedings of the
2007 IEEE Intelligent Vehicles Symposium
Istanbul, Turkey, June 13-15, 2007

The BMW SURF Project: A Contribution to the Research on Cognitive Vehicles

Stefan Hoch and Manfred Schweigert, *BMW Group Research and Technology*
Frank Althoff*, BMW Group*
Gerhard Rigoll, *Technical University of Munich*

***Abstract*— Within the last years, research in the field of cognitive vehicles has evolved to an important activity especially for car manufactures in the premium segment. Motivated by the robustness and flexibility of inter-human communication, highly integrated contextual awareness is being seen as a key component to provide an intuitive human-machine interface to the wide range of functions the driver is being confronted with in a modern driver's working place. In this work, we propose a highly flexible software architecture to develop and evaluate various services in the context of cognitive vehicle behaviour. The system has been implemented in an experimental car that is equipped with a wide range of sensors to acquire information from three sources: the car, the environment and the driver. Special efforts have been made to evaluate the performance and the potential of different driver monitoring approaches for the use in adaptive automotive applications, in both advanced driver assistance systems (ADAS) and in-vehicle information systems (IVIS). To evaluate the performance of our approach in detail, we have chosen an adaptive lane departure warning system. Different strategies for driver intention analysis are compared with regard to several criteria in a small evaluation study. The results clearly shows the importance and the potential of a complex contextual analysis for the emergence of future intelligent vehicle applications**.

I. Introduction

Vehicles with cognitive capabilities have increasingly become part of research topics of car manufacturers. Providing their customers with a maximum amount of security, luxury and comfort, future cars will be equipped with a wide range of advanced driver assistance (ADAS) and in-vehicle information systems (IVIS), which will have to react properly due to the current situation requirements and the drivers needs.

Today, this adaptive automotive system behaviour is mainly limited by the non-availability of context parameters. The development of context aware applications is often restricted by missing information about the behaviour of the driver and the complex traffic situation, leading to so-called nuisance alarms and a reduced system acceptance by the user [1]. But an increasing amount of contextual information provided by an emerging range of sensor technologies gives the possibility to obtain a more detailed view of in-car system status and the environmental situation. In addition to these "vehicle-related" information, driver monitoring sensors can yield new aspects about the driver's behaviour, which are left out in most actual systems. The knowledge of the complete situational context consisting of vehicle, environment and driver information offers extended possibilities for the development of enhanced automotive applications, both ADAS and IVIS, provided with parts of artificial intelligence and adapting to the situational context.

Cognitive and adaptive vehicle behaviour demands an interpretation of the contextual situation. The goal of the BMW Group Research and Technology project SURF (Situation and Driver Adaptive Functions) has been to build a prototype vehicle equipped with multiple sensors providing contextual information and a flexible software architecture for the development and evaluation of adaptive applications. Furthermore we have planned to explore the possibilities of advanced contextual information provided by driver monitoring sensors as well as an intelligent combination of all contextual information fields – environment, car and driver – with machine learning and artificial intelligence tools. We have decided to develop and implement a situation and driver adaptive lane departure warning system (LDWS) as a test scenario. We have chosen an ADAS application, which is already available on the market, to evaluate the impact in both intelligent system design and application transparency for the driver by introducing enhanced cognitive vehicle capabilities.

This paper gives an explicit overview about our system

Manuscript received April 15, 2007.
Stefan Hoch is with BMW Group Research and Technology, Hanauerstrasse 46, 80992 Munich, Germany. (Stefan.Hoch@bmw.de)
Dr. Manfred Schweigert is with BMW Group Research and Technology, Hanauerstrasse 46, 80992 Munich, Germany. (Manfred.Schweigert@bmw.de)
Dr. Frank Althoff is with BMW Group, Knorrstrasse 147, 80788 Munich , Germany. (Frank.Althoff@bmw.de)
Professor Dr.-Ing. habil. Gerhard Rigoll is Head of the Institute for Human-Machine Communication, Technical University of Munich, Theresienstraße 90, 80333 Munich, Germany. (Rigoll@tum.de)

software and hardware design and summarizes our experiences and first evaluation results with our cognitive vehicle.

II. RELATED WORK

Since the year 2000 BMW has promoted the development of active safety applications with its "ConnectedDrive" paradigm, which is based on the key idea of a symbiotic combination of driver, vehicle and environment information [2] [3].

The following chapter gives a brief excerpt of further related work on the central topics of our research project following the "ConnectedDrive" idea. We believe that the development of future cognitive vehicles has to deal with all the aspects, namely context awareness and adaptive interfaces, driver intention recognition and driver monitoring.

Several EU-co-funded projects have dealt with the idea of a context analysis component. The Project COMUNICAR (Communication Multimedia Unit Inside Car) proposed an Information Manager (IM), which collects the feedback information of assistance-, telematics- and entertainment functions. The IM estimates the driver's handicap resp. workload according to the impact of the current driving and environment situation. The derived information is combined with the priority of the information display requests to adapt the overall system-feedback to enable a safe interaction between the driver and a multitude of in-vehicle systems [4].

A maximization of the efficiency and the safety benefits of advanced driver assistance systems while minimizing the level of workload and distraction imposed by in-vehicle information systems are two main goals of the EU-Project AIDE (Adaptive Integrated Driver-vehicle InterfacE). The project architecture is based on a central high level perception instance, called the DVE, to sense and estimate the state of the driver, the vehicle and the environment. The DVE offers the contextual information to a Communication Assistant, which then controls a secure display of applications feedback information to the driver through the vehicle human-machine-interface [5].

The NHTSA project SAVE-IT (Safety Vehicles Using Adaptive Interface Technologies) has focussed on the development, demonstration, and evaluation of potential safety benefits of adaptive interface technologies that manage the information from various in-vehicle systems. This management process is based on a central real-time monitoring of the roadway conditions, vehicle's and driver's state parameters [6].

While the results and experiences of these projects have shown that there is a lot of potential in the idea of a central information processing unit, we think that context awareness is also strongly connected to the topic of knowledge analysis. Therefore we have decided to develop our own interpretation of a shared knowledge management instance with the focus of a flexible framework offering both powerful knowledge analysis services in the development phase and a generic usage of derived knowledge parameters in the application runtime phase.

Driver intention analysis has been a key component in this context analysis concept. The idea of a driver intention recognition and also driver monitoring components for adaptive user interfaces and ADAS have also been addressed in several projects. As part of the European PROMETHEUS program, Geiser and Nirschl [7] have suggested that the Driver Warning System (DWS) should analyse intention information for suppressing premature warnings. Likewise, in the United States the NHTSA Benefits Working Group [8] has recommended that an ADAS like a collision avoidance system should be intelligent enough to discern driver's intention (e.g., intent to change lanes, lane change start) though this is difficult to do. If available, such indicators might selectively alter the drive alerts or warnings" (e.g., thresholds, presentation mode, stimulus magnitude, etc.) and reduce the number of nuisance alerts thus increasing the driver's acceptance of these assistance systems [9]. The NHTSA-Project SAVE-IT therefore has proposed this development of intent inference algorithms to be a combination of theory-driven and data-driven approaches using methods from the machine learning sector and implying the use of driver monitoring information [9][10]. The EU-project AWAKE [11] also has focused on driver monitoring as a key information source to increase traffic safety by reducing the number and the consequences of traffic accidents caused by driver hypovigilance. The aim has been to develop an unobtrusive, reliable system based on a central component called Hierarchical Manager, which manages the information flow between the subsystems to monitor the driver and the environment and to detect hypo vigilance in real-time.

Summarizing these research projects, driver monitoring and driver intention analysis each have already shown strong potential for the development of intelligent vehicle behaviour. Therefore we have decided to explore the aspects of driver monitoring as an integrated component of a driver intention analysis, which is embedded in our context aware framework.

III. CONNECTEDDRIVE: THE CONTEXT

The central idea of the BMW "ConnectedDrive" paradigm is the combination of contextual information of an automotive framework according to the three autonomous fields illustrated in figure 1. The vehicle environment, its current driving state parameters and driver information can be interpreted as the vertices of a so-called contextual triangle, representing an image of the global situation context.

Environment information is acquired by sensors analyzing the traffic in the front, the side and the back of the car, including technologies like radar or lidar. The vehicle context describes the state of the car, like its velocity and

Fig. 1. Environment-Driver-Vehicle: These vertices stretch the context triangle and define the situation.

acceleration but also system status information. The last vertex of the context triangle includes the role and the influence of the driver on the situation, implying detailed information on his point of gaze and interactions with the vehicle.

Today's ADAS mainly have a lack of context information based on missing sensor availability to cover all three aspects of the situation. Nevertheless a perfect cognitive car should not miss one of the knowledge vertices in order to offer the next step of higher level context aware applications. One missing aspect, e.g. driver information, limits the intelligence of future ADAS and IVIS developments.

So the central idea of our project has been to build a car, which is able to cover all parts of the context triangle and bring these information sources together to one exemplary intelligent vehicle behaviour .

IV. Hardware-Framework: The Experimental Vehicle

Our experimental cognitive car has been equipped with sensors to cover parts of each of the above named contextual information fields. All vehicle state information like acceleration, speed and other common vehicle dynamics can be obtained from the Controller Area Network (CAN) bus. Environmental parameters like distance and relative speed to other cars are provided by a radar sensor of the built-in Adaptive Cruise Control system. For further processing of the road situation a camera based lane detection system has been installed behind the windshield.

A special focus regarding context sensors has been laid on driver monitoring approaches. A wide range of driver monitoring camera systems differing in their image capturing process, their refresh rate, their resolution and their colour space are currently available on the market. This hardware scale can also be matched to software scale regarding the information parameter, which has to be calculated by the driver monitoring system. While a head detection algorithm today can be implemented on a low-cost component, a system for the detection of the driver's point of gaze has even harder requirements on the software and hardware architecture of the driver monitoring system. The mapping of this hardware and software scale to the corresponding applications information parameter requirements has become a central research goal of our ongoing research project. We therefore have installed one version of a low-cost webcam-camera system and also one high-performance eye-tracking system. While the latter provides us with detailed information about the driver's gaze, head pose parameters and eye opening measurements, the webcam image has only been used to get raw information about the head position and head pose. Thus, by offering different levels of detail information, the systems offer different limits to the design process of adaptive strategies. Both systems work in the near-infrared, fulfilling the automotive constraint for a driver monitoring system to work at day- and night-time.

V. Software-Framework: The Idea of a Context Manager

A cognitive vehicle has to be aware of all past and current context information, so it has the possibility to interpret the situation and predict the right behaviour for upcoming situations. Therefore the two main aspects of the software architecture of our prototype car have been memory and intelligent reasoning capabilities, both integrated in one central knowledge managing component, the Connected-Drive Context Server (CDCS). The basic framework layout is shown in figure 2.

All sensor data is collected by a central I/O layer and forwarded to the core component of the server-client based context framework, the CDCS. It is processes and stores incoming raw contextual information from sensory components in a object-orientated knowledge structure. Each parameter object has its own value and timestamp history and is also analyzed over time to assign a quality value for each of them. The second key feature of the CDCS is its both flexible and powerful toolkit and interface

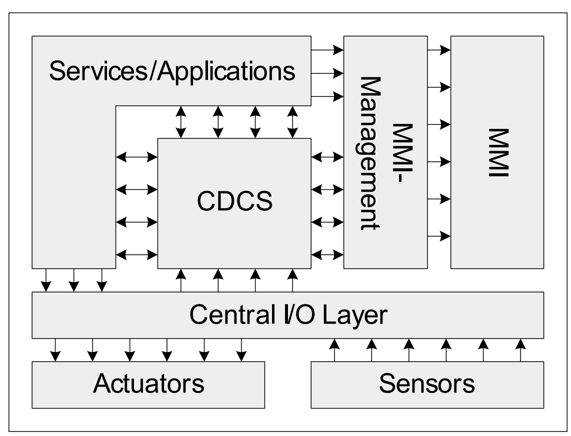

Fig. 2 Software Architecture of the Cognitive Vehicle

to generate high-level information parameters, so-called meta parameters, which can not be measured directly, but can be calculated through a intelligent combination of raw context values. For example the driver's intention could derivate out of his steering behaviour and his gaze parameters. For this purpose the CDCS is equipped with a flexible interface for the integration of statistical models for a knowledge generation process like Neural Networks, Support Vector Machines or other learning techniques known from the field of machine learning. Furthermore the CDCS not only offers the interface to integrate these models in a agent-based generation process of these meta parameters in real-time, but also provides the toolkit to record and prepare the data needed for the training of these models.

All applications can request their necessary context parameters from the shared context knowledge base of the CDCS and also insert parameters generated by their own behaviour.

Our framework also envisions one central Human-Machine-Interface (HMI) management instance, which, using the knowledge of the CDCS, arranges all applications' requests for MMI output modalities in a queue in order to guarantee an optimal context sensitive display of the applications outputs to the driver.

The following section deals with one example application, a context sensitive lane departure warning system, which has been completely developed using the tool chain of our context-aware framework.

VI. Context Sensitive Lane Departure Warning System (LDWS)

In our LDWS we have used data from the CDCS to adapt the application to the driver's intention. In case of departing the current lane, the system has to judge if this is intended or not. The criteria for departing the lane is fulfilled in case of the Time-To-Line-Crossing (TLC)-value [12] falling below a given threshold. The TLC-value indicates the preceding time until the lane marking will be crossed with the car's wheels under the assumption of a constant lateral velocity. Many of the intended lane departures consist of lane changes. In these cases, alarms of a LDWS could disturb the driver and are defined as nuisance alarms. To avoid this type of alarms, current systems on the market use the turn indicator for driver's lane change intention due to its mandatory use by law in these manoeuvres. On the other hand, drivers do not consequently use the turn indicator or sometimes too late [13] and there are further situations like for example slightly cutting corners, where there is no need to use the indicator and therefore nuisance alarms could occur, yielding a decrease in system acceptance [9].

Focussing on the avoidance of nuisance alarms, a huge potential could therefore result from additional indicators recognizing driver's intention. Regarding the whole system, there are further important aspects to be taken into account, like the transparency for the driver, acceptance and usefulness of the system, to mention only a few, but which are not addressed in this paper.

In our first research step we wanted to explore several strategies for the layout of a context-sensitive LDWS. Therefore we have accomplished a small user study, in which in total 10 subjects had to drive a test circuit, mainly on german highways with at least two lanes for each direction. They have not been told about the goal of the research to receive a naturalistic driving behaviour and a situation database for intended lane departures. All signals from the CDCS have been recorded and used for the design and off-line analysis of drivers´ intentions recognition strategies.

We have first decided to focus on realising a warning strategy, based on a pre-trained statistic classifier using only data from the car's common sensors like steering wheel angle, heading angle to the lane and related measures. Main advantage is the cost-efficient realisation of that strategy, as there would be no need for additional sensors in comparison to the series-production vehicle. A-priori a disadvantage consists of a difficult separation between in particular very smoothly and slowly accomplished intended lane changes to the right and unintended lane departures. Another difficulty yields in the fact, that it is hard to model unintended lane departures due to the lack of data availability. In comparison to the research work of Schmitz [14], we have voted against a provoked generation of these unintended lane departure scenarios in our data recording sessions with test persons, as we believe these experiments are too dangerous. Instead we have used data from our researchers simulating unintended lane departures (with full vehicle and environmental security control) to get relevant data for this scenario. Every time the TLC value hits the threshold, we have calculated a set of 112 features (raw parameters and some of their statistical derivates) and have labelled it according to the situation (intended ⬄ unintended, left ⬄

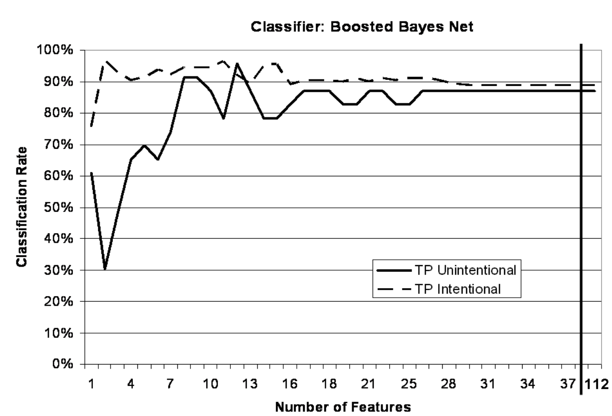

Fig. 3: Classification rates for the task "lane departure right"

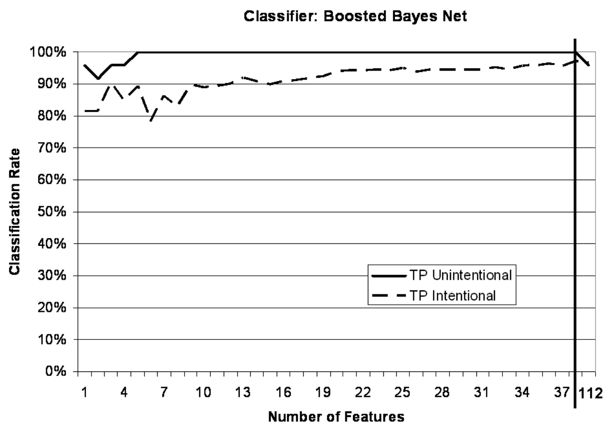

Fig. 4: Classification rates for the task “lane departure left”

right). Our training dataset has consisted of 70 instances for each left and right intended lane departures (only lane changes) and 12 instances for each left and right unintended lane departures. This data has been used to train multiple statistical classifiers (Bayes Net, Neural Net, Decision Tree and Support Vector Machine) and also enhanced classification schemes like AdaBoost. We have tested the models on our test corpus including lane change scenarios extracted from test persons driving on our test circuit. As there have not been any instances of unintended lane departures during these recording sessions, we again have chosen to record simulated data for the analysis of our classification models. The test corpus consists of 318 lane changes to the right, 12 simulated unintended lane departures to the right , 271 lane changes to the left and 12 simulated unintended lane departures to the left. Figure 3 and 4 show the results of the best classification model, a classifier based on boosted Bayes Nets, for both separated classification tasks: lane departure right and lane departure left. The dashed line shows the classification performance for unintended lane departures with a variation of the used

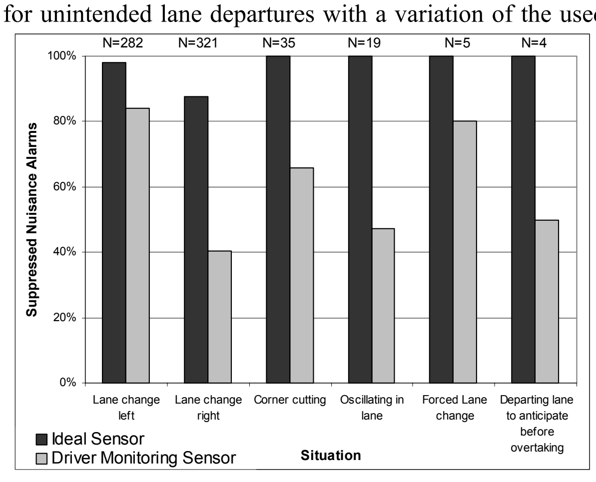

Fig. 5 Suppression of nuisance alarms: comparison between theoretical and practical results with strategy “gaze analysis”

number of features drawn on the x-axis. The same interrelation for intended lane departure scenarios has been drawn with a solid line. The order of the added features relates to their specific information gain, which has been calculated on the training corpus. Both figures show clearly, that the trained classifier can reach overall classification rates of over 90 percent. Although this encouraging first results, online testing of the algorithms in our test vehicle showed that transparency for the driver could nevertheless be poor, because it might be hard to understand the strategy of the statistical data mining model due to its black box character.

While modelling lane change behaviour with vehicle parameters has shown promising results, this method is not feasible for other intended lane departure scenarios, like curve cutting or lane departures for anticipation reasons before an overtaking scenario. These manoeuvres show to much diversity in their execution and occurrence. Therefore we have searched for alternative ways to infer this kind of intended lane departures. So In a next step we have decided to analyze the potential of our camera-based driver monitoring system for the intention analysis. We have developed a strategy called “gaze analysis”. This strategy in a first step uses the detailed gaze information of the high performance camera system and decides to warn the driver, if he departs the lane and his point of gaze is directed somewhere inside of the car. An exception is gazing at the speedometer, because many experimental drivers show that behaviour briefly when changing the lane. Therefore, a warning is only given, if the gaze at the speedometer exceeds an time threshold. Also in case of missing values from the driver monitoring sensor e.g. due to difficult light conditions, warnings are not suppressed, except the driver showed signs of intended lane departures, like looking into one of the mirrors or watching the blind spot in a given time span. Complementary, in case of looking outside of the car, all warnings are suppressed. A big advantage can be seen in a supposed low rate of wrongly suppressed warning, with only having little theoretical drawbacks of the strategy, e.g. the driver is looking outside of the car but is cognitive distracted or tired while having his eyes opened. Of course the ratio of correctly suppressed nuisance alarms should be high for a broad field of different driving situations. Figure 5 shows the comparison between the theoretical suppression power (assuming ideal driver monitoring values) and the actual values of this strategy from our evaluation rides. During the test drives the performance of the gaze analysis strategy for all scenarios was mostly limited by the overall availability of driver monitoring based parameters. The big difference between left and right lane changes emerged from a diverse tracking capability of the driver monitoring system. Compared with the left lane changes, our test drivers often moved their heads out of the tracking range of the used driver monitoring system when securing their lane changes to the right. This resulted in a higher tracking loss ratio.

However there's still a rather large gap between the theoretically achieved performance with an ideal driver monitoring sensor with no missing data and a real sensor with a higher drop rate.

VII. Future Work

Regarding the context sensitive lane departure warning system we plan to achieve several aspects in our future work.

The analysis of vehicle dynamics has to be optimized regarding an optimal feature setting, and classification scheme. We also think of designing a more reliable benchmark to produce a maximum suppression power of nuisance alarms, while ensuring that no unintended departure gets falsely rejected.

Our evaluations show clearly that thus indicating a high potential for driver intention recognition, driver monitoring sensors have to be optimized yielding less drop outs. We further plan to evaluate driver monitoring based strategies with less technical demand, e.g. based on raw head pose estimation instead of explicit gaze information as a driver distraction indicator.

Although both strategies, vehicle information and driver information based, have shown their potential independently during our evaluations, it is our opinion that one strategy alone will hardly suffice to reach the goal of complete driver intention detection in a lane departure warning scenario. Therefore we plan to further analyze a hybrid combination of statistical vehicle based data and rule based driver monitoring approaches in our next experiments.

Furthermore we plan to transfer our collected knowledge from the design and evaluation of the context sensitive LDWS to other possible driver assistance systems like lateral collision warning systems, which could also benefit from the introduced context architecture and tool chain.

VIII. Conclusion

In this paper we have introduced the project SURF, one of BMW's actual research works on the topic of cognitive vehicle behaviour. We have illustrated our experimental vehicle, which has been equipped with a range of sensors acquiring several information sources for the symbiotic analysis of the vehicle, the environment and the driver. We have shown our vision of a context analysis framework and have given a brief overview of our exemplary application, a context aware LDWS, which has the capability to offer cognitive behaviour based on intelligent situation information processing.

Our first results have indicated that it is possible to detect the intention to leave the lane only by monitoring vehicle status parameters. Nevertheless we think using this information in an intention analysis model of a LDWS can only be reasonable in a symbiotic analysis with other information sources like driver monitoring parameters, as it would lead to a more robust decision process for our problem.

The theoretical results of our evaluation with the driver monitoring information of our high-tech camera system have shown that there is a lot of potential in this new emerging technology. However we have also learned from our experiments that the technology for a precise gaze tracking is far away from being reliable to use in driver warning systems like a LDWS, as it's far too expensive and there has not emerged any usable technical approach for an automotive series development environment.

Nevertheless the experiments have enabled us to experience the possibility in an intelligent system design, which would open up having those detailed driver monitoring information. Therefore these first steps have encouraged us to further explore the chances of this technology track in our research projects. Our flexible hardware and generic software framework offers us large possibilities in a fast prototypical development, implementation and evaluation process on our way to future adaptive and context sensitive vehicle applications.

References

[1] S.N. de Ridder, J. Hogema, M. Hoedemaeker, "The Dutch experience with Lane Departure Warning Assistant systems: A Field Operational Test", TNO document TM-03-C048, 2003

[2] T. Bachmann, K. Naab, G. Reichart and M. Schraut," Enhancing Traffic Safety With BMW's Driver Assistance Approach Connected Drive", Torino, Intelligent Transportation Systems Conference, 2000

[3] D. Wisselmann, K. Gresser, H. Spannheimer and K. Bengler, „ConnectedDrive- ein methodischer Ansatz für die Entwicklung zukünftiger Fahrerassistenzsysteme", Garching, Tagung Aktive Sicherheit durch Fahrerassistenz , 2004

[4] R.Schindhelm, Christhard Gekau & Marika Hoedemaker , „COMUNICAR Information Manager: Ergebnisse der Felduntersuchungen", Garching, Tagung Aktive Sicherheit durch Fahrerassistenz , 2004

[5] H. Kussmann, H. Modler et. al, "System Architecture, data flow protocol definition and design and AIDE specifications", AIDE Deliverable D3.2.2, 2004

[6] J. Lee, M. Reyes et al., "Safety Vehicles using adaptive Interface Technology (Task 5) Final Report: Phase 1", 2004

[7] G. Geiser and G. Nirschl, "Towards a system architecture of driver's warning assistant". In Parkes, A.M. & Franzen, S. (Eds.) Driving future vehicles, London: Taylor & Francis (pp. 251-263), 1993

[8] NHTSA Benefits Working Group, "Preliminary Assessment of Crash Avoidance Systems Benefits", National Highway Traffic Safety Administration, Washington, DC, http://www. itsdocs.fhwa.dot.gov/ , 1996

[9] M. Smith and H. Zhang., "Safety Vehicles using adaptive Interface Technology (Task 8) A Literature Review of Intent Inference", 2004

[10] M. Smith and H. Zhang., "Safety Vehicles using adaptive Interface Technology (Task 9) A Literature Review of Safety Warning Countermeasures", 2004

[11] Awake Project Presentation, Deliverable 10.1, 2001

[12] D.A. Pomerleau, T. Jochem, et al., "Run-Off-Road Collision Avoidance Using IVHS Countermeasures", NHTSA Final Report, DOT HS 809 170, 1999

[13] S.E. Lee, E. Olsen and W. Wierwille, "A Comprehensive Examination of Naturalistic Lane-Changes", NHTSA, DOT HS 809 702, 2004

[14] C. Schmitz, "Adaptiver Spurverlassenswarner mit fahrerabsichts- und fahrerzustandsabhängiger Warnstrategie", Shaker Verlag, 2004