“Assistance-on-demand”: Development of a speech-based, personalized left-turning assistant

Martin Heckmann, Heiko Wersing, Dennis Orth, Dorothea Kolossa

Abstract

We recently developed the concept of “Assistance on Demand” [1]. This describes an advanced driver assistance system (ADAS) which supports a driver in an inner city scenario only if she asks for assistance. A key element is the control of the ADAS via speech which allows the driver to flexibly formulate her requests for assistance while the situation develops. Our application scenario is turning left at unsignalized urban intersections. After the driver has activated the system via a speech command it monitors the right side traffic and informs about suitable gaps to enter the intersection, just like a co-driver would do.

In a first user study together with the Würzburg Institute for Traffic Sciences we investigated this concept in a driving simulator [2]. The results showed that drivers clearly preferred our speech-based system to a visual system implemented via a HUD and to driving manually without system support.

We assume that drivers differ in what they perceive as a suitable gap to make the left turn. To test this hypothesis we have performed a second simulator study using CarMaker where 9 participants were turning left in crossing traffic from both sides. We deploy a probabilistic method to estimate the smallest accepted gap of each driver, so called critical gap. The results reveal that there is, as postulated, a significant inter-individual difference in the critical gap between the drivers. Next we investigate how well we can predict if a driver will accept a gap presented to him. We show that a prediction based on a driver’s personalized critical gap can achieve an accuracy of more than 90%.

Motivation

The increasing networking and digitization of vehicles offers a variety of new possibilities for driver assistance functions. Complex infotainment systems are also gaining ground in modern vehicle cockpits. Aside from the advantages of these functions, there is also an increased risk of drivers being distracted by the systems and their use. Using the systems may lead to a cognitive overload on the part of the drivers, keeping them from optimally performing the driving task. Manual operation and visual displays in particular compete against vehicle guidance and lane monitoring.

In order to counteract the distractions and cognitive overload, new control concepts for the human machine interface (HMI) are required. Future HMIs need to not only cover an increasing number of functions but also focus on intuitive and efficient use. Current HMIs are mostly based on a combination of haptic input and visual displays. The visual channel, on which driving already places a heavy strain, thus suffers from additional stress. The auditory channel, in contrast, is relatively open. The interaction with the in-

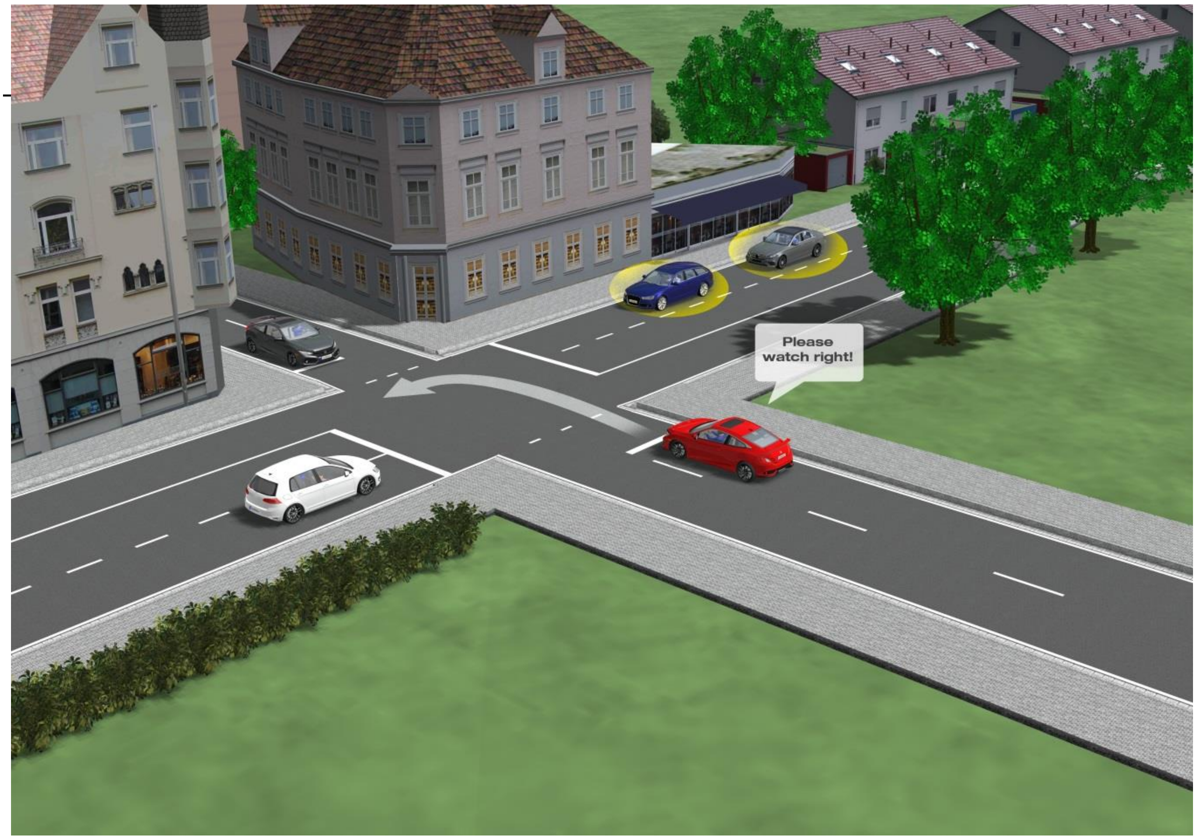

Figure 1: The speech-based on-demand intersection assistant (© IPG Automotive GmbH)

tersection assistant presented here is therefore speech-based. Another advantage of the speech-based interaction is its intuitiveness, as it is similar to the interaction with a human passenger. In addition, drivers do not need to take their hands off the wheel. Moreover, drivers are supplied with personalized recommendations that are based on their previous driving behaviour. The underlying assumption is an increased effectiveness and acceptance of the system due to the adaptation to drivers' preferences.

Description of the on-demand intersection assistant

The intersection assistant is developed by the Honda Research Institute Europe in cooperation with Ruhr University Bochum. The system offers support for the driver when turning left in inner-city traffic by monitoring the approaching traffic on the right and suggesting suitable gaps (see Figure 1). It is assumed that the suitability of a gap depends not only on the traffic situation but also on the individual drivers, in particular their driving behaviours.

Drivers activate the intersection assistant themselves when approaching an intersection, e.g. via the command "Please watch right!" In contrast to conventional advanced driver assistance systems that are constantly active, this assistance system is only activated on demand.

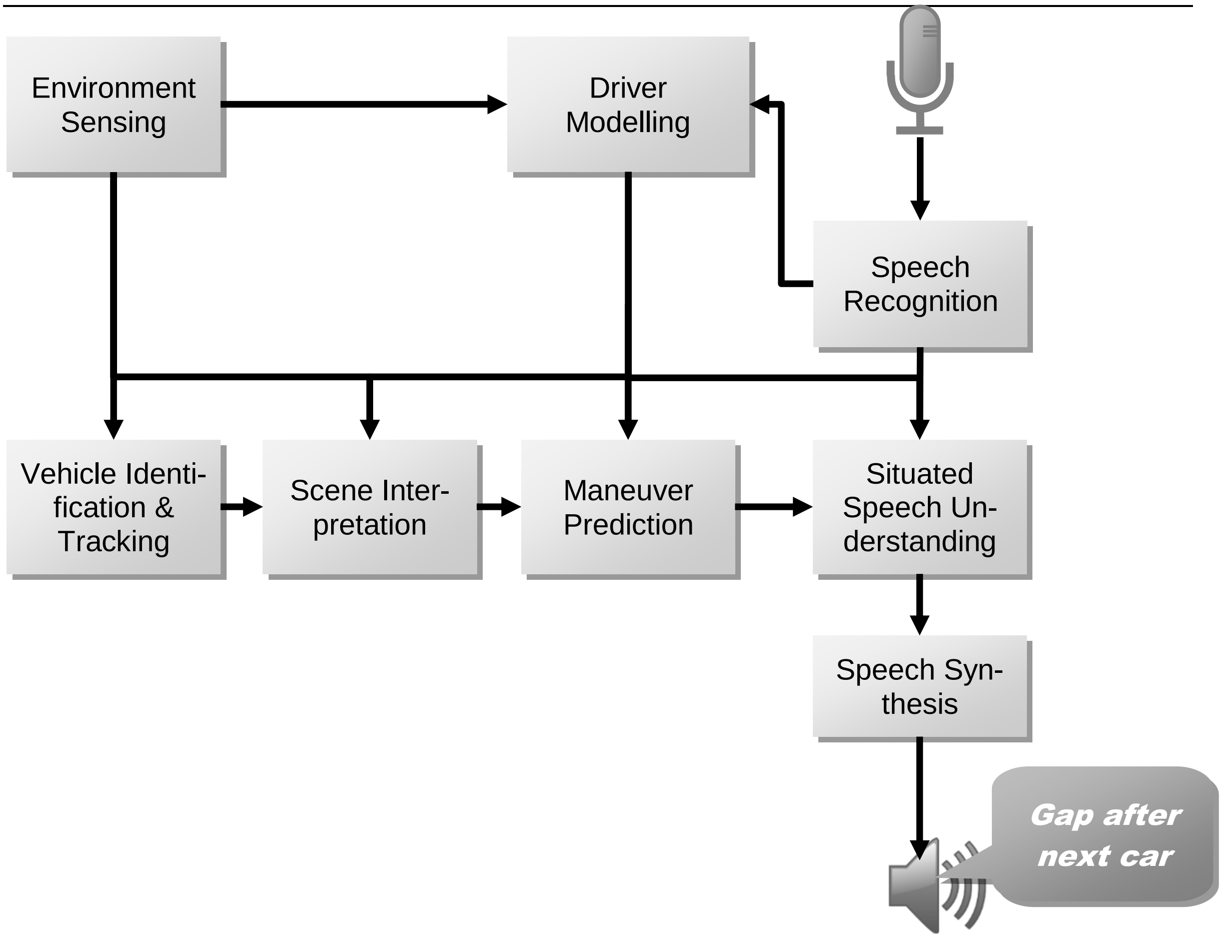

Figure 2: System graph

Systems that are permanently active have the drawback of potential false alarms and acoustic warning signals without apparent reason. False alarms are a distraction for drivers and may lead to drivers turning the system off. Drivers rejecting system warnings as unnecessary and ignoring them would also have a contra-productive effect on safety.

The system recognizes the environment and gives feedback on the traffic situation (e.g. “Gap after next car”). However, it does not give direct instructions with regard to the drivers’ actions. Drivers are not supposed to rely exclusively on the system but decide themselves whether the suggested gap is suitable.

Figure 2 illustrates the individual components of the intersection assistant. The system draws on a memory of previous perceptions of the environment and decisions made by the driver. Based on this data and the current interpretation of the scenario, the system informs the driver about the possibility of merging with the cross traffic.

Determining the individual gap acceptance

In a user study in cooperation with the Würzburg Institute for Traffic Sciences, the general acceptance of an on-demand system of this type was investigated in the Institute’s driving simulator. The study shows that test participants strongly preferred the assistance system to driving without support and that the system can facilitate the driving task when turning left [2]. This was especially evident in heavy traffic. When asked, however, many test participants noted that the suggested gaps did not correspond to their driving behaviour in many cases.

But which gaps are accepted as suitable by the individual driver? In order to answer this question, studies were carried out on a static driving simulator at the Honda Research Institute Europe. The simulator is embedded into the whole vehicle simulation of the open integration and test platform CarMaker with a steering wheel and pedals to steer a virtual vehicle. This simulation environment enables the detailed transfer of real test scenarios, including the entire environment, into the virtual world. Virtual test driving requires detailed, real-time capable models of the vehicle, road, driver and traffic. When using the simulation environment in combination with driving simulators, the virtual vehicle is driven by the test participant instead of the driver model. The traffic situation is visualized by means of the 3D visualization tool in CarMaker and displayed on screens that are placed around the driver.

This simulation platform allows for the generation of reproducible test scenarios with other road users and static objects such as parked cars, buildings or construction sites. An inner-city intersection scenario was generated in order to investigate which gaps are accepted by the individual driver. The scenario consisted of four successive three-way intersections. Buildings were placed on the side of the road obstructing the view of the intersecting road, which forced the driver to stop at the intersections. The layout of the four intersections was identical.

The test participants first acquainted themselves with the driving simulator in a training session. In short intervals, all nine participants then completed three scenarios with 16 intersections each, i.e. each participant turned left a total of 48 times and spent approximately one hour in total in the simulator. The traffic situation was identical for all drivers. During the drives in the simulator, each participant was offered the same gaps, which were between two and eight seconds. All cars were travelling at 50 km/h, i.e. the inner-city speed limit.

After the test drives, the so-called “critical gap” was determined for each driver. The critical gap (t_c) describes those gaps whose size is just large enough for the driver to accept them. Naturally, this parameter cannot be observed directly. Instead, the largest rejected gap (r_i) and the gap actually taken (a_i) were recorded for each intersection. It is assumed that the individual critical gap for each driver lies between these two values. Eq. 1 describes these facts by means of a probabilistic model. The model assumes a log-normal distribution of the critical gap. According to Troutbeck [3], the largest rejected gap (r_i) and the gap actually taken (a_i) are first log-transformed to simplify the calculation. This allows for the critical gap to be described by means of the cumulative distribution function F of a normal distribution with mean μ and variance σ^2. The joint probability of all drives through the intersections of one driver is the product of the single probabilities:

$$L(\mu, \sigma) = \prod_{i=1}^{N} [F(\ln a_i \,|\mu, \sigma) - F(\ln r_i \,|\mu, \sigma)] \qquad \text{Eq. 1}$$

In Eq. 2, the estimations $\hat{\mu}$ and $\hat{\sigma}$ are subsequently calculated by maximizing the function L:

$$(\hat{\mu}, \hat{\sigma}) = \arg\max_{\mu,\sigma} L(\mu, \sigma) \qquad \text{Eq. 2}$$

The values thus calculated are valid in the logarithmic range. Finally, the estimations are converted from the logarithmic to the linear time range (Eq. 3). From this, the driver-specific critical gap t_c can be obtained.

$$t_c = e^{\hat{\mu}+0.5\hat{\sigma}}, s^2 = t_c^2\left(e^{\hat{\sigma}^2} - 1\right) \qquad \text{Eq. 3}$$

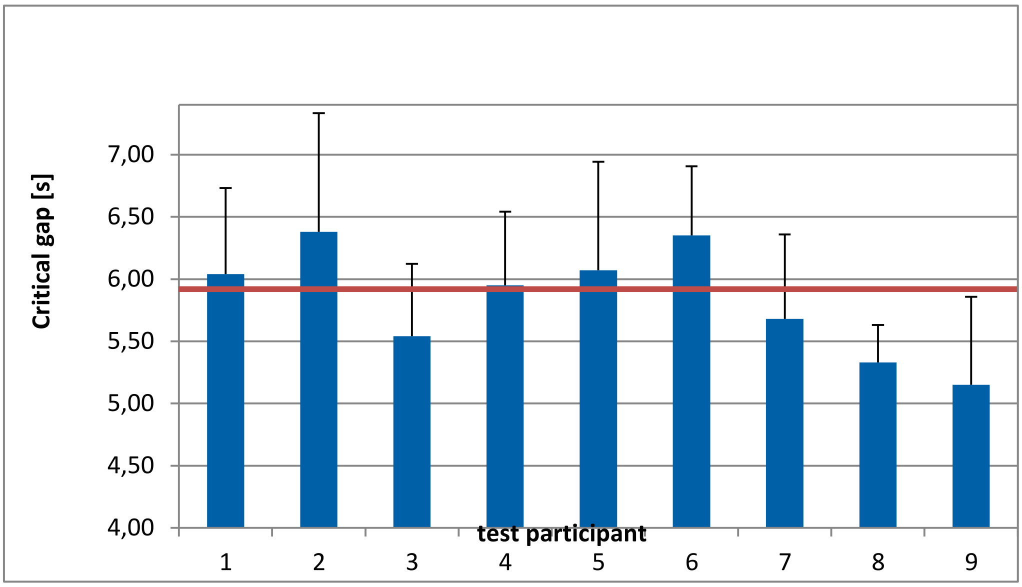

Figure 3: Gaps taken by drivers

Figure 3 illustrates the estimated critical gaps of each test participant and the corresponding standard deviations. The red line indicates the estimated critical gap ($t_c = 6.1\ s$) for the measured values of all participants. Similar values of the critical gaps on a priority road were published previously in traffic research (e.g. in [4]). It can therefore be assumed that the results are plausible and that the behaviour of the test participants in the driving simulator is sufficiently realistic. Furthermore, a statistical analysis was carried out focusing on the extent to which the values differ between the individual drivers. This analysis of variance (ANOVA) showed a high probability of distinct critical gaps for different drivers (ANOVA $F = 17.6$, $p \ll 0.01$).

Prediction of gap accepted by driver

The previous results indicated that there is a significant variation between the gaps the drivers accept. The task of the intersection assistant is to predict if the driver will accept the gap he is currently observing and support him accordingly. Hence in a next step we investigated how well we can predict if the driver will accept a gap. In particular we want to determine if a prediction based on a driver’s individual critical gap yields better results than one based on an identical critical gap for all drivers. We calculate the accuracy of the prediction by determining for each gap we observe if it is a least as large as the critical gap. If this is the case we predict that the driver will accept it. Otherwise we predict that he will reject it. First we use as critical gap the one we determined when we pooled the data of all drivers while leaving out the data from the driver we are currently investigating. We repeated this process for all drivers (leave one driver out). In a second experiment we determined the critical gap for an individual driver only from the data recorded from this driver. Here we used all but one intersection to determine this personalized critical gap and then evaluate it on the remaining intersection. Likewise, we repeated this process for all intersections (leave one intersection out) and all drivers and calculate the average.

Table 1: Prediction error rates for driver independent (leave one driver out) and driver dependent (leave one intersection out) prediction.

Error Rates		Relative improvement
Leave one driver out	**Leave one intersection out**	
12.3%	**9.8%**	**20.1%**

Already a first look on the results in Table 1 reveals that the prediction is in general quite accurate [5]. Even when we use a driver independent model we observe only

12.3% errors. When we replace this model with the driver dependent model obtained from the individual driver’s critical gap estimation the error rate drops to 9.8%.

Acceptance enhancement by personalization

The results of this study show that there are obvious differences in the gaps accepted by the drivers when completing identical intersection scenarios. The gap size varies between 5.1 and 6.4 seconds. Assuming a speed of 50 km/h, this corresponds to a difference of 18 m between the driver who shows the smallest critical gap and the driver for whom the largest critical gap was calculated.

We assume that some drivers prefer larger gaps because they exhibit a more defensive driving style in general, while drivers that tend to be more dynamic also accept smaller gaps. A prerequisite for the success of the personalization is that it allows the system to work more effectively. It will not suggest smaller gaps to a defensive driver who considers these too small, possibly resulting in frustration. Dynamic drivers, in contrast, will not receive suggestions for very large gaps as this could result in doubts about the usefulness of the system. The analysis of the prediction accuracy of the gap acceptance showed that using a driver’s personalized critical gap for this prediction reduces the errors by 20% and hence clearly increases the efficiency of the system.

In a further step, the investigation is intended to focus on the extent to which the personalization of the suggested gap actually increases the acceptance of the on-demand intersection assistant and whether this contributes to an improved interaction of the driver with the system. We are currently performing a second user study with the Würzburg Institute for Traffic Sciences to answer this question. In this study we will first estimate the personalized critical gap of the participants. Next they will experience a system variant using a driver independent critical gap and one using their personal critical gap. The results will show if the personalization is, as postulated, able to further increase the user acceptance. If the concept of a personalized, speech-based assistance system continues to prove successful, the system is planned to be subsequently tested on-road in real traffic.

Summary

Advanced driver assistance systems have great potential to significantly improve the safety and comfort of driving. Yet this can only be achieved if the systems can be operated by the driver without distracting him and furthermore drivers accept the support offered to them. Personalization of the systems can benefit both factors. Until now, the personalization approach was particularly relevant in the field of infotainment and navigation. The ideas presented here take the concept one step further and apply it to the area of advanced driver assistance systems. The intersection assistant can be requested on demand, intuitively and speech-based. It offers suitable recommendations for drivers to merge with crossing traffic. In a first driving simulator study it was shown that this way of supporting the driver was very well received. A following driving simulator study revealed that there are significant inter-individual differences in terms of the gap sizes which are considered just sufficiently large. The simulation platform allowed for an easy reproduction of the corresponding test scenario. A more detailed analysis of the results also showed that the efficiency of the system can be further increased when the recommendations to the driver are based on personalized driver models. In the continued development of the intersection assistant, there will be additional user studies. This will ensure that this novel, speech-based HMI concept can be geared to the drivers’ needs. The research results thus yield new approaches to the intelligent integration and usability of assistance functions.

Acknowledgements

We are much obliged to Mr Kersten Schaller and Prof. Andreas Pech for their contribution to the development and implementation of the intersection scenario and a first draft of the estimation of the critical gap. We also thank Mr Milton Orlando Sarria Paja for revising the estimation of the critical cap and for the statistical analysis. Furthermore, we would like to thank Ms Ramona Bach and Mr Christian Walter of IPG Automotive for their assistance in realizing the scenario in CarMaker and Ms Birgit Kremer for her help with previous versions of the manuscript.

References

[1] M. Heckmann, D. Orth, H. Wersing, D. Kolossa: Development of a Personalised Intersection Assistant, to appear in ATZ worldwide, May 2017

[2] N. Schömig, M. Heckmann, H. Wersing, C. Maag, und A. Neukum, „Assistance-on-demand: A speech-based assistance system for urban intersections“, adjunct proceedings of Automotive’UI 16, Ann Arbor, MI, USA, 2016.

[2] R. J. Troutbeck (1992): Estimating the critical acceptance gap from traffic movements. Physical Infrastructure Centre, Queensland University of Technology. Research Report 92-5

[3] A. Weinert, „Grenz- und Folgezeitlücken an Knotenpunkten ohne Lichtsignalanlagen“, Ruhr-Universität Bochum, 2001.

[4] D. Orth, D. Kolossa, M. Sarria-Paja, K. Schaller, A. Pech, M. Heckmann: A Maximum Likelihood Method for Driver-Specific Critical-Gap Estimation, submitted to IEEE Intelligent Vehicles Symposium, 2017