FAKULTÄT FÜR INFORMATIK

DER TECHNISCHEN UNIVERSITÄT MÜNCHEN

Increasing the maturity of the Augmented Reality Head-Up-Display

Christian A. Wiesner

FAKULTÄT FÜR INFORMATIK

DER TECHNISCHEN UNIVERSITÄT MÜNCHEN

FACHGEBIET AUGMENTED REALITY

Increasing the maturity of the Augmented Reality Head-Up-Display

Christian A. Wiesner

Vollständiger Abdruck der von der Fakultät für Informatik der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitzender: Prof. Dr. Michael Bader

Prüfende der Dissertation:

1. Prof. Gudrun J. Klinker, Ph.D.
2. Prof. Dr. Albrecht Schmidt, Ludwig-Maximilians-Universität München

Die Dissertation wurde am 29.07.2019 bei der Technischen Universität München eingereicht und durch die Fakultät für Informatik am 09.11.2019 angenommen.

Abstract

With the introduction of Head-Up-Displays (HUDs) into the vehicle the possibility to create Augmented Reality (AR) applications emerged. The development of the AR HUD remains a challenging field, as knowledge from several research areas, such as computer science, engineering and psychology is required to create a sophisticated product.

This work assesses the accuracy of current Global Navigation Satellite System (GNSS) sensors in series vehicles. The determined positional error of the GNSS sensors in series vehicles is still too high in order to steadily position AR visualisations on the road. Therefore we developed a visualisation concept taking this error into account. This "Sails" visualisation is solely dependent on the distance to the next turn and changes its appearance accordingly, creating a dynamic visualisation. The visualisation was well perceived by participants in a comparative user study.

Furthermore this work presents the complete development of a visualisation of the future road course, displayed in a three-dimensional manner (3D-FRC). The 3D-FRC is based on the electronic horizon which is already available in current premium series vehicles. This visualisation overcomes the limited field of view (FoV) of the HUD by presenting the information in an *AR-like* fashion. It has been successfully implemented in a vehicle prototype and the conducted user study found an overall improvement of the participants' braking behaviour around sharp corners.

Another area investigated in this work is the assessment of the change in gaze behaviour by the introduction of HUD visualisations. The results show that *AR-like* visualisations significantly decrease the visual attention towards the vehicle's instruments, therefore potentially decreasing the time the driver spends not looking at the road. However, the visual attention towards the HUD area increased. Therefore a further investigation was conducted into the gaze behaviour introduced by conventional HUD representations. The results show that conventional HUD visualisations already increase the attention towards the HUD area significantly and by adding *AR-like* visualisations the attention is not further increased.

Lastly, the investigation of the temporal aspect of HUD visualisations was conducted. The end-to-end latency measurement revealed that the latency is still too high to display seamless full-AR visualisations. Furthermore the perception threshold of AR HUD applications was determined through user studies. This revealed comparable results in the case of following an object to touchscreen applications, but found different results in the case of appearing objects.

This work presents the assessment of sensors (GNSS), the development of complete visualisation concepts (Sails and 3D-FRC) and their thorough investigation through user studies. Overall, this work demonstrates the strong potential of the AR HUD in increasing the trust in assistance systems, on the path towards automated driving.

Zusammenfassung

Mit der Einführung von Head-Up-Displays (HUD) im Fahrzeug wurde es möglich Augmented Reality (AR) Applikationen zu kreieren. Die Entwicklung des AR-HUD bleibt ein herausforderndes Feld, da verschiedene Disziplinen, wie Informatik, Ingenieurswessen und Psychologie notwendig sind um ein ansprechendes Produkt zu gestalten.

Diese Arbeit beurteilt die Genauigkeit von derzeitigen Global Navigation Satellite System (GNSS) Sensoren in Serienfahrzeugen. Der bestimmte Positionierungsfehler von GNSS Sensoren in Serienfahrzeugen ist immer noch zu hoch um AR-Visualisierungen stabil auf der Straße zu positionieren. Deshalb haben wir ein Visualisierungskonzept entwickelt, welches diesen Fehler berücksichtigt. Diese "Segel"-Visualisierung ist lediglich von der Distanz zur nächsten Abbiegung abhängig und verändert ihr Aussehen entsprechend, was eine dynamische Visualisierung erstellt. Die Visualisierung wurde gut angenommen von Teilnehmern einer vergleichenden Nutzerstudie.

Außerdem wird in dieser Arbeit die vollständige Entwicklung einer Visualisierung des zukünftigen Straßenverlaufs vorgestellt, welche in einer dreidimensionalen Art und Weise dargestellt wird (3D-FRC). Der 3D-FRC basiert auf dem elektronischen Horizont, welcher bereits in derzeitigen Premium-Serienfahrzeugen vorhanden ist. Die Visualisierung überwindet das begrenzte Field of View (FoV) des HUDs, indem die Information in einer *AR-like* Art und Weise dargestellt wird. Sie wurde erfolgreich in einem Fahrzeug-Prototypen implementiert und die durchgeführte Nutzerstudie stellte insgesamt eine Verbesserung des Bremsverhalten der Teilnehmer bei scharfen Kurven.

Ein anderer Bereich welcher in dieser Arbeit betrachtet wurde, ist die Beurteilung der Änderung des Blickverhaltens durch die Einführung von HUD-Visualisierungen. Die Ergebnisse zeigen, dass *AR-like* Visualisierungen die visuelle Aufmerksamkeit auf die Fahrzeug-Instrumente signifikant verringert und deshalb potentiell die Zeit verringern können, in welcher der/die Fahrer(in) nicht auf die Straße schaut. Allerdings ist die visuelle Aufmerksamkeit auf die Fläche des HUD erhöht. Deshalb wurde eine weitergehende Analyse des Blickverhaltens durch konventionelle HUD Repräsentationen durchgeführt. Die Ergebnisse zeigen, dass bereits durch konventionelle HUD-Visualisierungen die Aufmerksamkeit auf den Bereich des HUD signifikant erhöht wird und *AR-like* Visualisierungen diese nicht weiter erhöhen.

Abschließend wurde der zeitliche Aspekt von HUD-Visualisierungen betrachtet. Die Ende-zu-Ende-Latenz Messung zeigte, dass die Latenz immer noch zu hoch ist um nahtlose volle AR-Visualisierungen dazustellen. Außerdem wurde die Wahrnehmungsgrenze von AR-HUD Visualisierungen durch Nutzerstudien festgestellt. Diese zeigten vergleichbare Ergebnisse in dem Fall eines verfolgenden Objekts zu Touchscreen-Applikationen, aber stellten unterschiedliche Ergebnisse in dem Fall von erscheinenden Objekten fest.

Diese Arbeit präsentiert die Beurteilung von Sensoren (GNSS), die Entwicklung von kompletten Visualisierungskonzepten ("Segel" und 3D-FRC) und deren gründliche Untersuchung mit Nutzerstudien. Insgesamt zeigt diese Arbeit das enorme Potential des AR-HUD das Vertrauen in Assistenzsysteme zu erhöhen auf dem Weg zum automatisierten Fahren.

Acknowledgements

This work would not have been possible without the support by several people. First of all I want to thank Prof. Gudrun Klinker for her guidance in my research and for the fruitful discussions. I really enjoyed the friendly atmosphere in her lab and the times I spent there in person, sparked many interesting ideas. Furthermore, I am grateful for the professional setting, provided by the TU Munich, for conducting my PhD studies. I found the courses at the TUM Graduate School especially interesting. I would also like to thank Prof. Albrecht Schmidt for agreeing to serve as my second examiner.

I would like to express my gratitude to the Robert Bosch GmbH for providing a platform to conduct my thesis. Foremost, I want to thank my advisor Dr. Demet Sirim for supporting my research and for her continuous encouragement. Additionally, I would like to thank the development team for their help with the prototype vehicle.

Above all, I am indebted to my family for their infallible belief in me throughout my PhD. I wouldn't have made it through without their encouragement. Last but not least, I want to thank my wife Dr. Ratna Tripathy for her support in good times, as well as the more difficult ones.

Contents

Part I.

Introduction

This part introduces the concept of Augmented Reality (AR), its definition, applications and current limitations. A general overview of existing and future Head-Up-Displays (HUDs) is provided. The method of user studies is presented in the context of User Experience (UX). Benefits and limitations of user studies in a car prototype are compared to user studies in a car simulator. After a general overview of the existing related work the main contributions of this thesis are pointed out.

1. Motivation

The introduction of Head-Up-Displays (HUDs) help the driver in keeping the gaze towards the road while perceiving important information about the vehicle, e.g. the current speed. It furthermore enabled several new applications, because the information is displayed right in the area where the street is visible. Amongst these potential applications are applications involving Augmented Reality (AR). Examples for these applications involve showing navigational hints embedded into the real environment. Another potential is to display which other vehicles are recognised by the car. This can especially increase the trust in the advanced driving assistance systems (ADAS).

Developing AR visualisations for series production is still a challenging field. The field of view (FoV) of current HUDs still very limited as the available construction space behind the dashboard is quite narrow limiting the size of the used mirrors. Another challenge is to display the graphics accurately and steadily in the real environment. To achieve this accurate sensing of the environment is necessary.

Therefore, this work investigates the accuracy of the used sensors in vehicles, in particular the accuracy of current series Global Navigation Satellite System (GNSS) sensors. Upon the found restrictions by the sensors and the available display area Human Machine Interface (HMI) concepts are developed and the derived HUD visualisations are implemented in a car prototype.

After successful implementation of these visualisation concepts their perception by users is investigated in user studies. Besides the subjective perception in particular the change of the gaze behaviour is investigated by eye tracking systems.

Furthermore the temporal aspect was investigated by determining the found latency of the current AR HUD prototype system. Besides determining the current latency it is important to determine which level of latency is perceivable by humans. Therefore, this work also investigates the perceivable latency through user studies.

This work provides an overview of the current accuracy of the current sensors in the car. It then describes the development of visualisation concepts handling the technical restrictions of the system. This work takes the important step of porting AR applications, which have been thoroughly studied in simulators, into a real car prototype and assesses their effect on the drivers.

2. Augmented Reality

This chapter introduces the definition of Augmented Reality (AR) and surveys a few selected applications. Practical limitations of current AR applications are also pointed out in this chapter.

2.1. Definition of Augmented Reality

AR applications are characterised by embedding virtual objects into the real environment. Applications of AR have existed since the introduction of the Head-Mounted-Display (HMD) by Sutherland in 1968 [Sut68] and AR was only later defined by Ronald T. Azuma in 1997 [Azu97]. According to the definition, AR is characterised by the following:

1. Combines real and virtual objects

2. Is interactive in real-time

3. Is registered in three-dimensions

The applications of AR can be categorised into Video See-Through AR (VST-AR) and Optical See-Through AR (OST-AR) applications. The category of VST-AR displays the merged view on a monitor. Examples can range from ordinary desktop monitors over smartphone displays to video HMDs in which the merged view is displayed like the AR Rift by Steptoe et al [SJS14]. The AR Rift is basically an Oculus Rift [Rif] combined with consumer webcams. OST-AR applications are characterised by not obstructing the view onto the real environment. The concept of OST-AR can be seen in Figure 2.1. Examples for OST-AR are HMDs which enable the user to perceive both, the real environment as well as the virtual images. Another example is the technology of the Head-Up-Display (HUD). The HUD projects the virtual image onto the windshield, where the images can be perceived by the viewer without any additional equipment.

The reality-virtuality continuum which can be seen in Figure 2.2 was defined by Milgram et al. [MTUK94]. The ends of the reality-virtuality continuum represent the completely real and completely virtual environment respectively. The term Mixed Reality (MR) describes all the applications between these ends of the continuum and AR describes in the continuum the applications being closer to reality.

The definition of Azuma made it possible to classify the existing MR applications into pure AR applications and other MR applications. With the increase of computing capabilities by mobile processors and the introduction of sophisticated and cost-efficient displays, the number of AR applications has seen a corresponding expansion.

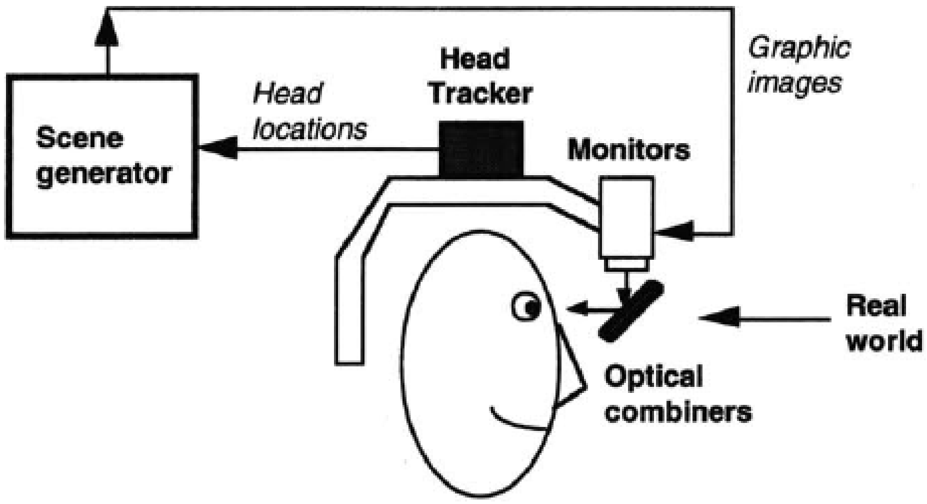

Figure 2.1.: Optical See-Through AR [Azu97]

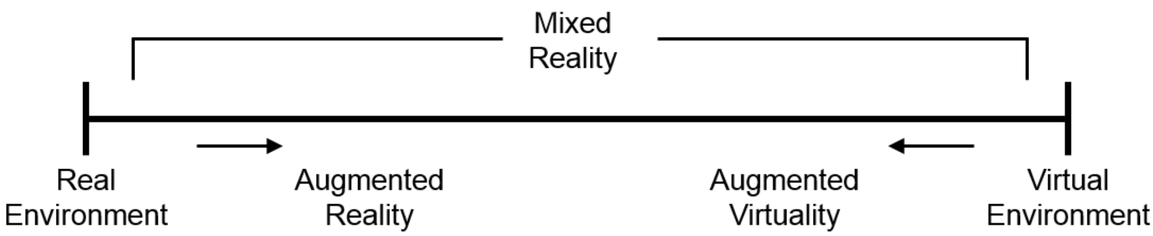

Figure 2.2.: The reality-virtuality continuum according to Milgram et al. [MTUK94]

2.2. Applications of Augmented Reality

The development of AR applications saw a rise in the 90s, among them the Knowledge-based Augmented Reality for Maintenance Assistance (KARMA) [FMS93] which embeds the visualisation of maintenance steps into the real environment. Another historic example is the Touring Machine by Feiner et al. [FMHY97] through which it was possible to wander around the university's campus and see information about the surrounding buildings, embedded in the viewer's field of view (FoV). Also after the 2000s the number of AR applications increased, especially through the introduction of smartphones. In the following we will highlight in particular the developments of automotive AR applications.

2.2.1. Automotive applications

Most of the research in the automotive sector of AR applications was conducted so far in car simulators. Tönnis and Klinker presented a visualisation (see Figure 2.3 a) for directing the attention of a driver towards an object not visible in the field of view, such as other vehicles or cyclists behind the car [TK06]. The visualisation outperformed a conceptual visualisation of the vehicle's bird eye view.

Kim et al. presented a visualisation (see Figure 2.3 a) in which a virtual two-dimensional map merges into the real environment [KD09]. The visualisation led to significantly fewer

(a) A three-dimensional arrow. [TK06]

(b) Virtual map merging into real environment. [KD09]

Figure 2.3.: Automotive AR applications in car simulators.

navigational errors and distraction-related measures compared to a standard navigation visualisation in the central display of the car.

In another application Charissis et al. presented a visualisation (see Figure 2.4 a) wherein other cars and lanes were highlighted to support the driver in low visibility situations [CPMA11]. The visualisation resulted in significantly lower reaction times and decreased the collision occurrences.

Bolton et al. compared different visualisations (see Figure 2.4 b) concerning the navigation task. They highlighted existing objects, such as houses, by the use of a virtual arrow or a frame box around the object [BBL15]. The virtual frame box had significantly lower reaction times and higher success rates.

All of the afore mentioned applications have been developed in car simulators. The results show a strong potential of AR applications in the vehicle, but it is close to impossible to be able to simulate all the possible factors, which can occur during real driving situations. One necessary step is to implement the AR visualisations in a car prototype to analyse the effect of the these visualisations on the drivers under real circumstances.

2.3. Practical Limitations of Augmented Reality

There are several limitations hindering the introduction of AR in several areas. They are most often restricted by price or by technology. One major hindrance is the limited computing power. However, in spite of the increase in the embedded and mobile computing devices, the manufacture of high performance cost effective solutions still remains a challenge. With an increase of computing power the additional weight for batteries as well as heat generation remain to be tackled. This becomes clearer by the examples of HMDs. Google Glass can be hardly considered as an AR application, but they managed to incorporate a small projecting unit on a relatively light-weight HMD. The next step was the development of true AR glasses like the Microsoft HoloLens [Hol] and the Meta Glasses [Met]. Whereas the HoloLens computes everything on the embedded hardware, Meta does the expensive computation on the connected PC. The newly introduced Magic Leap One [One] performs the calculations and stores the electrical power in a portable device.

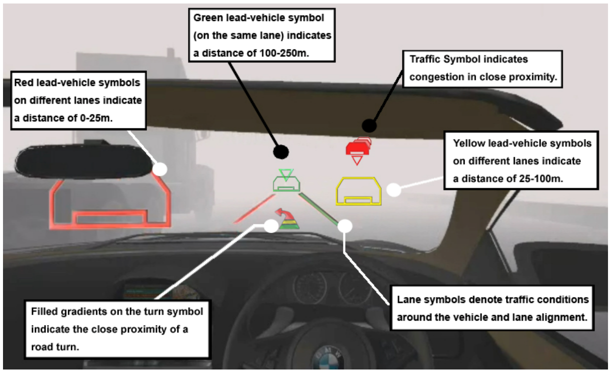

(a) HUD visualisation for low visibility conditions. [CPMA11]

(b) Highlighting of existing landmarks. [BBL15]

Figure 2.4.: Further automotive AR applications in car simulators.

When designing AR applications one has to take the display size into account. If the display is increased the resolution per viewing angle decreases. One major criticism of the HoloLens is its limited FoV, but the resolution per viewing angle enables it to display the virtual objects in accordance with the real environment.

Another important factor is the latency of the system. Here again, the HoloLens is renown for its stable and fast tracking of the virtual objects. To be able to position the virtual objects within the real environment the sensors to register the real environment need to have a minimum accuracy as well as providing the data fast enough.

2.3.1. Spatial limitations

In AR virtual objects are superimposed onto the real environment with a minimum required accuracy, such that the user has the impression that the virtual object has its position in the real environment. To achieve this the sensors for registering the real environment need to meet a certain accuracy. Examples for these sensors in the vehicle are the Global Navigation Satellite Systems (GNSS) sensors, the RADAR and the built-in cameras. Especially in the automotive sector the sensors have to endure extreme temperatures induced by the weather and shocks due to the condition of the roads. To increase the accuracy of positioning virtual objects in the environment the registration by the sensors can be fused together.

2.3.2. Temporal limitations

Even with sufficiently accurate sensors, another criterion to place virtual objects seemingly stable at a certain position, is to update the position of the virtual object fast enough happening due to changes of the viewpoint or if the real objects are being moved. This can either be achieved with a sufficient update rate of the sensors or the change of the position is predicted through techniques, such as the Kalman filter, when no new data is available from the sensors.

2.3.3. Display limitations

Even when the position of a virtual object is correctly determined and is updated sufficiently fast, it is still important that the object incorporates the lighting conditions of the real environment, such as shadows. Furthermore, it is necessary that the actual displayed colours match the intended colours in the real environment [IDAK15].

2.3.4. Summary

The aforementioned limitations are still the challenges we face today for creating convincing AR visualisations. This work investigates some of these challenges and points out concepts of dealing with a limited accuracy or a limited display area. Its focus is to create visualisation concepts which are already achievable with the current hardware and can be ported to series production in a relatively short timeframe.

3. Head-Up-Displays

This chapter introduces the technology of Head-Up-Displays (HUDs). The name is derived by the fact that users can keep their head up to perceive information. The user is not required to tilt the head downwards to read information from instruments. First, we describe the basic elements of an HUD and introduce common characteristics of an HUD. We continue with an overview of the different types of HUDs and close with an outlook to the AR HUD.

3.1. History of Head-Up-Displays

The first HUDs were developed in the avionic sector in the 1940s. The main reason for developing HUDs was to enable jet fighter pilots to perceive current information about the state of the aircraft and the weapon systems while still keeping their gaze towards possible targets through the windshield.

The first HUDs in the automotive sector were introduced by the end of the 1980s. The first European car manufacturer which introduced HUDs in series production was BMW with the BMW 5 series in 2003. Nowadays HUDs are available from almost all car manufacturers and across all classes of vehicles.

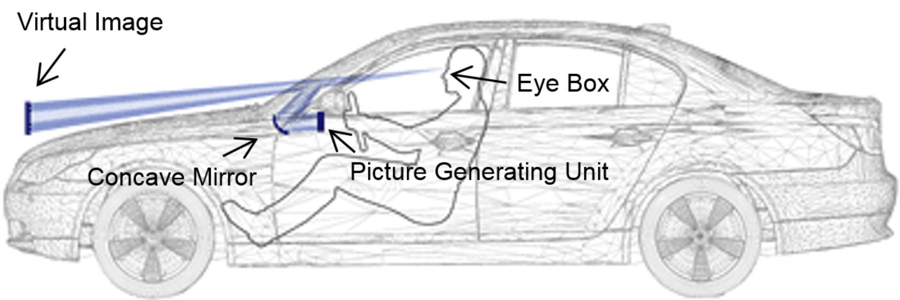

Figure 3.1.: Schematic of the construction of a HUD. The PGU generates the image and is then reflected via the mirror onto the windshield, where the image is perceived by the driver. (Courtesy of BMW)

3.2. Optical Elements of an HUD

An HUD is in general composed of the following parts: a Picture Generating Unit (PGU) which creates the displayed image, a projector unit (composed mainly of lenses and mirrors) and the combiner. The PGU generates the displayed image and the lenses and mirrors are mainly used to adjust the Field of View (FoV) and the Virtual Image Distance (VID). The combiner can be either the windshield of the vehicle or a combiner specifically designed for the HUD. An overview of the concept can be seen in Figure 3.1. We will describe the components in the following in more detail.

3.2.1. Picture Generating Unit

The Picture Generating Unit (PGU) generates the image being displayed on the HUD. Different technologies exist for these systems. The most common one is a Thin-Film Transistor (TFT) display, as they are cheap and reliable. The TFT is lit up by a backlight which is in current systems a Light-Emitting Diode (LED) as they employ a low power consumption. The PGU is required to display the virtual image with a sufficient brightness, in order to make the virtual image perceivable also under bright lighting conditions which occur on a sunny day. Furthermore, the backlight needs to be dimmable to not blind the user in low lighting conditions occurring while driving through a tunnel or by night. The technology of Digital Micromirror Devices (DMDs) is currently under investigation in the automotive sector, as their reaction time is much faster and they produce not as much heat. However, their integration into the package of the HUD is a lot more challenging.

3.2.2. Optical lens

The lens is used in the system to magnify the virtual image. Its basic principle can be seen in Figure 3.2. In the Figure f describes the distance to the focal point of the lens, $s1$ describes the distance of the object, the display on the PGU in our case and $s2$ describes the distance of the virtual image. The virtual image is then perceived by a second lens, the lens of the eye which creates again a real image on the retina. Through this magnifying lens it is possible to change the distance of the virtual image.

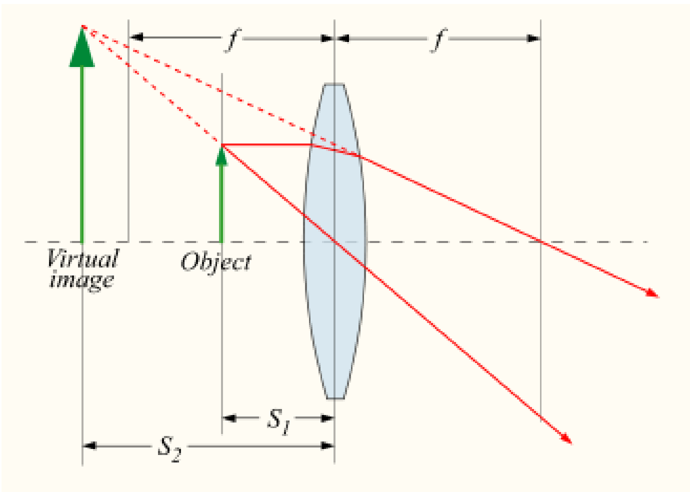

Figure 3.2.: Formation of a virtual image. [Com15]

3.2.3. Concave mirror

Through the use of a concave mirror it is possible to increase the size of the virtual image and also correct the shape of the windshield. The mirror needs to be adapted to the shape of the windshield, so the virtual image is perceived as one planar image. For more details see [Neu12] and [Jan18].

3.2.4. Optical reflection

The principle of an optical reflection describes that a virtual object is perceived at the distance of the length of the optical path, as can be seen in Figure 3.3.Through the use of additional mirrors in the HUD the distance of the virtual image can be increased further which is needed for AR visualisations.

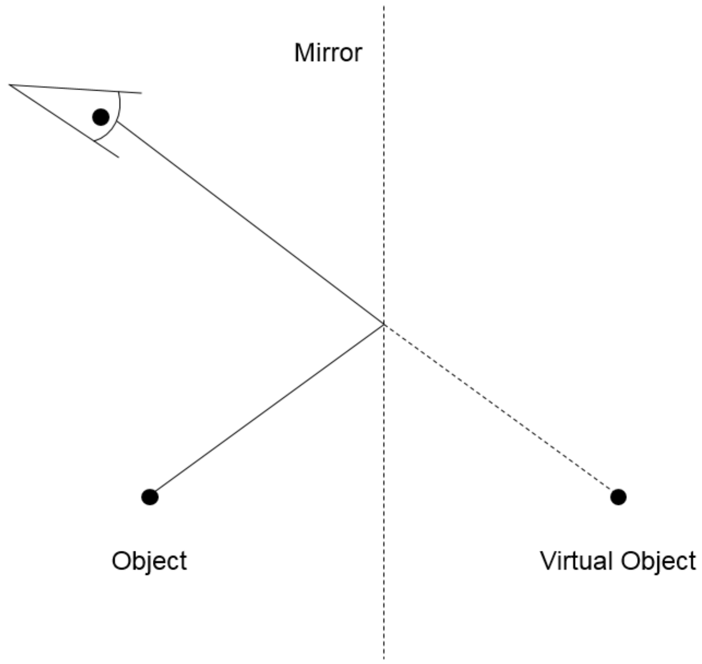

Figure 3.3.: Principle of an optical reflection.

3.3. Characteristics of an HUD

In the following we will describe the basic characteristics related to an HUD.

3.3.1. Field of View and Eye Box

The Field of View (FoV) describes the area in an HUD, in which it is possible to display visualisations. The size of the FoV is measured in degrees of the horizontal and vertical aperture angle. The Eye Box (EB) describes the area, from which it is possible to see the virtual image. In general the EB is designed in a way so that the width is approximately double the size of the mean interpupillary distance and the height is usually a few centimetres high so slight movements of the eyes are possible [Neu12].

The FoV and the EB are mutually dependent, as can be seen in Figure 3.4. The bigger the EB is designed, the smaller the usable FoV becomes and vice versa. Therefore, it is important to account for both factors while designing an HUD. Current conventional HUDs feature a FoV of usually 3°x 2°. For AR visualisations the visualisation itself and the objects one wants to augment require a much larger FoV of at least 8°x 4°. Through the larger FoV it is possible to augment the vehicle ahead in most cases. Naturally, the resolution needs to be high enough to display the virtual objects in a similar perceivable resolution as the real objects. The *fovea centralis* of the human eye has a density of 147,000 cone cells per square millimeter (mm) [Shr11], or 383 cones per mm. As the average distance between the fovea and the lens is 17.1 mm [SP11], the average perceivable angle of view is around 0,00875°as given by the following formula:

$$\alpha = 2arctan\frac{d}{2f} \tag{3.1}$$

where d denotes the size of the film, in our case the distance between two cones, and f denotes the focal distance, the distance between the fovea and the lens in our case. Therefore, one should make sure that one is displaying at least one pixel per 0,00875°. In our example this yields a resolution of at least 914 x 457 pixels.

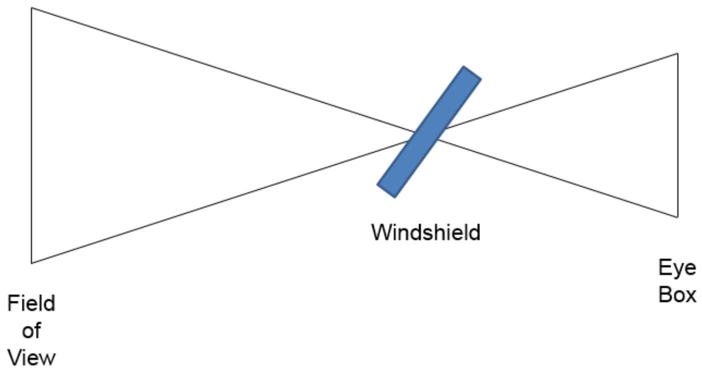

Figure 3.4.: Relation between the Field of View and the Eye Box.

3.3.2. Virtual Image Distance

The Virtual Image Distance (VID) describes the distance of the virtual image from the driver's point of view. It is determined by the optical path created by the mirrors and the magnification by both, the concave mirror and the optical lens. In conventional HUDs this distance is roughly around 2 metres. When displaying AR visualisations the objects that are to be augmented are usually a few hundred metres away. The VID affects the accommodation of the human eye: car drivers can either adjust their eye lenses to see sharply at 2 metres or at far distances. According to Cutting and Vishton [CV95] the *personal space* extends to 2 metres, whereas the so-called *vista space* begins at roughly 30 metres and the *action space* resides between the two. The different spaces and their interaction with depth cues can be seen in Figure 3.5. Thus, drivers cannot see both the real objects and their associated augmentations sharply at the same time. This has several affects: 1) Depth perception for real and virtual objects may differ. 2) This is tiring on the driver's eyes since the lenses need to shift back and forth between different distances. 3) Drivers are distracted from the real objects while looking at the virtual augmentations. Therefore, it is important to display the AR visualisations at a reasonable distance, so the integration of the visualisations is perceived as natural. To overcome these issues displays are developed, such as stereoscopic displays or holographic displays, to create the illusion of an image being perceived at an arbitrary distance. So far, these displays are far from entering series production and therefore the current HUDs display visualisations at a fixed VID which is higher than the afore mentioned 2 metres.

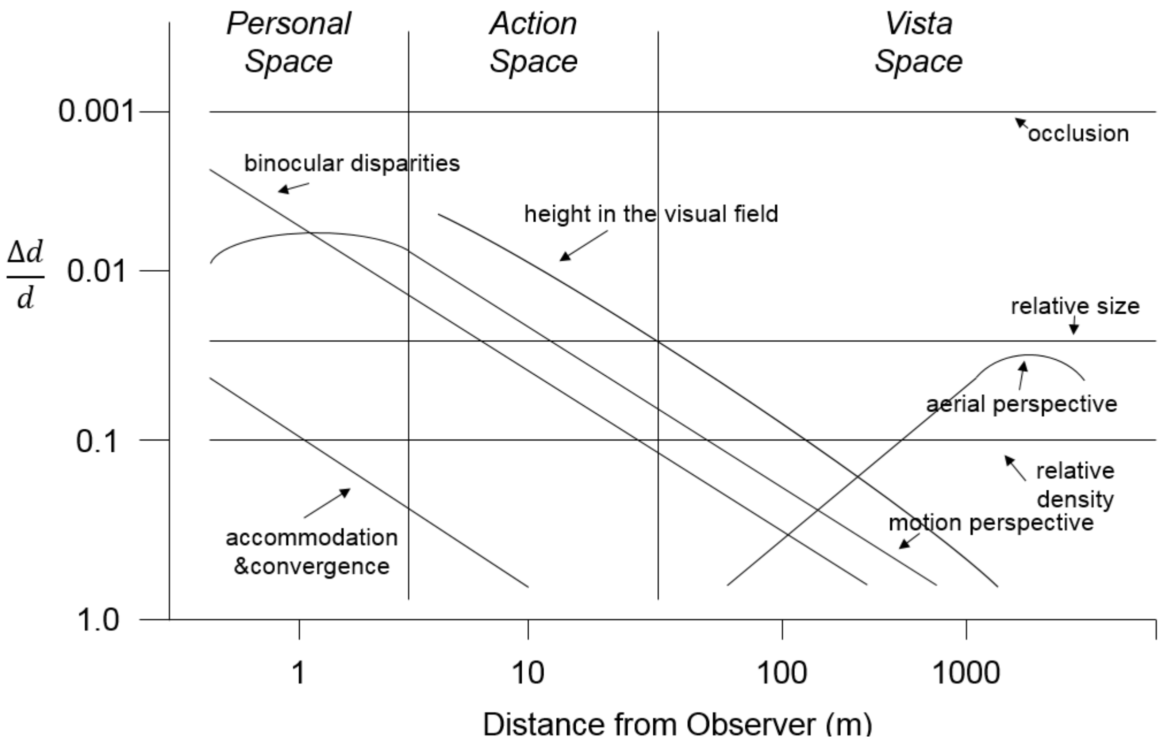

Figure 3.5.: The different depth cues and their relation to the different spaces according to Cutting and Vishton [CV95]. The functions describe just-discriminable depth thresholds according to the distance of the observer. The different depth cues each contribute differently in the personal space, the action space and the vista space.

3.3.3. Look Down Angle

The Look Down Angle (LDA) describes the angle between the horizon (0°) and the centre of the displayed image. According to law the display of static information in the HUD, such as speed information, should be displayed below 5°as it might obstruct the view on other vehicles [Isr12]. As motorists and cyclists are usually visible in the upper part of the image, it is still possible to augment the real objects in the driver's FoV. It is important that developers think about what information should be shown above 5 °to not create a distraction in the car.

3.3.4. Construction Space

The construction space behind the dashboard is extremely limited, therefore the design of the HUD system is a crucial part [Isr12]. As mentioned beforehand the size of the FoV and the VID can be adapted by the use of a magnifying lens, a concave mirror and by altering the optical path through the use of additional mirrors. Also as mentioned beforehand a larger FoV of at least 8°x 4°and a VID above 2 metres is needed to display AR visualisations. To increase the size of the FoV and the VID a larger mirror as well as a higher number of mirrors is needed. Packaging all these components behind the dashboard and under the windshield is an extremely challenging task.

3.4. Types of Head-Up-Displays

This section gives an overview of currently existing conventional displays and an outlook to the new generation of the AR HUDs.

3.4.1. Conventional Displays

The conventional automotive HUDs, currently available, can be mainly separated into two different categories: the Windshield and the Combiner HUDs.

Figure 3.6.: Example of a windshield HUD showing speed information, speed limit, lane departure warning as well as navigational information. (Courtesy of BMW)

Windshield HUD

Windshield HUDs are projecting the image directly onto the windshield. An example for this can be seen in Figure 3.6. To be able to project onto the windshield the windshield has to be specifically designed to match the corresponding HUD and vice versa to prevent distortion or double images [Neu12]. Double images can occur because the projected image is reflected twice by the windshield. To prevent this the angle of the two glass plates has to be designed in a way that the two reflected images are overlapping when they enter the drivers' eyes. Figure 3.7 exemplifies this situation.

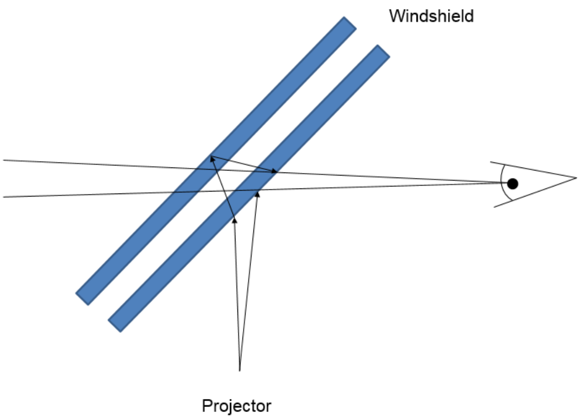

(a)

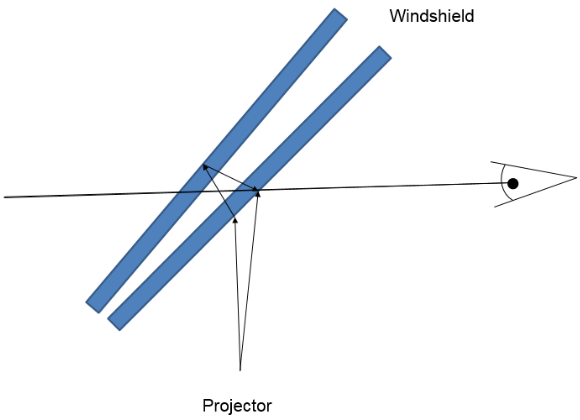

(b)

Figure 3.7.: Occurring double images and how to solve them. (a) Double image occurs, as the reflection by the glass itself does not match the projected image. (b) Both reflections are overlaid by a wedge-shaped foil to match the two reflections.

Combiner HUD

Combiner HUDs project the image onto an optical combiner showing the image without distortion or double image. The optical combiner is specifically designed to the corresponding HUD. Therefore the integration of these devices into a vehicle is in general cheaper. However, in most cases the virtual image size and the VID is in general smaller compared to windshield HUDs as the additional combiner usually needs to be incorporated into the HUD package. An example can be seen in Figure 3.8.

3.4.2. AR-HUD

Because the information is displayed right in the area where the street is visible, the possibility arises to display AR visualisations on the HUD. Early prototypes, such as the one by Sato et al. in [SKKO06] used regular projectors to be able to project an image with a large FoV onto the windhshield (see Figure 3.9). The main challenge is to integrate such an optical system into the construction space available behind the dashboard. Various car manufacturers studied the potential of AR visualisations in the HUD. [Sch08] [Isr12] However, until AR visualisations will enter series production there are a few challenges to be tackled. First of all, the image size needs to be a lot bigger than current ones. Current windshield HUDs feature a FoV of 3°x 2°and for AR visualisations a FoV of at least 8°x 4°is necessary. The VID of current HUDs is around 2 metres. The problem with displaying AR visualisations at this distance is that the eyes are focused on the near-field, whereas the objects one usually tries to augment are in the far-field. Therefore, a VID of at least 10 metres is recommended. Another solution to this problem would also be the incorporation of a stereoscopic or holographic HUD which make it possible to display the virtual image

Figure 3.8.: Example of a combiner HUD showing speed information, speed limit, road sign as well as navigational information. (Courtesy of Bosch)

Figure 3.9.: Prototypical implementation of a large FoV HUD enabling AR visualisations. [SKKO06]

at an arbitrary distance. To increase the FoV and VID the construction space is largely increased. Current HUDs feature a construction space of 2-3 litres, whereas the proposed AR HUDs would feature a construction space of 12 litres and more. As this is a large increase of the construction size, the benefits of the proposed visualisations need to outperform the required larger size.

4. Usability Testing and Perceptual Psychology

In this work we used the method of user studies to evaluate the visualisation concepts in the vehicle. This chapter introduces the method of user studies as part of the User Experience (UX). According to ISO standard 9241-210 UX is defined as ”person’s perceptions and responses resulting from the use and/or anticipated use of a product, system or service” [ISO10]. The User-Centred Design (UCD) process describes the assessment of UX and addresses issues of Human Factors and Ergonomics (HF&E). All of the afore mentioned fields are closely interconnected. Therefore this chapter introduces these fields and shows where they overlap.

Firstly the concept of the UCD process is introduced. Afterwards the term Usability and its testing are explained. Consecutively the field of Human Factors Engineering (HFE) is described and its relations to UCD is pointed out. As our studies incorporated the assessment of subjective measures and the change in gaze behaviour these terms are described in detail.

We then continue with the description of perceptual limits in psychology, in particular the phenomenon of Cognitive Capture is described. The chapter concludes with the advantages and disadvantages of conducting user studies in a real prototype compared to a simulator.

4.1. User-Centred Design

ISO 9241-210 defines human-centred design and UCD which are, also according to the standard, used interchangeably. As described in the standard, UCD does not assume a particular design process and is therefore complementary to existing design methods. According to the standard it comprises of the following steps:

1. Specify the context of use
2. Specify requirements
3. Create design solutions
4. Evaluate designs

After the last step of evaluating the application design, the found adaptations can be performed at the respective step and the application can be evaluated again. An overview of the UCD process is shown in Figure 4.1.

The last step of evaluating the design can be performed by user-based testing which can assess the usability of the design. The term Usability is described in the next section.

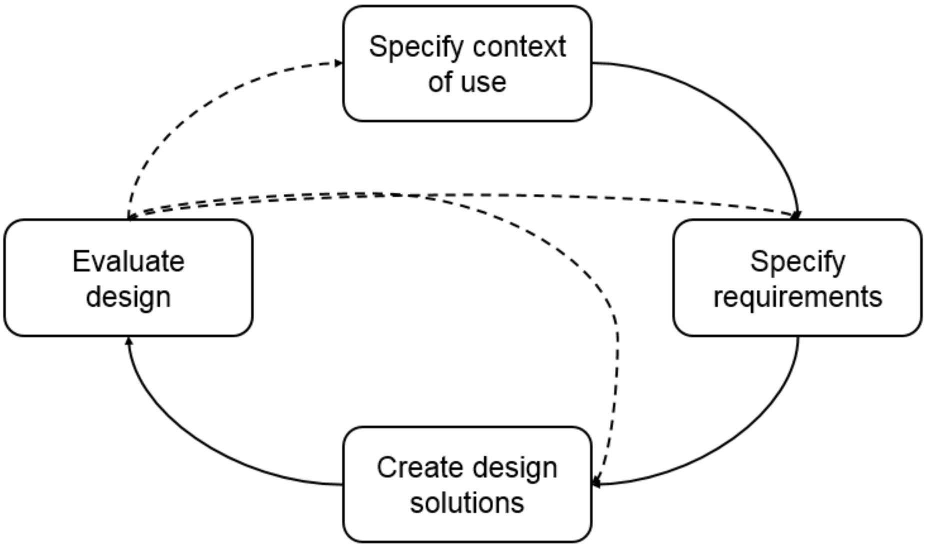

Figure 4.1.: Overview of the User Centred Design (UCD) process.

4.2. Usability Testing

According to ISO9241-11 Usability is defined as the ”extent to which a product can be used by specified users to achieve specified goals with effectiveness, efficiency and satisfaction in a specified context of use” [ISO98]. Donald A. Nielsen defines Usability by the following quality components: [Nie12]

- **Learnability:** How easy is it for users to accomplish basic tasks the first time they encounter the design?
- **Efficiency:** Once users have learned the design, how quickly can they perform tasks?
- **Memorability:** When users return to the design after a period of not using it, how easily can they reestablish proficiency?
- **Errors:** How many errors do users make, how severe are these errors, and how easily can they recover from the errors?
- **Satisfaction:** How pleasant is it to use the design?

The Usability Testing investigates these quality components. In our studies we particularly investigated the quality components of errors and satisfaction, as the goal of our visualisations is to create applications increasing the drivers' comfort and supporting them in their task, e.g. wayfinding. In the final step of the UCD process the quality components of the Usability are investigated. This is achieved by testing the application prototype on users. Usability Testing is an important step of the UCD as it is the crucial moment showing if the application is perceived by the users as intended.

As mentioned by Wickens et al. [WLLGB03] in context of the general engineering task and by Gabbard and Swan [GS08] in the specific engineering of AR applications, usability testing should be performed iteratively throughout the development process. In particular Gabbard and Swan provide guidelines and lessons learned from performing user-based studies to inform design. They point out that user-based studies are especially necessary for evaluating optical see-through devices which could also be used outdoors. The reasoning behind this is that the effects of the lighting conditions, the display technology and the context can vary and should therefore be assessed through user studies. This is one of the main reasons why we conducted user studies in the development of our new HUD visualisations.

4.2.1. Different Types of Experiments

Usability Testing can be conducted in a controlled lab experiment or in a field-study, testing the application in the intended environment. It involves the assessment of qualitative, e.g. task analysis, as well as quantitative, e.g. eye tracking, measures. By involving the assessment of these measures, an overall impression of the perception of the product can be obtained. As already pointed out for optical see-through devices, it is important to test the applications under real environment conditions.

4.3. Human Factors Engineering

Human Factors Engineering (HFE) describes the process of developing systems according to the user's needs. HFE is described in more detail by Wickens et al. in [WLLGB03]. According to Wickens et al. the main goal of human factors is to design systems which:

- Enhance(s) performance
- Increase(s) safety
- Increase(s) user satisfaction

HFE studies the factors and develops tools to facilitate the achievement of these goals. These goals lead to the concept of Usability, described in the previous section.

Related disciplines of HFE are Engineering psychology, Cognitive Engineering and Ergonomics. Part of HFE is also the UCD which describes the early focus on the user and the tasks. It also advises to use empirical measurements by the use of questionnaires, usability studies and usage studies which focus on the quantitative performance data. As in Usability Testing it is using an iterative design, making rapid changes to the design. In the design process the users are involved as part of the design team.

A major distinction is made between bottom-up vs. top-down processing. In bottom-up the perception is formed by external stimuli from the outside world. Top-down processing describes the formulation of the perception by already experienced stimuli, which result in knowledge.

The history of HFE is a lot older than the history of Human Computer Interaction (HCI) dating back to the 19th century. Frederick Winslow Taylor pioneered a method to increase the amount of coal workers could shovel more by designing the shovels accordingly. Therefore these areas emerged from different fields of psychology and now come together into a common field, addressing the user's needs at an early stage of the design process.

4.4. Subjective measures in user studies

To get an initial impression of the effect of our visualisations on the user we assessed the following measures in the user study of [WK17].

4.4.1. System Usability Scale

The System Usability Scale (SUS) was introduced by John Brooke in 1996 [Bro96]. It gives a first impression of the usability of a product. The test is easy to conduct and evaluate as it consists of 10 questions with a 5-point Likert scale. The resulting SUS score is between 0 and 100 but the Scale is not normally distributed, its overall mean is about 70 [BKM09]. As easy it is to conduct and evaluate as limited are the insights of the test. Therefore it should also be supported by other measures.

4.4.2. NASA-TLX

The NASA Task Load Index (TLX) was developed by the NASA in 1988 [HS88]. Besides the overall workload it assesses the mental, physical and temporal workload, as well as the performance, the effort and the frustration of the subject. First, the test subject rates each dimension personally after experiencing the application. In its full form each dimension is compared to each other and the test subject indicates which dimension is more important to him/her. Again, the measure alone should be supported by other measures to get a meaningful impression of the tested application.

4.4.3. AttracDiff

The last subjective measure we assessed in the user study was the AttracDiff [Has04]. It consists of 28 word pairs which are opposite to each other and the test subject rates the application on a 7-point Likert scale accordingly. Afterwards the software calculates the pragmatic and the hedonic quality, as well as the application's overall attractiveness. The pragmatic quality describes the overall usabilty and how effectively the application supports the user in achieving his/her goal. The hedonic quality describes how stimulating the application is to use and how much the user identifies him-/herself with the application. The attractiveness describes the overall quality of the application based on the quality perception. The execution of the AttracDiff is rather simple and the results are not shown in only one dimension, giving a differentiated view of the results.

4.5. Objective Measure: Fundamentals of Gaze Tracking

As an objective measure we incorporated gaze tracking in some of our user studies to gain important insights about the change in gaze behaviour. When analysing the gaze behaviour of users it is important to be familiar with the general terms connected to gaze tracking. The human perceives a situation by manually scanning a situation visually. The human eye only sees sharp images and details in the area of the *fovea centralis*. It is surrounded by the *parafovea* and the *perifovea*, in which the density of cones decreases from 50 cones/100μm to 12 cones/100μm. To better understand how humans perceive a situation visually and how this can be measured, the fundamental terms of gaze tracking are introduced, as defined in the ISO 15007. [ISO14]

Fixation. *A fixation describes the alignment of the eyes in a way so that the fixated area of interest falls on the fovea.*

Saccade. *A saccade describes the quick movement of the eyeball from one fixation to the next.*

Glance. *A glance is described by maintaining the gaze within an area of interest. It may be comprised of several fixations and saccades. The time inside the area of interest is described as glance duration.*

Area of Interest (AOI). *A predefined area in the environment to compute metrics towards it.*

With these initial definitions it is possible to define metrics which can be used to assess the visual attention towards an Area Of Interest (AOI). The following metrics were assessed in our experiments:

- **Number of Glances:** Describes the number of times an AOI was entered and exited by the visual gaze.
- **Glance Duration:** Describes the duration between the moment the visual gaze enters an AOI and it exits the AOI.
- **Total Glance Duration:** Describes the accumulated duration over the whole measurement of a particular AOI.
- **Mean Glance Duration:** Describes the arithmetic mean of the glance duration over all glances towards an AOI.

4.6. Perceptual Limits

This section introduces the effect of Perceptual Tunneling and more specifically the effect of Cognitive Capture, a phenomenon we assessed in our studies.

4.6.1. Perceptual Tunneling

Perceptual Tunneling or Inattentional Blindness describes the effect when a user is too focused on a particular event, not recognising the environment around him/her. A famous experiment condcuted by Daniel J. Simons and Christopher F. Chabris [SC99] investigates that participants did not recognise a gorilla walking through a scene, because they were too much distracted by a visual task. In our work we addressed a certain type of Perceptual Tunneling described in the next section.

4.6.2. Cognitive Capture

Cognitive Capture is a certain type of Perceptual Tunneling, in which the user is too much involved in a certain task and therefore not present in the environment. This type of Perceptual Tunneling is in particular interesting when designing new HUD visualisations in a way that they support the user and do not create another distraction inside the vehicle. As the visualisations investigated in this work are of dynamic nature, we investigated the Perceptual Tunneling taking place by investigating the user's gaze behaviour and later also measuring the user's distraction by the involvement of a Detection Response Task described in chapter 11.

4.7. Car prototype vs. simulator

Quite a number of studies have been researched in the area of automotive AR (e.g., [KD09]; [MKPP11]; [SSBG15]). Most of the research so far was conducted in car simulators. The use of simulators has several advantages and disadvantages. First of all their development is usually more price efficient as the development of the simulator itself and the development of the applications is a lot cheaper. Another advantage of car simulators is that it is possible to control the virtual environment. This is important in case one wants to control the error of a particular application or if it is dangerous situations which are to be investigated. But even with car simulators becoming more sophisticated and a lot cheaper, they will never be able to simulate all the effects of a real driving situation. Therefore at some point of the development stage it is necessary to implement the applications in a real-life car prototype to be able to assess all the effects, occurring in real-driving situations. This moment has arrived for the AR HUD visualisations. The control of the environment is of course not possible in real traffic situations. As AR applications have already been thoroughly investigated in car simulators, this work takes the important steps of porting the AR applications into a real car prototype and performing user studies of participants using the visualisations in real traffic situations.

5. Related Work

This chapter introduces general related work in the field. More specific related work is described in the respective chapters.

5.1. AR in vehicles

This section surveys the different papers on AR in vehicles. In the vehicle the driver has to attend to different types of tasks which can be divided in primary, secondary and tertiary tasks as described by Tönnis et al. [TBK06]. With the introduction of AR a whole range of opportunities as well as challenges arose for using it inside a vehicle. A good summary of these are described by Gabbard et al. [GFK14]. The topic of AR in the vehicle has been around for some time. A rather traditional AR visualisation regarding navigation is the three-dimensional arrow on the road surface, as described by Tönnis et al. [TKK08]. A number of interesting navigation visualisations arose making use of AR, such as the visualisation presented by Kim et al. [KD09]. In this visualisation a two-dimensional map is merging into the three-dimensional environment. Conventional navigation was compared to the mentioned AR arrow visualisation as well as AR highlighting of landmarks in the real environment, as desribed by Bolton et al. [BBL15]. All of the mentioned user studies were performed inside car simulators. A necessary step is to perform user studies in actual car prototypes to be able to take effects into account which affect the perception of AR visualisations.

One possibility of displaying AR visualisations in the car is the Head-Up-Display (HUD). The HUD projects the virtual information on the windshield giving the driver the perception the virtual information is displayed embedded into the real environment. Common principles for presenting information inside HUDs is described by Tönnis et al. [TKP09]. The needed GNSS accuracy for displaying a three-dimensional arrow correctly at an intersection was determined in a car simulator by Pfannmüller et al. [PWB15]. As the field of view of HUDs is relatively limited Tönnis et al. presented an AR representation for directing the driver's attention towards hazards outside the visible field of view [TK06]. The effects of HUDs on the driving and task performance have been analysed in a simulator based user study by Smith et al. [SSBG15]. Kim et al. compared different hazard warning cues in the HUD [KWGP13]. Charissis et al. presented an interface for informing drivers about the upcoming traffic under adverse weather conditions [CPMA11]. Medenica et al. compared different visualisation methods and the AR visualisation showed the least negative impact on driving [MKPP11]. A laser display projecting onto the windshield was investigated by Doshi et al. [DCT09]. Kawamata et al. presented a depiction of the future road course in a three-dimensional way in the car [KKKO13].

5.2. Perception of AR visualisations

The basic sources of depth perception were surveyed by Cutting and Vishton [CV95]. Several studies have been conducted to determine the perception of virtual objects in AR. [SLS+06] [SSE15] [DWSS17] The depth of the virtual objects is usually underestimated. Furthermore, with the introduction of stereoscopic 3D interfaces new possibilities also in the automotive sector arose. The potential and the perception of stereoscopic 3D interfaces has been analysed by Broy. [Bro16] Also with AR it becomes possible to create a new kind of X-ray vision, giving the user the ability to see objects occluded by other real objects. As this is both, a very helpful feature of AR but also very unnatural for users' perception, studies have also been conducted studying the perception of these visualisations. [LSG+03] Creating AR visualisations in a manner such that they are perceived as natural objects is still a challenging task. Reasons for this include that it is necessary to position the visualisations precisely and steadily in the environment. For this the positioning as well as the display latency need to be improved.

5.3. Latency in the AR context

The temporal aspect is crucial for optical see-through AR solutions, as the image of the real environment cannot be delayed in order for the virtual object to be aligned correctly with the real environment. Therefore it is important to always keep the end-to-end latency of the system in mind and find ways of measuring it. One measuring technique used for determining the overall latency of a system is described in Billeter et al. [BRW+16]. When using several sensors it is important to temporally calibrate these sensors to each other, as described by Huber et al. [HSK14]. A prototype of rendering AR visualisations with off-the-shelf hardware with latency below 1 ms is described by Itoh et al. [IOH+16].

5.4. Visualisations based on the electronic horizon

The electronic horizon describes what the vehicle is able to perceive with its sensors and enabled a whole range of new driver assistance systems. By knowing the future road course these assistance systems can react in a more advanced way. Its main concept is described by Ress et al. [REKB06]. Different visualisation concepts were shown for smart deceleration by Nestler et al. [NDPR09]. Duschl described various visualisations of the electronic horizon in the instrument cluster [Dus15].

6. Contributions of this Dissertation

This chapter summarises the main contributions of this work.

1. Analysis of current GNSS systems in vehicles: The positional error of current navigation systems in cars is measured by comparing it to two high-precision GNSS systems. After selecting a reference route the measurement is performed and afterwards an analysis is made to investigate the root mean square (RMS) error is calculated between the measurements.

2. Development of a visualisation concept taking the positional error into account: Based on the measurement of the positional error a new visualisation scheme is developed. This visualisation scheme does not require high-precision global localisation, yet conveying a visualisation interacting with the environment and taking place in three-dimensional space.

3. Development of a visualisation of the future road course based on the electronic horizon: A visualisation of the future road course in the HUD is developed. This visualisation should help drivers in foreseeing the road geometry of the future path, which can be obstructed due to weather conditions or sharp bends of the road. Because of the limited Field of View (FoV) of the HUD a symbolic representation was chosen, which interacts with the environment dynamically. The visualisation is based on the electronic horizon, therefore not requiring additional data to generate the visualisation.

4. Assessment of the change in gaze behaviour due to HUD visualisations: The change in gaze behaviour is assessed in a user study through the use of eye tracking glasses (ETG). In a first experiment the change in gaze distribution of the head unit, the instrument cluster and the area of the HUD between the HUD being switched on or off is investigated. After an initial assessment a refined experiment is conducted to investigate in particular the change introduced by dynamic visualisations in the HUD.

5. Latency analysis of the current HUD visualisation and experimental assessment of perceivable latency: The end-to-end latency of the current HUD is measured by displaying and processing a timestamp in an image. After the overall latency is determined the perceivable latency in the context of moving virtual objects and the appearing of virtual objects is investigated through user studies.

Part II.

Location inaccuracies in AR navigation

This part first assesses the accuracy of current Global Navigation Satellite Systems (GNSS) sensors used in series production vehicles. The use of GNSS sensors is one of the techniques needed to steadily position an AR visualisation in the real environment. Furthermore a high-precision map, as well as other techniques such as terrain recognition are needed in order to position a virtual object in the environment in a sophisticated manner. The results of the assessment show that the current GNSS sensors do not achieve the required accuracy to place an AR visualisation steadily on the road. Based on the determined accuracy a novel AR visualisation is developed, which is able to handle the determined errors of the global positioning, yet conveying the impression to the user that the visualisation is taking place in three dimensions. This novel AR visualisation as well as a traditional AR visualisation were implemented in a car prototype in order to compare these two visualisations in a user study. The user study revealed that, overall the novel visualisation was preferred without adding additional workload on the driver or decreasing its usability.

7. GNSS Analysis

This chapter describes the analysis of the global positioning accuracy by current car navigation systems. Parts of the analysis were published in [WK17]. The global position is acquired by the use of a satellite based positioning system, for which Global Navigation Satellite Systems (GNSS) is the general term. Examples for these systems are the American Global Positioning System (GPS) and the Russian GLONASS system. Many others, like the European Galileo system are currently in development.

The chapter begins with an initial analysis of the GNSS accuracy. In this initial analysis two test drives were conducted, one with a high-precision GNSS sensor and a second one consisting of another high-precision GNSS sensor and the vehicle's built-in GNSS sensor. The initial analysis introduced inevitable errors due to the difficulty in reproduction of the exact same conditions, e.g. weather conditions or small deviations when driving the route a second time. Therefore we conducted a refined analysis, in which all the GNSS devices were recording the same route at the same time. With these measurements we gained a better understanding of the actual accuracy of current GNSS sensors used in series production vehicles.

7.1. Introduction

The global position is determined by measuring the time a signal needs to travel from the satellite to the receiver on the ground. Hence for each satellite the possible locations of the receiver are on a line of a circle on the earth's surface and a minimum of three satellites is needed to successfully determine the position of a receiver on ground. The quality of the signal may be distorted by weather conditions and reflections or absorptions by nearby buildings. So far the American GPS and the Russian GLONASS system are in use. The European Galileo system is currently under development.

In order to display AR navigational visualisations correctly in the environment, accurate knowledge of the vehicle's surroundings is needed. Besides a precise assessment of the vehicle's global position, accurate maps as well as recognition of the terrain are also required. This part assesses the accuracy of current GNSS sensors used in cars. Pfannmüller et al. determined in [PWB15] through a user study conducted in a driving simulator the acceptable threshold of the misplacement of an AR visualisation for navigation purposes. They found that a global deviation of 6 metres led to higher navigation errors when used with an AR visualisation. Therefore, the global deviation is assessed according to this threshold.

7.2. Initial Analysis

To determine the accuracy of current GNSS sensors we first identified a route consisting of the standard characteristics of a common route. The route consisted of highways, rural as well as city roads. Furthermore it consisted of turns, roundabouts and changing slopes. The route was about 10 minutes in driving time. In a first step the route was measured with a high-precision differential GPS (dGPS) system by JAVAD.

A dGPS system improves the precision of GPS receivers through the support of ground stations, of which the precise location is known beforehand and previously determined by more classical methods. As the position of the ground station is known beforehand, the deviations in the timings of the satellites and the delay, introduced by atmospheric changes and deviations in the clocks of the satellites, can be calculated. In Europe the EGNOS system has been in operation since October 2009. The EGNOS system consists of three geostationary satellites and a network of ground stations. [EGN]

The first measurement served as a reference to establish the accuracy of another high-precision GNSS system, which was the racetracker Racelogic VBSS100_V3. After the accuracy of the VBSS100_V3 was established, we used it as a reference for determining the accuracy of the vehicle's built-in navigation system. Through these steps we determined a first estimate of the accuracy of nowadays GNSS systems used in current series production vehicles.

7.2.1. First Measurement with differential GPS

The first measurement was performed using a dGPS system by JAVAD. For increasing the position accuracy it uses the Real-Time Kinematic (RTK) technology. RTK works on the principle of having stationary broadcast stations, whose precise positions are known beforehand. The station measures the time required for the signal to travel from the satellite to the station. As the position of the station is known in advance, the delay of the signal due to current atmospheric circumstances or offsets of the satellites' clocks can be calculated. This delay is then distributed to GNSS sensors in the region which can use these correction factors to improve their position.

The result of this measurement was a precise measurement of the reference route, which can be seen in Figure 7.1.

7.2.2. Second Measurement

For the second measurement we used the Racelogic VBSS100_V3, which is a commercially available racetracker offering a higher precision compared to conventional vehicles' GNSS sensors. At the same time the measurement of the car's navigation system was recorded. After completion we compared the Racelogic VBSS100_V3 measurement to the previously recorded route by the JAVAD system, to investigate its validity of serving as an internal reference.

In order to compare the measurements to each other a number of necessary steps were performed. First, the geographic coordinates, given in degrees of latitude and longitude, needed to be transformed into a Cartesian coordinate system, in our case measured in metres. To do so a reference point for the Cartesian coordinate system needed to be de-

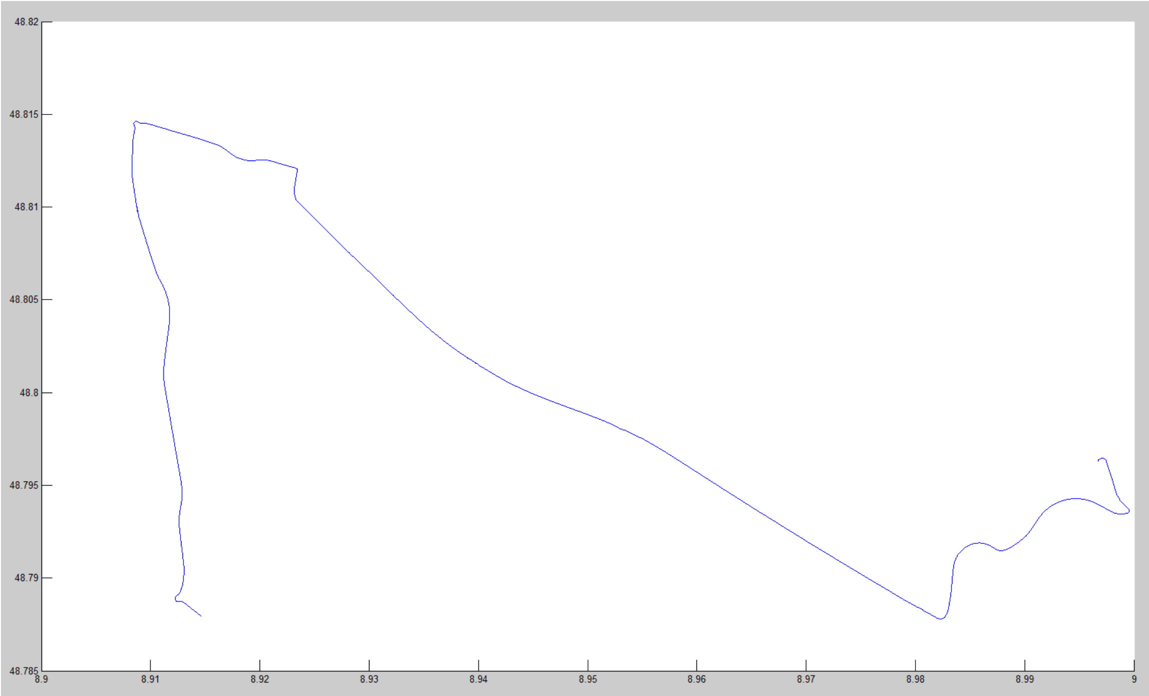

Figure 7.1.: Measured route (x-axis: longitude, y-axis: latitude)

termined, around which the Cartesian coordinate system was constructed. In order to be able to compare the measurements an equal number of data points is needed. According to the method of [HSK09] it is advisable to interpolate the dataset with the higher number of data points to the number of points in the smaller dataset.

Because these were two independent measurements, the minimal distance for each data point of the VBSS100_V3 dataset had to be calculated to all data points in the JAVAD dataset. The resulting root mean square error (RMSE) for the minimal distances was derived as 2.84 m (arithmetic mean=2.4 m, std.dev.=1.52 m).

However, the route taken in the two measurements was not precisely the same. The small deviations while driving the route a second time and changes in the weather conditions, resulted in different travel times of the satellite signals. Due to the above calculation, we considered the VBSS100_V3 measurement as a precise internal reference for the vehicle's built-in navigation system.

7.2.3. Comparison to the conventional Car GPS

As the last step we compared the measurement of the Racelogic VBSS100_V3 with the vehicle's navigation system. As the measurements were taken at the same time we could measure the distance between two data points with the same timestamp. However, the sampling frequency of the two devices was different. Therefore, we interpolated the dataset with the higher number of data points (the VBSS100_V3) to match the number of data points in the smaller dataset (the vehicle's navigation system). We found the RMSE of the distances between the measurements to be 16.13 m (mean=14.92 m, std.dev.=6.14 m) as

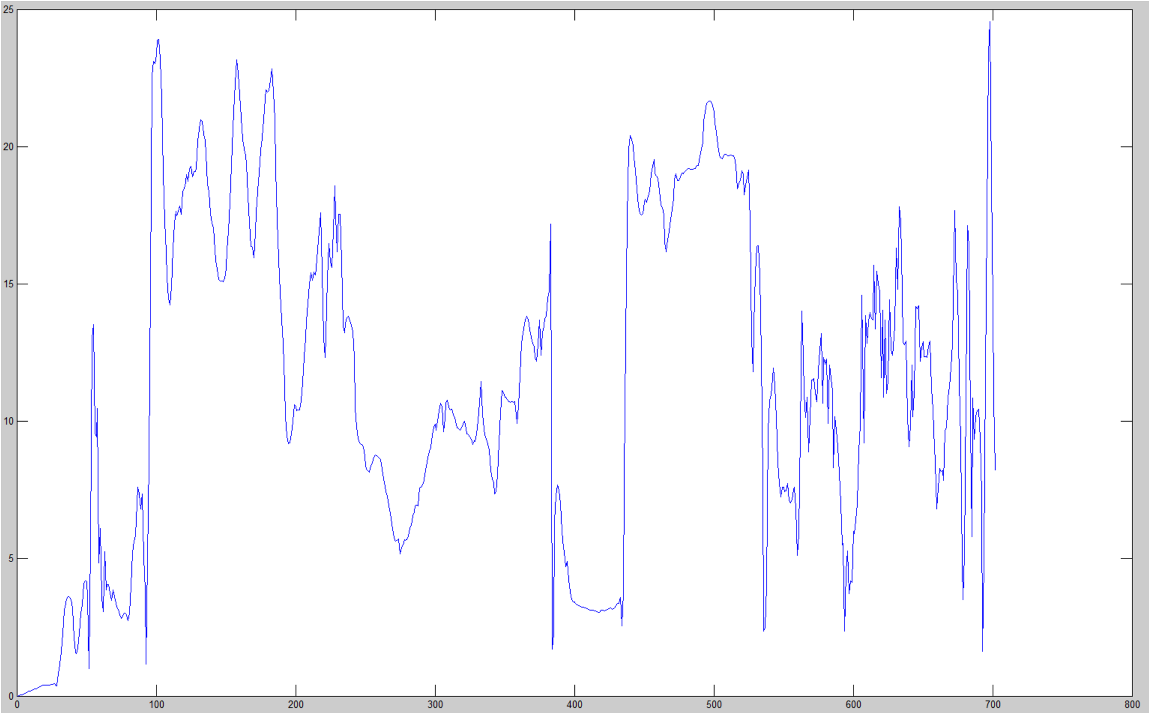

Figure 7.2.: Positional error between the VBSS100_V3 and the car's navigation system (x-axis: timestamps, y-axis: distance between the two measurements)

can be seen in Figure 7.2.

This measurement gave us a first hint at the actual GNSS accuracy encountered during driving with a conventional vehicle navigation system. We investigated the GNSS accuracy further which is described in the next section.

7.3. Refined Analysis

Through the initial measurement we gained a first understanding of the actual global positioning error introduced by current GNSS sensors used in series production vehicles. However, the initial measurement comprised of a few flaws, the most significant was that the measurements of the GNSS devices took place in two independent test drives. For this reason we conducted a refined measurement, this time the devices recording at the same time.

Similar to the previous measurement the dGPS system by JAVAD was the one with the highest accuracy, the VBSS100_V3 the second highest accuracy and the car navigation system was the system under test. We drove the same route as in the previous measurement with the three different GNSS sensors recording the route at the same time.

7.3.1. Analysis of the GNSS Measurement

After the measurement, the recordings of the different GNSS devices were compared to each other. Similarly to the previous measurement the geographic coordinates needed to

Table 7.1.: RMSE between measurements

	RMSE (in metres)
VBSS100_V3/ JAVAD dGPS	3.0397
JAVAD dGPS/Car navigation system	6.0332
VBSS100_V3/Car navigation system	5.533

be transformed into a Cartesian coordinate system. As the original measurement results were given in degrees of latitude and longitude, the measurement point having the minimal Euclidean distance between the two measurements was determined. This measurement point then served as the reference point for the Cartesian coordinate system. Furthermore, as in the previous measurement the dataset with the higher number of data points needed to be interpolated to the one with the fewer number of data points. As the measurements were taken at the same time we interpolated it to the respective timestamp. Finally we calculated the RMSE between the GNSS sensors and obtained the results which can be seen in Table 7.1 and in Figures 7.3, 7.4a and 7.4b. The RMSE between the JAVAD dGPS and the VBS100_V3 was about 3 metres, whereas the the RMSE between the JAVAD dGPS and the vehicle's navigation system was about 6 metres. The RMSE between the VBSS100_V3 and the vehicle's navigation system was lower with about 5.5 metres.

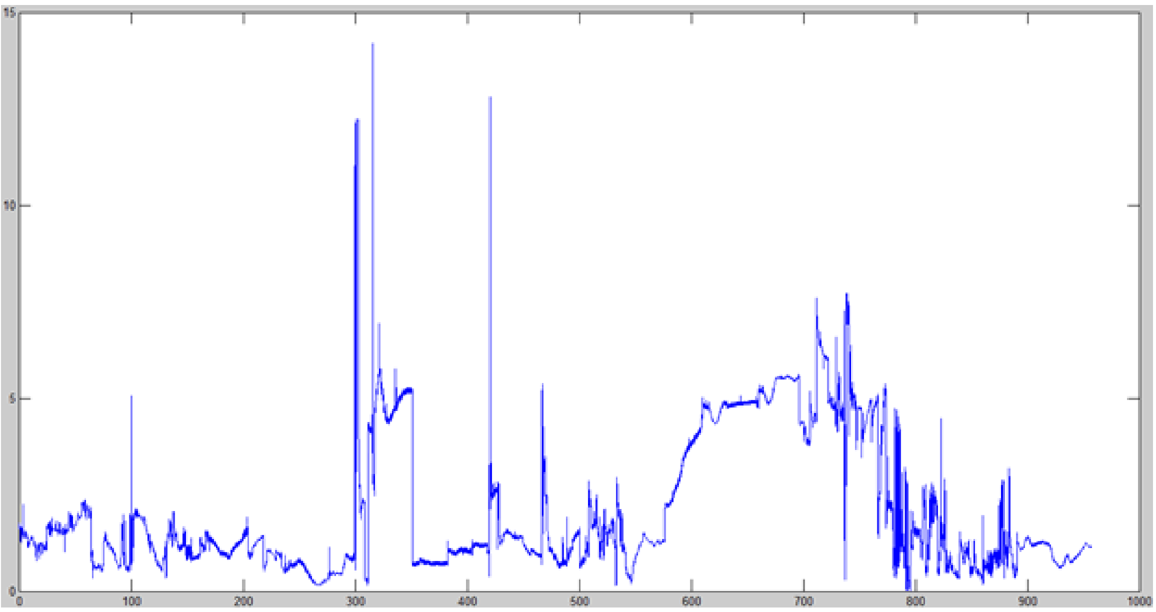

Figure 7.3.: Positional error between the VBSS100_V3 and the JAVAD dGPS in the improved measurement (x-axis: timestamps, y-axis: distance between the two measurements)

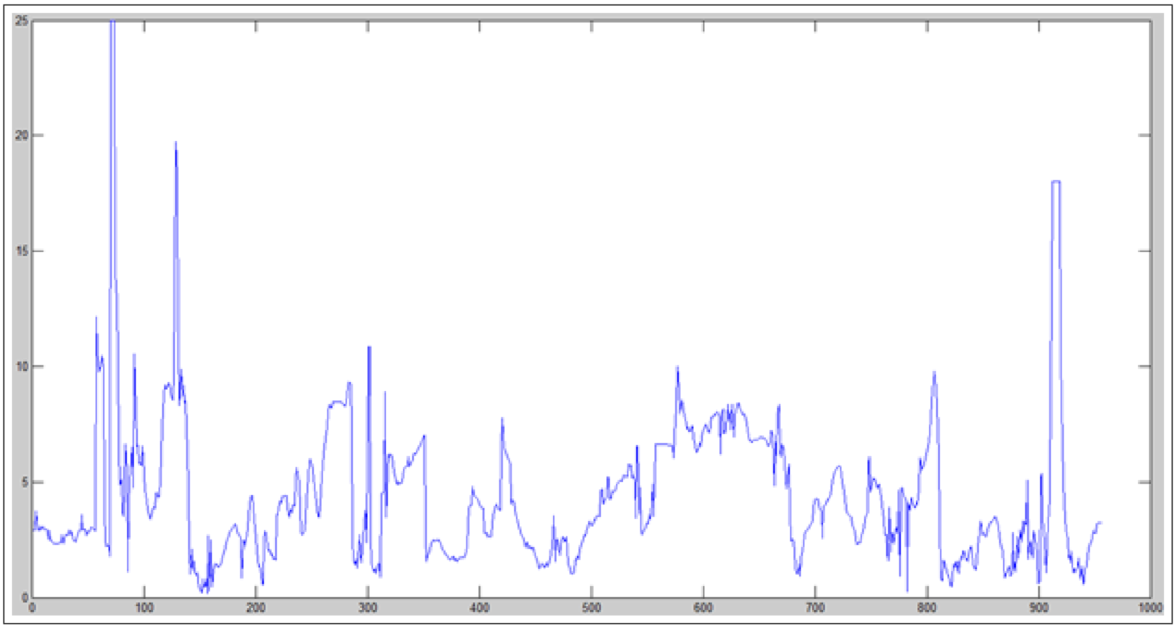

(a) Positional error between the JAVAD dGPS and the car's navigation system in the improved measurement

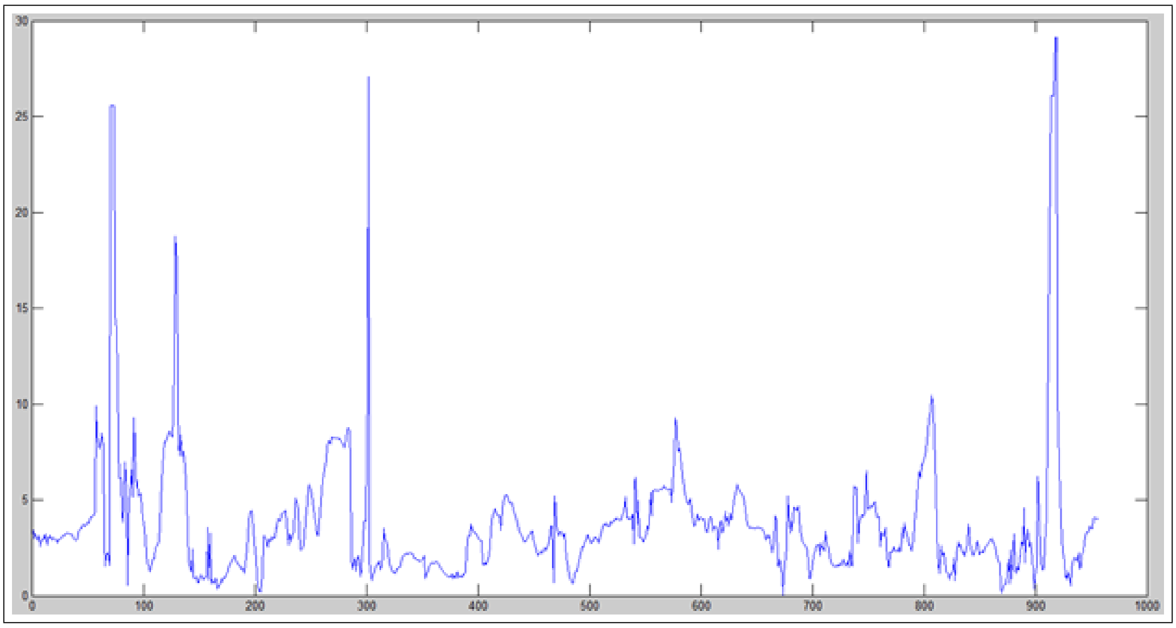

(b) Positional error between the VBSS100_V3 and the car's navigation system in the improved measurement

Figure 7.4.: Positional errors between the measurements (x-axis: timestamps, y-axis: distance between the two measurements). The large errors between the measurements especially take place in areas the reception of the satellites' signals are sparse and JAVAD dGPS and VBSS100_V3 are superior to the car's navigation system

7.3.2. Discussion

The deviation between the VBSS100_V3 and the car's navigation system was found to be lower than the deviation between the JAVAD dGPS and the car's navigation system. Therefore, we looked for an explanation for this phenomenon. First, we investigated the specifications of the sensors. The specification of the JAVAD dGPS system specified an accuracy of below 0.5 metres (see Figure 7.5). According to the specification of the VBSS100_V3 system the accuracy is about 3 metres if the dGPS mode is disabled and below 1 metre if it is enabled (see Figure 7.6). For the car's built-in navigation system we didn't have the specifications, but we assumed this device to have the least accuracy as it was a cost-efficient solution.

Features/Receiver Type	Sigma			SigmaD			SigmaQ
	G2T	G3T	G3TAJT	G2	G2D	G3D	
Autonomous Accuracy	<2m						
Static, Fast Static Accuracy	Horizontal: 0.3 cm + 0.5 ppm * base_line_length Vertical: 0.5 cm + 0.5 ppm * base_line_length						
Kinematic Accuracy	Horizontal: 1 cm + 1 ppm * base_line_length Vertical: 1.5 cm + 1.5 ppm * base_line_length						
RTK (OTF) Accuracy	Horizontal: 1 cm + 1 ppm * base_line_length Vertical: 1.5 cm + 1.5 ppm * base_line_length						
Real-time heading accuracy	-			~ 0.004/L [rad] RMS*			
Roll/Pitch	-						~ 0.008/L [rad] RMS*
DGPS Accuracy	< 0.25 m Post Processing, < 0.5 m Real Time						

Figure 7.5.: Specifications of the JAVAD dGPS.

Absolute Positioning				
Accuracy	5m*	3m*	3m*	3m*
Accuracy with SBAS DGPS				
-Europe (EGNOS)	N/A	<1m*	<1m*	<1m*

Figure 7.6.: Specifications of the VBSS100_V3.

For this reason we assumed that the dGPS mode of the VBSS100_V3 was disabled. We then confirmed this by investigating the CAN trace of the vehicle. Figure 7.7 shows that the dGPS mode of the VBSS100_V3 was in fact disabled. Therefore the RMSE of 3.04 metres between the VBSS100_V3 and the JAVAD dGPS is according to the specifications. The lower RMSE of 5.53 metres between the VBSS100_V3 and the car's navigation system compared to the RMSE of 6.03 metres can be explained due to the fact that the JAVAD dGPS system is the only system using a correction service. Figure 7.8 shows an example

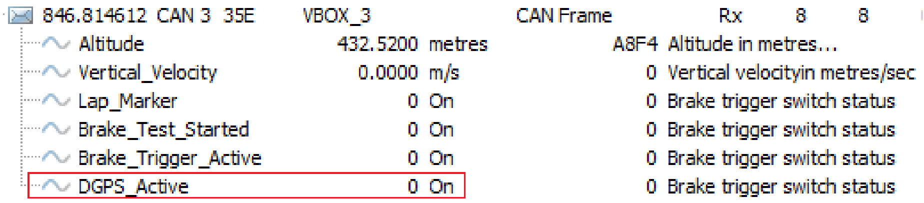

Figure 7.7.: CAN trace of VBSS100_V3 measurement.

in which the determined position by the VBSS100_V3 and the car's navigation system is closer than the determined position by the dGPS and the car's navigation system. As both, the car's navigation system and the VBSS100_V3 are not using any correction service both systems are introducing the same type of error, resulting in a closer distance compared to the JAVAD dGPS system.

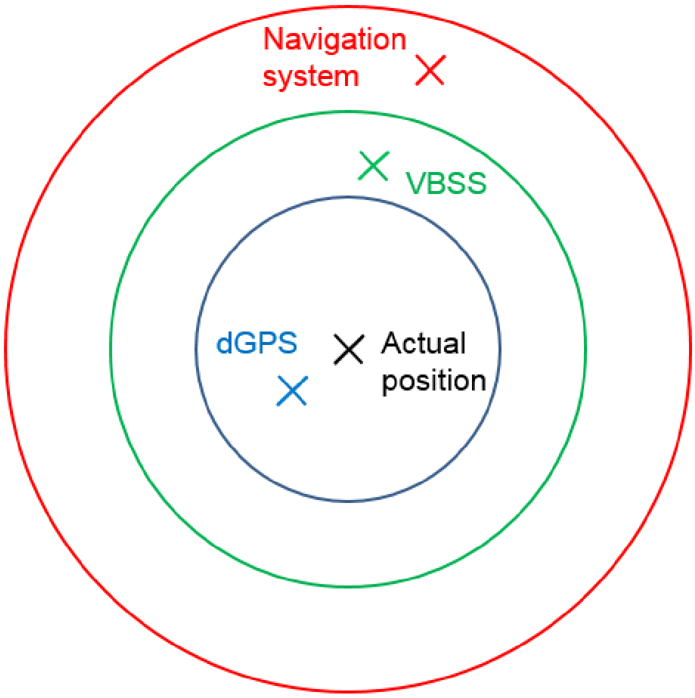

Figure 7.8.: The determined positions of the JAVAD dGPS (blue), the VBSS100_V3 (green) and the car's navigation system (red). The sketch exemplifies if the VBSS100_V3 and the car's navigation system introduce the same kind of error the distance can be actually lower than the distance between the dGPS system the car's navigation system.

7.4. Conclusion

In this chapter we described the assessment of GNSS sensors' accuracy currently used in series production vehicles. The initial measurement found a RMSE of 16.13 metres. As the measurements were taken in two different test drives, we conducted a refined measurement recording all GNSS devices at the same time. The RMSE of this measurement was found to be 6.03 metres. In combination with the results of [PWB15] this shows that the global positioning error is currently too high to be able to position a navigational hint, such as a three-dimensional arrow, precisely at an intersection for the user to have the impression it is positioned there and to be able to successfully navigate through traffic. Therefore, the next chapter will introduce a novel visualisation concept for the HUD which is able to handle the determined GNSS error.

8. Augmented Reality Visualisation overcoming location inaccuracies

The findings of the GNSS analysis showed that the accuracy of current car navigation systems is not sufficient for the positioning of AR graphics at a particular global position. For this reason we developed a novel visualisation concept which does not require a precise global localisation, and yet conveys to the driver the visualisation is part of their environment. This visualisation is described in detail in this chapter, first results of the visualisation were published in [WK17].

After designing the appearance and the behaviour of the visualisation, we implemented this visualisation in a car prototype. In the car prototype we had also implemented a rather traditional AR visualisation. We compared the two visualisations in a user study to gain insight about the advantages and disadvantages of the two visualisations.

8.1. Introduction

A common problem with path finding is that it is problematic to figure out which turn to take, as numeric distance depictions are difficult to relate to actual distances. This is especially the case if several intersections are close to each other. As identified by Gabbard et al. [GFK14] AR can help the driver in these cases in choosing the correct turn.

As described in chapter 7, the GNSS error of current series vehicles is too high to be able to place an AR visualisation steadily at an intersection. Furthermore, as revealed by Pfannmüller et al. [PWB15] the deviation of 6 metres led to a significantly higher number of navigation errors. For these reasons we developed a new AR visualisation which takes the positional error into account. We developed the visualisation in a way, it does not require a precise localisation and still conveys to the driver it is registered in three dimensions and correlates with the environment.

We implemented this visualisation in a car prototype equipped with a prototypical version of an HUD. We also implemented the traditional three-dimensional arrow on the road in the car prototype. With these two visualisations we performed a user study in our engineering process, as suggested by Gabbard and Swan [GS08]. The user study was performed to investigate the effect of the visualisations on drivers in actual traffic situations. It should be noted that we used a GNSS system with a higher accuracy to focus on the perception of the visualisations.

8.2. Related Work

Tönnis et al. distinguished between primary, secondary and tertiary tasks [TBK06]. Navigational information is usually displayed mainly in the area of the tertiary tasks, which leads to a relatively high distraction. Displaying and conveying the navigational information in the area of the primary tasks, where the display area of the HUD is located, has the potential to decrease the inattention by the driver.

Wayfinding and navigational aids are opportunities for the utilization of AR, as identified by Gabbard et al. [GFK14]. However, localising the absolute position and the orientation of the car remains a major challenge.

Different visualisations for AR wayfinding have already been proposed. Examples include a three-dimensional arrow [TKK08] of which we implemented a similar version in our car prototype. Kim et al. described a navigation visualisation which was implemented in a car simulator [KD09]. This visualisation shows the route on a two-dimensional map which merged into a three-dimensional representation of the road. Although this visualisation significantly reduced the number of navigational errors, it would need a much bigger field of view than currently available in HUDs. Bolton et al. compared, using a driving simulator, conventional navigational hints with an AR arrow visualisation and an AR highlighting of landmarks in the real environment (e.g. a public house) [BBL15]. The AR visualisation outperformed the conventional navigation hints, but the proposed AR landmark visualisation performed the best. As these studies showed an improvement in the navigational task, we implemented our visualisation in a real car prototype to investigate the effect on drivers under real traffic situations.

Tönnis et al. identified common principles for presenting information in HUDs [TKP09]. They classified them in dimensions and used pair-wise combinations of these to illustrate examples. One of the examples was the comparison between contact-analog (i.e. correctly world-registered) and an unregistered representation. As they pointed out, AR visualisations should never be distorted, but correct registration is not always necessary. Pfannmüller et al. analysed the needed GNSS accuracy for displaying an arrow correctly on the road and its effect on users' ability to successfully navigate through traffic [PWB15]. They concluded a positional error of 6 metres led to significantly higher errors in navigation. We developed an AR navigational visualisation which is not registered in the world environment to cope with the localisation inaccuracies.

8.3. Navigation Concept taking positional Errors into Account

We developed a visualisation which is able to cope with the reported GNSS inaccuracy and does not require terrain detection. Therefore the visualisation is able to handle a lower GNSS accuracy and does not require additional sensors, such as cameras or radar, creating a low-cost and robust implementation. We call this visualisation the Sails visualisation. This visualisation is not shown continuously, instead the "Sails" are appearing and disappearing as the vehicle moves through the intersection, creating the impression of a virtual crash barrier. The visualisation is only dependent on the distance to the next turn.

A rough overview of the system can be seen in Figure 8.1. The global position, as well as the heading of the car is received by a GNSS unit. This is then combined with the map data in the main computation unit. In our prototype we used a car PC for this purpose, while in the case of a series vehicle this would be an Electronic Control Unit (ECU) using embedded hardware. After the visualisation is rendered it is displayed on the HUD of the vehicle. The

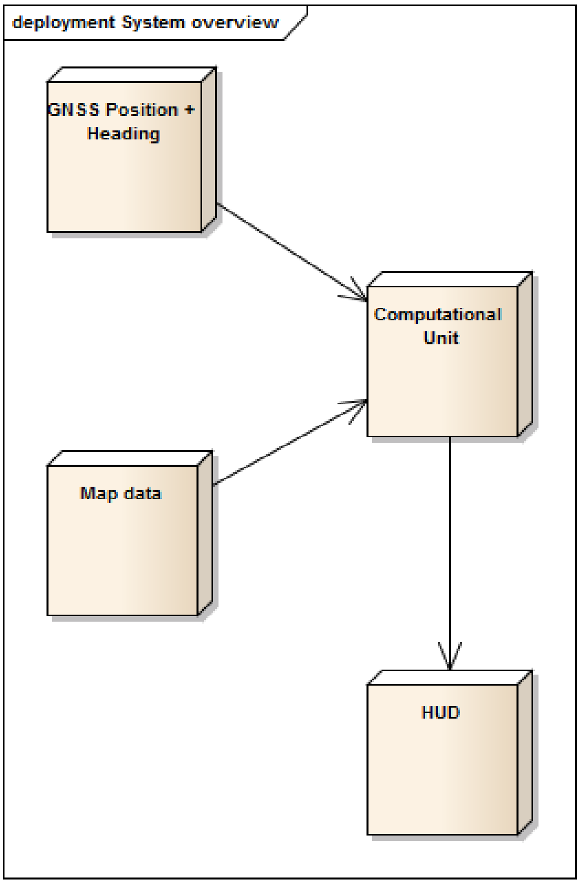

Figure 8.1.: System overview of the involved systems in the Sails visualisation.

visualisation does not require additional hardware besides the afore mentioned.

Very importantly, the visualisation is registered with respect to the car coordinate system rather than with respect to the world coordinate system. Therefore the visualisation creates the impression the visualisation is part of the vehicle. The spatial relationship graph (SRG) in Figure 8.2 illustrates this relation. This visualisation conveys a three-dimensional impression to the user. In the following the design of this visualisation is described in detail. [1]

When the car is approaching the intersection the Sails appear on the side opposite to the turning direction. As the car further advances towards the intersection, the Sails are spreading out and bending in a curve to convey to the driver in which direction and when to turn the car. Once the car reaches the intersection, the Sails are fully bent. At the exact location of the intersection the Sails change their colour to convey to the driver the exact moment of the turn. This is especially important if there are several intersections next to each other. In a last step the Sails disappear to the side opposing the turning direction, giving the driver the impression of driving past them. The different states and the transition can be seen in Figures 8.3 and 8.4. In case the driver does not steer in the direction of the indicated direction, the Sails simply disappear.

[1] Filed for patent DE102016203080 [KMW16]

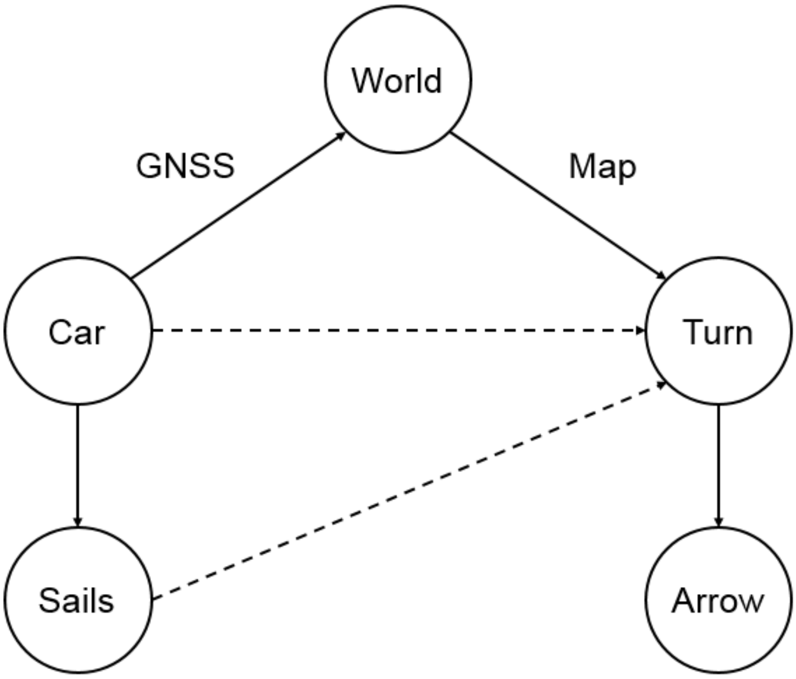

Figure 8.2.: Spatial relationship graph for Sails/Arrow representation.

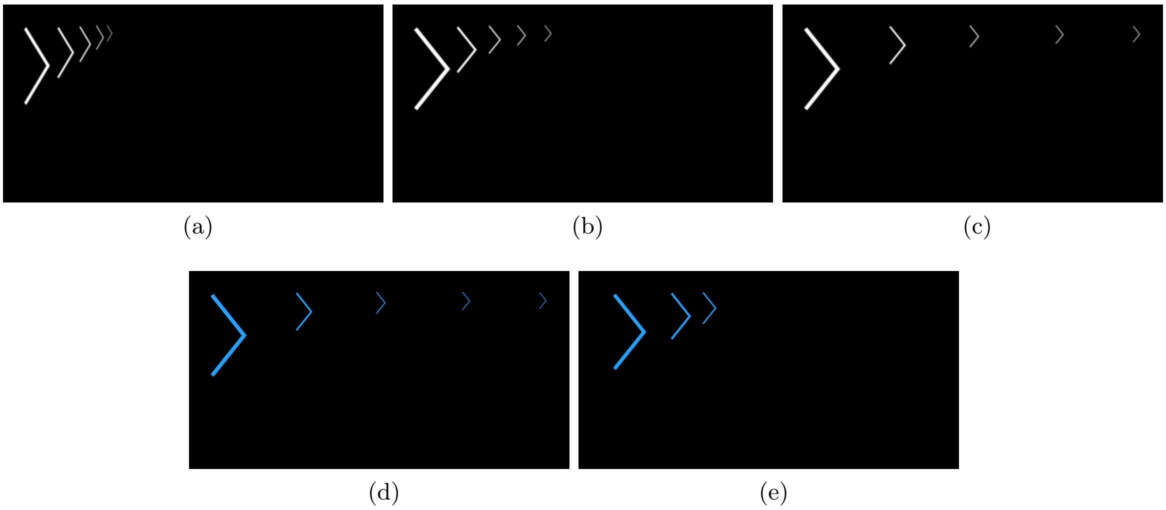

Figure 8.3.: The five different states of the ”Sails” representation. (a) The sails appearing on the side opposing the turning direction to notify the driver of an upcoming turn. (b) The sails bend towards the turning direction while approaching the intersection. (c) The sails are fully bent to convey the turning direction. (d) The sails change their colour to indicate an immediate turn. (e) The sails are disappearing to the side opposing the driving direction, giving the driver the impression of driving past them.

We implemented this visualisation in our car prototype, which is equipped with a prototypical version of an AR HUD. We also implemented a rather traditional AR visualisation, a three-dimensional arrow, in our car prototype to compare it against our novel visualisation. The two different visualisations can be seen in Figure 8.5. These two visualisations were compared in a user study which is described in the next section.

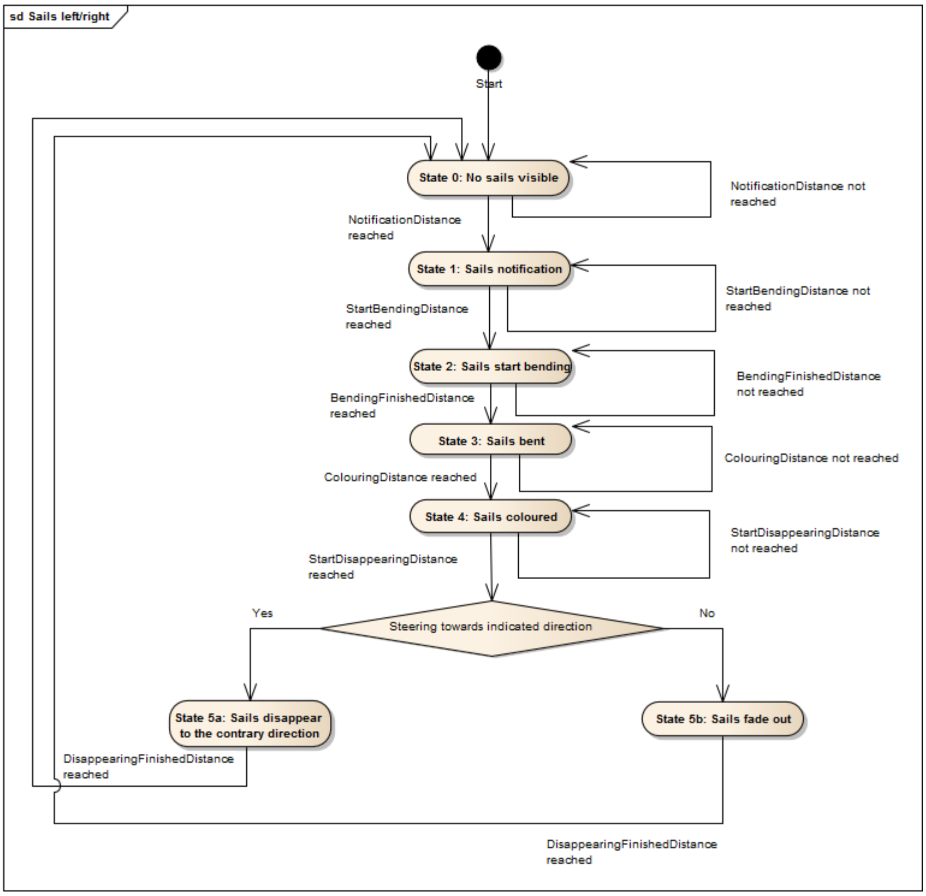

Figure 8.4.: Transition between the different states of the Sails visualisation.

8.4. Comparative User Study

In the user study we compared the novel Sails visualisation to a traditional three-dimensional arrow AR visualisation. The study was conducted in actual traffic situations.

8.4.1. Setup of the User Study

We asked the participants to drive a predefined route three times. The route consisted of six right and one left turn. Furthermore it included a roundabout and two successive turns to investigate the behaviour of the visualisations in these special cases.

(a) A 3D arrow located at the intersection pointing to the left.

(b) Sails representation, where the sails are pointing to the right.

Figure 8.5.: The two visualisations shown in the Head-Up-Display.

8.4.2. Participants

The user study was conducted with 16 test subjects, 10 of them being male and 6 of them being female. The age stretched from 23 to 45 with an average age of 32.4 years (standard deviation of 6.0 years). All of the participants had driven a car with a navigation system before, while 8 of the participants had driven a car equipped with a HUD before.

8.4.3. Test Procedure

The route took about 10 minutes driving time. The setup was a within-subjects design. In the first test drive the visualisations were disabled to give users the opportunity to familiarise themselves with the car and the route. This would allow them to focus in the remaining test drives on the visualisations. The users then drove the route a second time with one of the visualisations shown in the HUD and the last time they drove the route with the remaining visualisation. The order of the visualisations was equally distributed amongst the test subjects to counteract effects due to the order of the visualisations. After each test drive a short break took place, in which the participants filled out the questionnaires. The questionnaires are attached in the Appendix A and are described in the section of Dependent Variables.

After all three test drives, we conducted a qualitative interview to record feedback by the test subjects on how the two visualisations were perceived and how they could be improved.

8.4.4. Independent variables

The independent variable in this user study was the HUD visualisation. All participants drove without any visualisation, as well as with the ”Arrow” and the ”Sails” visualisation.

8.4.5. Dependent variables

We used the NASA Task Load Index (NASA-TLX) [HS88], the System Usability Scale (SUS) [Bro96] and the AttracDiff [Has04] in our study as described in chapter 4.4. The NASA-TLX gives an indication of the perceived workload. The NASA-TLX was filled out after all three test drives. The SUS gives an initial impression of the perceived usability of the visualisation.

Table 8.1.: Overview of the SUS and TLX scores

Rating	SUS		TLX	
	Mean	Std.dev	Mean	Std.dev
None	-	-	16.51	12.93
Arrow	76.875	17.26	17.15	8.78
Sails	77.03	20.05	19.49	11.77

The AttracDiff is a popular measure in industrial design, as it provides insights about the hedonic and pragmatic quality of an application. To develop an attractive product both dimensions should be maximised. The SUS and the AttracDiff were only filled out by the participants after a test drive with a visualisation, as they can only be judged on an application.

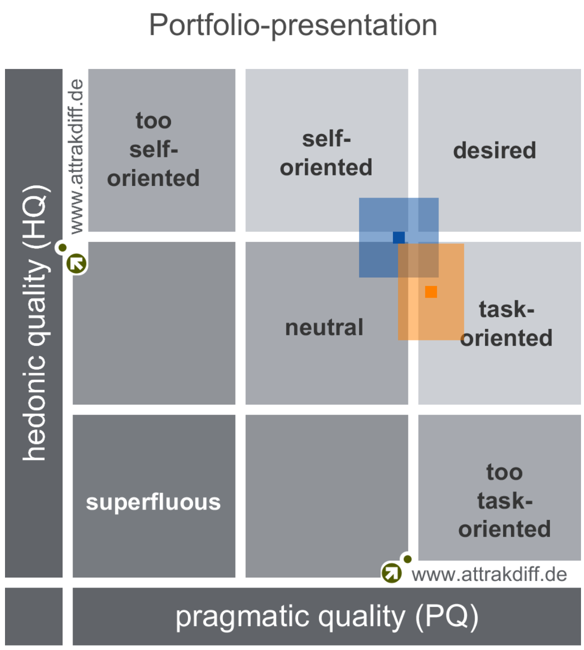

Figure 8.6.: The results of the AttracDiff. Orange depicts the "Arrow" visualisation and blue the "Sails" visualisation.

8.5. Results and Discussion

The results for the SUS and the TLX scores can be seen in Table 8.1. As the SUS scores were not normally distributed, we performed a paired signed-rank Wilcoxon test and did not find significant differences ($p\text{-}value = 1$) between the Sails and the Arrow visualisations. As the TLX scores were not normally distributed, we performed a Friedman's test and did not find significant differences ($\chi^2 = 2$, $p\text{-}value = 0.3679$) between the two visualisations. These results show that the new Sails visualisation did not significantly increase the workload on the participants and is similar to the traditional visualisation in terms of usability.

Figure 8.6 shows the mean values for the AttracDiff where the rectangles depict the confidence intervals for the visualisations. Overall, the Sails representation scored a higher hedonic quality, whereas the arrow scored a higher pragmatic quality. The results for each individual dimension can be seen in Figure 8.7.

The arrow's higher score in the pragmatic quality can be interpreted as that the visualisation was perceived as an easier way of route guidance. In both hedonic dimensions the Sails scored higher which means users identified themselves more with this visualisation and it was more stimulating to use. Overall, the attractiveness, i.e. how much people desire a product, was higher of the Sails representation.

We observed an overall preference for the Sails representation. Results of the subjective preference can be seen in Table 8.2. In the interview three subjects mentioned that the Sails can be distracting due to their relatively complex animation. However, seven of the participants mentioned that the arrow representation was not always visible, this is especially the case while stopping at an intersection. This is due to the limited field of view (FoV) of the HUD. Inspite of using a high-precision GNSS system, four participants nevertheless noted that the arrow is not steadily attached to the road.

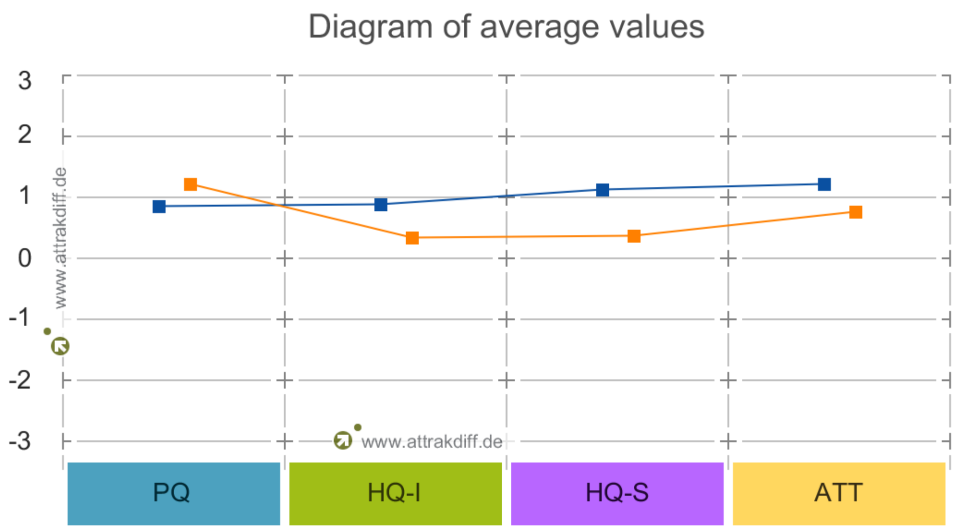

Figure 8.7.: The results for each dimension of the AttracDiff. Orange depicts the "Arrow" visualisation and blue the "Sails" visualisation. (PQ=pragmatic quality, HQ=hedonic quality, I=identification, S=stimulation, ATT=attractiveness)

Overall, the Sails representation was perceived as more capturing and the arrow visualisation as more simplistic. This interpretation correlates with the outcome of the AttracDiff. The Sails representation was therefore the preferred visualisation, without adding extra workload or being less usable as indicated by the TLX and SUS scores.

Table 8.2.: Preference of the visualisations

Preference	No. of people	Percentage
Sails	9	56.25%
Rather Sails	1	6.25%
Rather Arrow	0	0%
Arrow	5	31.25%
None	1	6.25%

8.6. Summary

In this chapter, we described the development of a navigational AR visualisation which overcomes the global positioning error of current navigation systems. Our visualisation is only dependent on the distance to the next turn and gives users the impression the visualisation is taking place in their environment. The visualisation is registered to the vehicle's coordinate system and therefore does not need to compensate for fast movements due to yawing or pitching of the vehicle.

We implemented the proposed AR visualisation and the traditional three-dimensional arrow in our car prototype. We then compared these two visualisations in a user study. We measured the usability, the experienced workload as well as the pragmatic and the hedonic quality. There were no significant differences in the usability and the experienced workload. The arrow visualisation scored in a higher pragmatic quality whereas the Sails visualisation achieved a higher hedonic quality as well as an overall higher attractiveness.

For the user study we chose a high-precision GNSS system to concentrate on the perception of the visualisations. Even though the three-dimensional arrow's location was much more precise than a conventional GNSS system, the Sails visualisation was still the preferred visualisation. Hence this visualisation is one necessary step closer in integrating AR visualisations into series-production vehicles.

The main contribution of this work is a visualisation concept which is able to cope with errors in global localisation. As we expect positional errors to be inevitable in the coming years, this visualisation is a visualisation which functions steadily and is not too sensitive to the errors. Therefore, it is a visualisation which could be easily implemented in current vehicles and can make car navigation easier.

Part III.

Visualisation of the future road course

This part describes the generation of a visualisation of the future road course in the HUD. Because of the limited Field of View (FoV) of the HUD a visualisation scheme was chosen which does not fully register the visualisation in the real environment, but is still registered in three dimensions, resulting in a highly dynamic visualisation. We call this type of visualisation an *AR-like* visualisation. The visualisation is based on the electronic horizon, which is already present in modern cars, therefore does not require any additional data. After successful implementation of the visualisation in a car prototype, the effects of the visualisation on the driving performance were assessed in a user study. The visualisation led to an improvement of drivers' braking behaviour and was well received by participants.

9. Depiction of the Future Road Course in the Head-Up-Display

In addition to supporting the driver in the navigational task, another opportunity of the AR HUD is to enable the driver in foreseeing the future road course. This is especially important in situations with bad visibility conditions due to weather conditions or the road geometry, i.e. sharp bends of the road. It would be advantageous in such situations, for the driver to know the future road course, and the HUD has the potential of informing the driver about the future road without taking the gaze away from the windshield.

This chapter describes the development of a visualisation about the future road course in the HUD, which uses the electronic horizon already available in modern cars. It therefore, does not need any additional data. Because of the limited Field of View (FoV) of HUDs (as described in chapter 3), we decided to display the visualisation in a symbolic manner. After successful implementation in the car prototype, we conducted a user study investigating the change in the braking behaviour by participants. The implementation of the visualisation was conducted in the bachelor's thesis by Mike Ruf [Ruf16]. The concept of the visualisation, as well as the user study were published in [WRSK16] and [WRSK17].

9.1. Introduction

One of the advantages of the HUD technology is that drivers are not required to take their gaze off the windshield in order to perceive important information and can therefore continue to observe the traffic situation behind the windshield. Especially in situations of poor visibility, it is important for the drivers to keep their gaze towards the road. By visualising the future road course in the HUD, the safety and comfort of the driver can be increased. Such a visualisation is already possible by the electronic horizon available in current premium series cars.

The electronic horizon was mainly developed for Advanced Driver Assistance Systems (ADAS) and was standardised by the ADASIS forum. [ADA] The standard describes the administration and the distribution of the electronic horizon. The forum consists of several car manufacturers, car suppliers, map providers, ADAS suppliers and car navigation manufacturers. The visualisation proposed in this chapter uses the information of the electronic horizon to visualise the future road course in the HUD. By following automotive standards the visualisation shows the feasibility of implementing such a visualisation in a series production vehicle.

Current HUDs feature only static representations of navigational information. These representations depict distances usually as numerical values. There is the hypothesis that by replacing these numerical values with spatial illustrations, the cognitive effort of the drivers can be reduced to comprehend the shown information.

We therefore propose an *AR-like* presentation of the future road course in the HUD. Due to the limited FoV and since it is currently not yet possible to fully track the outside environment with sufficient precision the *AR-like* presentation does not embed the visualisation perfectly into the real scene. Our system shows a dynamically adapting, abstracted view of the future road course in a dedicated area of the HUD. The visualisation fulfils two of the three requirements of the definition by Azuma [Azu97]. Firstly, by displaying the future road course in the HUD, the visualisation is displaying virtual content in the real environment. Secondly, the visualisation is rendered continuously in real-time and reacts to changes of the car's position and heading. However, it does not fully comply with the third requirement of being perfectly registered with the real environment. Yet, it is closer to a full AR setting than the presentation of a navigational map on the head unit, the instrument cluster or the presentation of navigational hints in the HUD. We call our novel visualisation the three-dimensional future road course (3D-FRC).

3D-FRC makes the driver aware of the road geometry of the future road course, enabling him/her to adapt his/her speed in advance. As it is difficult to enforce bad visibility situations due to weather conditions we investigated the acceleration behaviour around sharp bends. The results of this investigation are presented in this chapter.

9.2. Related Work

The use of AR has been well studied in car simulators. Tönnis et al. describe in [TK06] a visualisation for directing the driver's attention towards a hazard outside the visible field of view of the driver. Smith et al. compared in [SSBG15] the HUD to a Head-Down-Display (HDD), where the HUD was associated with significantly faster task performance. Kim et al. described in [KWGP13] a comparative user study showing that AR HUD interfaces have the potential to increase the driver's safety. These studies show promising results for AR visualisations in the car. But it is important to assess the effect of these visualisations in a real car prototype. For this reason we implemented an *AR-like* visualisation to assess the effects of these visualisations already today, before such applications are even available.

More specifically visualisations of the future road course have been studied before. Kim et al. showed in [KD09] a visualisation concept, informing the driver about the future road course. The visualisation shows a 2D-map merging into the three-dimensional environment. The results of the user study showed significantly less errors in navigation. Charissis presented in [CPMA11] a visualisation of supporting the driver in bad visibility situations. The results showed an improvement in braking behaviour, as well as collision avoidance. While both visualisations showed potential of supporting the driver in driving safely, they would need a much larger Field of View (FoV) than available to present the information on the windshield.

Medenica et al. compared different navigational aids in [MKPP11] and AR showed the least negative impact on driving behaviour. This study was performed in a car simulator, indicating a high potential for AR visualisations in the vehicle. Doshi et al. presented in [DCT09] an investigation of a laser display projecting onto the windshield with some rudimentary form of eye tracking. This underlines the fact that visualisations need to be tested inside a car prototype in order to encompass the effects taking place in real driving situations.

Nestler et al. presented in [NDPR09] visualisation concepts for smart deceleration. These visualisation concepts were shown in the instrument cluster of the car. However, a gaze towards the instrument cluster requires the driver to take the gaze off the road. To overcome this shortcoming we implemented a similar visualisation in the HUD.

The data for our visualisation is provided by the electronic horizon. Ress et al. described in [REKB06] the ADAS Horizon, interchangeably used with the electronic horizon. Furthermore, Duschl presented in [Dus15] different visualisations of the electronic horizon in the instrument cluster. These visualisations were investigated in a car simulator. Our HUD visualisation also builds upon the electronic horizon, however showing the visualisation in the HUD instead of the instrument cluster. As pointed out earlier, it is expected to decrease the time the driver's gaze is not on the road.

There are other existing HUD visualisations informing the driver about the furture road course. The most common ones are symbolic navigational hints, usually using numbers to indicate the distance to a turn. Our visualisation intentionally omits the use of numbers and conveys to the the driver the future road course in a continously changing 3D course. HUDway [HUD] is also presenting the information in the windhsield, but along with lacking optical brilliance the visualisation does not use the electronic horizon. The electronic horizon provides much more information than just the road geometry, such as road condition and traffic jams.

In [KKKO13] Kawamata describes also the generation of the future road course in an HUD. Apart from not having one integrated HUD component the visualisation is, to our knowledge, not building upon the electronic horizon.

Müller and Dauenhauer suggested a taxonomy that allows to classify different ways of displaying information as connected to a physical anchor, i.e. a point or an object, in AR [MD16]. In their taxonomy, the depiction of the future road course is spectator coordinate system positioned and world coordinate system oriented, symbolically connected to its anchor in the physical world and is displayed with context. They make a clear recommendation to use context which we followed.

The visualisation of the 3D-FRC was presented as a concept in [WRSK16] and the evaluation of the user study was presented in [WRSK17].

9.3. Visualisation Generation

This section describes the steps taken to generate our 3D-FRC visualisation. We begin with the conception of the visualisation and continue with the description of the electronic horizon which is the main data source for the 3D-FRC visualisation. After that a brief system overview is given and the software architecture of the application is described. We continue with a description of the details of the development of the visualisation. Consequently we describe the waypoint extraction and describe how the information from the electronic horizon is fused with a high-precision GNSS provider. Finally the results in the vehicle prototype are presented.

9.3.1. Conception of the Visualisation

We aimed to reduce the time the driver is not gazing towards the road by visualising the future road course in the HUD instead of the head unit. The visualisation by Kim and

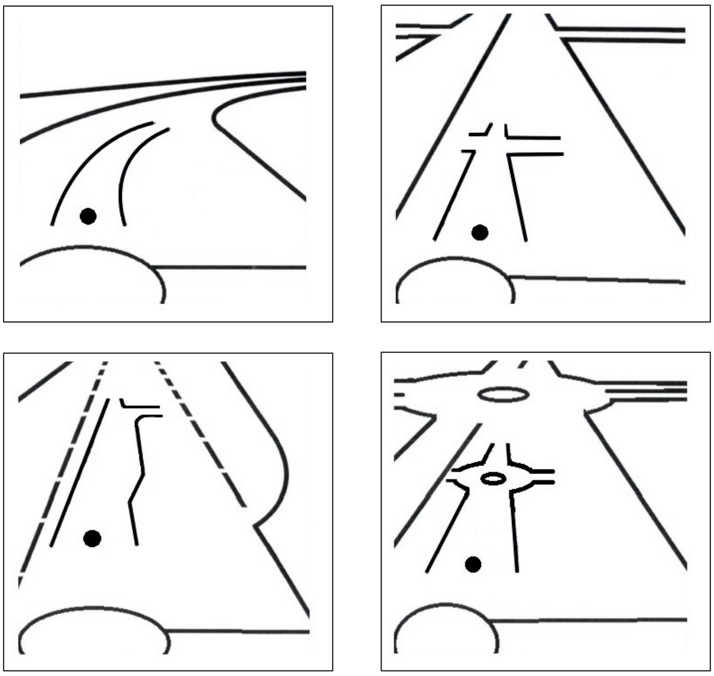

Figure 9.1.: Proposed visualisations of the future road course at special cases for intersections, highway exits and roundabouts.

Dey [KD09] aimed at reducing the distraction-related measures of the driver. To transfer such a visualisation from a simulator to a car remains challenging for several reasons. 1) A sophisticated recognition of the terrain is required in order to correctly align the road with the visualisation. 2) Even with perfect registration of the environment, a larger field of view (FoV) is needed for displaying the visualisation.

Before we started with an implementation of the visualisation we thought of special cases which may arise while using the application. The most common cases are bending curves and intersections. More special cases however, include roundabouts and highway exits. For these cases we created conceptions (see Figure 9.1).

While designing the visualisation we took precautions to not occlude real objects, as suggested by Gabbard et al. in [GFK14]. For this reason we designed the 3D-FRC visualisation as a minimal line, occluding real objects or hazards in the environment in the least possible way.

An application similar to ours is the HUDway app [HUD]. Besides lacking optical brilliance the app has other shortcomings, as it is difficult to locate oneself within the visualisation. At the time of the investigation, the visualisation did not display the paths leaving the main route. In our visualisation we show the leaving paths as "stubs" off the main path and the visualisation follows the vehicle's heading. The vehicle's position is depicted as a circle, making it easier for the driver to localise his-/herself within the path.

We displayed the 3D-FRC visualisation in a bird's eye view manner, as depicted in Figure 9.2. It is similar to the visualisation suggested by Duschl in [Dus15], but taking place in a three-dimensional environment and dynamically adapting to the real environment. By changing the orientation of the 3D-FRC visualisation according to the vehicle's orientation the cognitive effort to comprehend the visualisation is kept to a minimum. The visualisation's frame of reference is the one of egomotion [TPK13] by following the car's movement from an exocentric view.

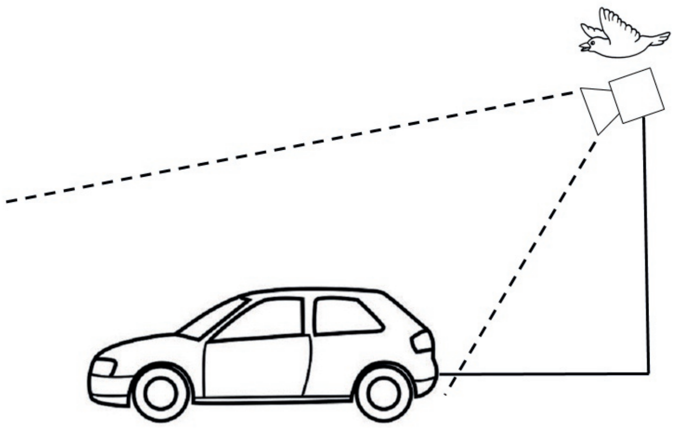

Figure 9.2.: Depiction of the bird's eye view.

9.3.2. Electronic Horizon

As the main data source for information about the future road course, we used the electronic horizon. The name is derived from the visual horizon as it describes what the vehicle "sees" with its sensors. Examples for these sensors are global positioning sensors, Radar, cameras and Lidar. Some information as static speed limits and the number of lanes might be known in advance, whereas other dynamic information such as the condition of the road or car accidents are only known when registered by local sensors.

Ress et al. [REKB06] define the electronic horizon as an "enhanced map data as well as vehicle route prediction for the road segments ahead of the vehicle". An exemplary overview of the electronic horizon can be seen in Figure 9.3. Current digital maps already include various information about the route. Such information is so far only processed in the navigation system of the vehicle. This information also becomes relevant for Advanced Driver Assistance Systems (ADAS), such as an intelligent form of an Adaptive Cruise Control (ACC). If the ACC would also know the information about a speed limit ahead of the vehicle, it could decrease the speed in advance. This information can be distributed amongst the ADAS by the concept of an electronic horizon.

The Advanced Driver Assistance Interface Systems Specification (ADASIS) aims to standardise the construction and the distribution of the electronic horizon inside the vehicle. ADASIS is hosted by ERTICO [ADA] and consists of several car manufacturers, car suppliers, map providers and navigation system manufacturers. By standardising the construction and exchange of the electronic horizon the modularity of the components increases.

Besides the already mentioned intelligent form of an ACC, other applications as increasing the efficiency of driving behaviour, become more and more important in times of increasing electrical cars. By knowing the future road course in advance, the system can adapt the

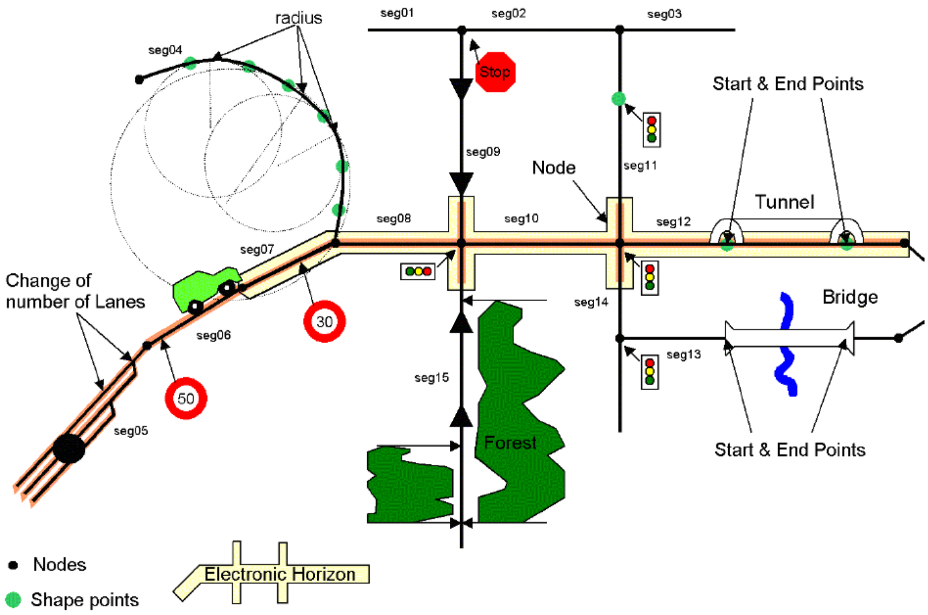

Figure 9.3.: Exemplary electronic horizon; Ress et al. [REKB06]

driving strategy early on to increase the efficiency.

In the ADASIS standard all the information from the map data as well as the vehicle's local sensors are aggregated in the ADAS Horizon. This ADAS Horizon is constructed and maintained by the ADAS Horizon Provider, the central unit controlling the ADAS Horizon in the vehicle. By having a central unit controlling the state of the electronic horizon, all applications are accessing the same state of the electronic horizon. In the ADASIS v2.0 standard the information of the ADAS Horizon is distributed via the Controller Area Network (CAN), a standard protocol in the automotive industry.

The electronic horizon provides information about the most likely path. This includes information about the road geometry of the road ahead, as well as turning angles of paths leaving the most likely path. It also provides the vehicle's position within the path.

From the CAN frames the ADAS Horizon can be rebuild by the ADAS Horizon Reconstructor. The ADAS Horizon Reconstructor then offers information about the electronic horizon via the ADASIS v2.0 API, from which all receiving ADAS can access this information. We chose to develop our application according to the ADASIS standard to be able to easily adapt our application for series production.

9.3.3. System Overview

As described earlier, the electronic horizon is maintained by the ADAS Horizon Provider. We integrated a prototype compliant to the ADASIS v2.0 protocol into our vehicle. The ADAS Horizon Provider runs as an app on an iPad. The ADAS Horizon Provider then sends the information about the electronic horizon to an Avisaro box which then translates the TCP messages received via WiFi to CAN messages. The Avisaro box can be programmed by a rudimentary programming interface. After this the ADASIS messages are available on the CAN bus. We confirmed the existence and accuracy of the CAN messages via the tool

CANoe by Vector.

The CAN messages are then received by the car PC, on which we render our 3D-FRC visualisation. Via the ADASIS Reconstructor the electronic horizon is rebuilt. We integrated the reference implementation of the ADAS Research Project (RP) into our visualisation software. The future waypoints as well as the information about intersections can then be accessed via the ADASIS v2.0 API.

After a first implementation and a test in the car prototype, we realised that the visualisation was not updated often enough. The reason was that the ADAS Horizon Provider updates the position only once per second, which is certainly not high enough to display the visualisation in a smooth manner. For this reason we used a high precision GNSS provider, which was already present in the car. This GNSS provider updates the position at a rate of 100 Hz which is sufficient to render the visualisation smoothly.

Sometimes the GNSS position is not as accurate as the map data due to environmental changes, such as changes in weather or buildings blocking the communication path to the satellites. For this reason we performed an orthogonal projection in order to match the GNSS position of the vehicle to the route of the map. The final result of our rendering is then projected onto the HUD which offers a larger field of view (FoV) and a higher virtual image distance (VID) than common HUDs. A complete overview of the system can be seen in Figure 9.4.

To generate the HMI visualisation on the car PC we used the CGI Studio by Fujitsu [Stu]. It is a development platform to develop automotive 2D/3D graphical interfaces. It is based on the Candera engine which renders the two- and three-dimensional scenes. The application widgets run on top of it and transform the data into a graphical representation. By developing the HMI in CGI Studio we enable a close to series development.

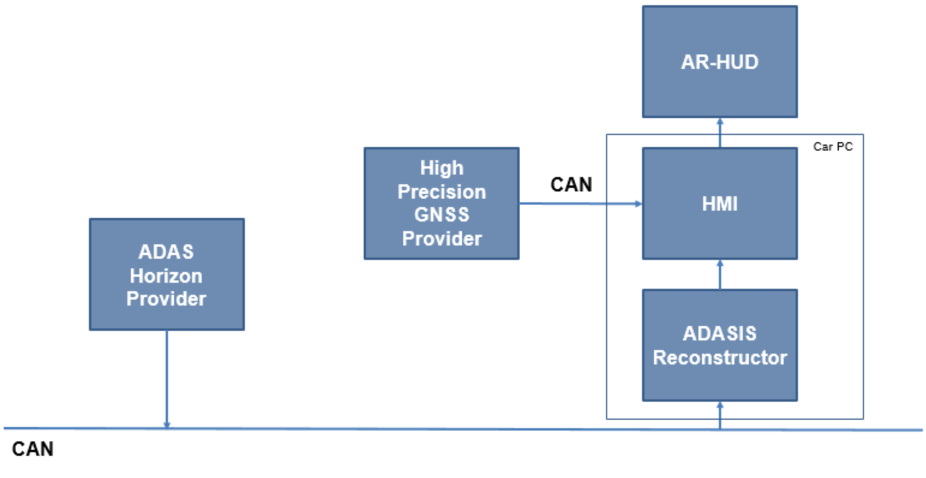

Figure 9.4.: System overview

9.3.4. Software Architecture

In the following we describe parts of the software architecture we developed the 3D-FRC visualisation according to. A class diagram of the visualisation can be seen in Figure 9.5.

The class *FutureRoadCourse* accesses the ADAS Horizon Reconstructor via the ADASIS

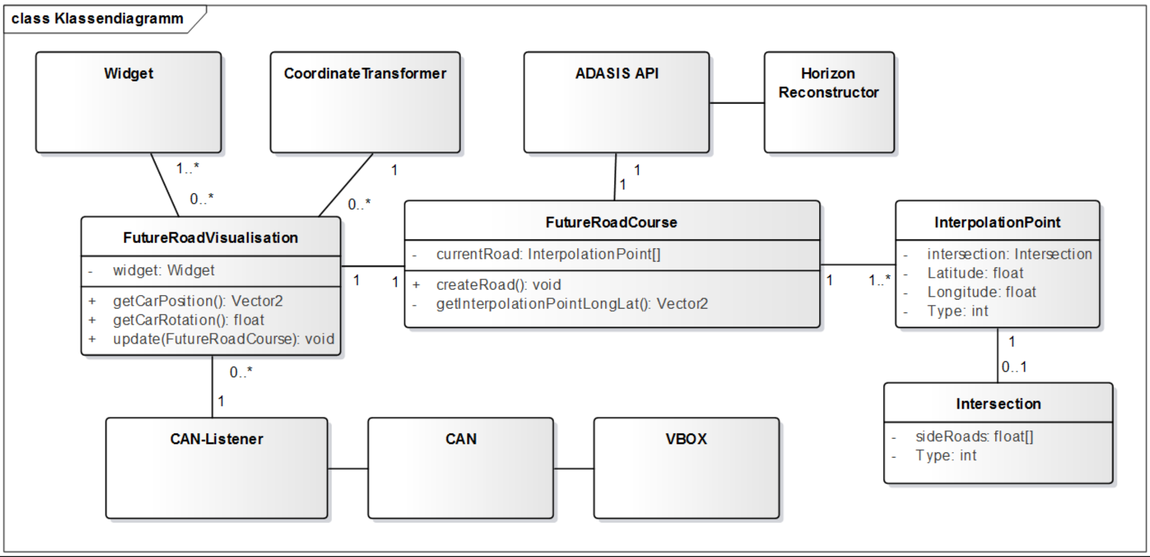

Figure 9.5.: Class diagram of the visualisation [Ruf16]

API. It consists of an array of *InterpolationPoints*. Each *InterpolationPoint* consists of the geographical coordinates, represented in latitude and longitude, and if it is an intersection, information about this *Intersection*. An *Intersection* consists of the angles of the side roads leaving the main path and the type of the intersection. This can be a roundabout, a highway exit or an ordinary intersection.

The *FutureRoadVisualisation* receives the car position and orientation from the GNSS provider *VBOX* via CAN. The interpolation points are transformed via the *Coordinate-Transformer*. The transformed interpolation points are then displayed with *Widgets* in the scene.

9.3.5. Development of the Visualisation

In the following we describe the algorithms we developed to extract the future road course from the ADAS Horizon, as well as the transformation we used to display the 3D-FRC visualisation.

Algorithm 1 describes the extraction of the future road course from the ADAS Horizon. Function calls to the ADASIS API are indicated by the ADASIS prefix. First the position of the vehicle within the path is determined. C gives in this case the offset (in metres) from the start of the current path. In a next step the next intersection is determined and its information stored in I. This information will be used later on. We then set the upper threshold as the current offset added by a variable parameter how much information about the future road course should be extracted in advance. In the ADASIS API the interpolation points are represented as profiles. A profile has a start point P_{start} and an end point P_{end}, as well as an offset and a value. Possible values depend on the type of profile and can be, e.g. longitude, latitude or a speed limit. Between the start and end point the values are interpolated. As the values are not all linearly interpolated, e.g. in the case of a speed limit, the type of interpolation is also supplied. Another characteristic of the P_{end} and P_{start} point

Algorithm 1: Extraction of the future road course from the ADAS Horizon [Ruf16]

Data: upper offset threshold max
Result: array of interpolation points Q

`// Get position of the car`
1 C $\longleftarrow$ ADASIS_GET-POSITION-OF-CAR()
`// Current offset is the offset of the car`
2 $offset \longleftarrow C_{offset}$
`// Get intersection next to car`
3 I $\longleftarrow$ ADASIS_GET-NEXT-INTERSECTION($offset$)
`// Set upper threshold`
4 $upperthreshold \longleftarrow offset + max$
`// Read out all the interpolation points with offset smaller than the upper threshold`
5 **while** $offset \leq upperthreshold$ **do**
    `// Get profile at offset`
6   P $\longleftarrow$ ADASIS_GET-PROFILE($offset$)
    `// Check if interpolation point is an intersection`
7   **if** $P_{end_{offset}} = I_{offset}$ **then**
        `// Push back new interpolation point with intersection`
8         $Q_{last+1} \longleftarrow$ ADD($P_{end_{offset}}, P_{end_{lon}}, P_{end_{lat}}$,I)
        `// Get next intersection`
9   I $\longleftarrow$ ADASIS_GET-NEXT-INTERSECTION(I_{offset})
10   **else**
        `// Push back new interpolation point without intersection`
11         $Q_{last+1} \longleftarrow$ ADD($P_{end_{offset}}, P_{end_{lon}}, P_{end_{lat}}$)
12     $offset \longleftarrow P_{end_{offset}}$

is that each end point is the next profiles start point. In the algorithm at line 7 it is checked if P_{end} is an intersection. In case it is the interpolation point with the information about the intersection is added at the end of Q by Q_{last+1}. Subsequently the next intersection is determined (if within limits of the upper threshold). In case the interpolation point is not an intersection the point is added without an intersection to Q. At the end the current offset is to the one of P_{end}. This information can then be used to visualise the future road course.

9.3.6. Waypoint Extraction

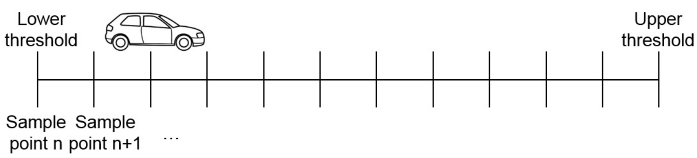

Figure 9.6.: Extracted sample points.

In this section we describe the extraction of the future waypoints and the construction of the 3D-FRC visualisation. The future waypoints are extracted from a lower threshold to an upper threshold ahead of the car. Figure 9.6 visualises this principle. The future waypoints are extracted for the next 2000 metres in our case. The length of the electronic horizon can be adapted as desired. For each waypoint we receive the latitude and the longitude, which then need to be projected into a Cartesian coordinate system in order to render them. There exist several possible projections, introducing errors in size or angle. We chose the equirectangular projection for this purpose, as the error in size is minimal when adjusting the standard parallel relative to the vehicle's position. The coordinates are calculated as follows:

$$\begin{aligned} x &= r\lambda cos(\varphi_0) \\ y &= r\varphi \end{aligned} \tag{9.1}$$

where λ denotes the longitude and φ the latitude. φ_0 is the standard parallel and to convert the angles to metres they are multiplied by the radius r of the earth. The further away the points are from the standard parallel the stronger the error introduced by the projection. We therefore averaged the latitude of the vehicle and the waypoint.

After extracting the waypoints and projecting them into a two-dimensional Cartesian coordinate system, the waypoints are connected with a line. As it is difficult to locate oneself in the environment, we also added the information of intersections to the visualisation. The electronic horizon provides us with the angle of paths leaving the main path. We visualise these roads as stubs in order to not clutter the visualisation, yet enabling the driver to locate his-/herself within the visualisation. We then rotate the visualisation according to the vehicle's orientation, so it is facing forward. If waypoints are behind the vehicle they are excluded from rendering and we stop the extraction of waypoints at this waypoint. As a final step the visualisation is tilted slightly upwards to provide the driver with a better overview of the situation.

The algorithm for transforming the visualisation is described in Algorithm 2.

In the first step, the standard parallel φ_0 is calculated by averaging the latitude of the vehicle and the latitude of the last interpolation point to minimise the error of the equirectangular projection. $VBOX_{lat}$ is the latitude as the information is received from the VBOX, the high-precision GNSS provider. Q is the result of Algorithm 1. The obvious advantage of the equirectangular projection is that the standard parallel needs to be calculated only once. The functions with the prefix CT are functions of the coordinate transformer in which we also implemented the equirectangular projection. In the next step the transformation matrix is calculated, where $VBOX_\alpha$ denotes the heading of the vehicle. The transformation matrix $M_{transformation}$ is then applied to all interpolation points in Q. To apply the transformation matrix the coordinates are transformed to homogeneous coordinates and back to 2D to apply the result to the widget. If the interpolation point is an intersection, the angles of the paths are read out and are applied to the widget.

9.3.7. Fusion with high-precision GNSS Provider

As mentioned before the position of the high-precision GNSS provider was sometimes not as accurate as the map data. The main reason for this is that the current position could be distorted by nearby buildings reflecting and absorbing the satellite signal, whereas the map data is generated by iteratively measuring and filtering of the route. The consequence

Algorithm 2: Transformation of interpolation points for displaying within scene [Ruf16]

Data: array of interpolation points Q

`// Calculate` φ_0

$\varphi_0 \longleftarrow \frac{VBOX_{lat}+Q_{last_{lat}}}{2}$

`// Equirectangular projection of car position from VBOX`

$C \longleftarrow$ CT_EQUIRECTANGULAR-PROJECTION$(VBOX_{lat}, VBOX_{lon}, \varphi_0)$

`// Create transformation matrix`

$M_{translation} \longleftarrow$ CT_TRANSLATION-MATRIX(C_x, C_y)

$M_{rotation} \longleftarrow$ CT_ROTATION-MATRIX$(VBOX_\alpha)$

$M_{scale} \longleftarrow$ CT_SCALE-MATRIX$(scale)$

$M_{transformation} \longleftarrow M_{scale} \cdot M_{rotation} \cdot M_{translation}$

`// For each interpolation point in Q`

foreach $Q_i \in Q$ **do**

- `// Equirectangular projection of interpolation point`

  $P \longleftarrow$ CT_EQUIRECTANGULAR-PROJECTION$(Q_{i_{lat}}, Q_{i_{lon}}, \varphi_0)$

  `// Convert to homogeneous coordinates for transformation`

  $P \longleftarrow$ 3D$(P_x, P_y, 1.0)$

  `// Apply transformation matrix`

  $P \longleftarrow M_{transformation} \cdot P$

  `// Convert back`

  $P \longleftarrow$ 2D(P_x, P_y)

  `// Check if interpolation point is intersection`

  if Q_i *is intersection* **then**

  - `// Add information about intersection to widget`

    WIDGET_ADD-INTERSECTION(P, Q_i)

  `// Add transformed interpolation point to widget`

  WIDGET_ADD-INTERPOLATION-POINT(P)

`// Update widget in scene`

WIDGET_UPDATE()

of this is that the first element of the route is not parallel to the route which is exemplified in Figure 9.7. Therefore, it is necessary to map the initial position of the vehicle onto the route.

For this purpose we used the orthogonal projection. The vehicle's GNSS position is between two interpolation points. Let $\vec{a}$ denote the vector from P_{start} to P_{end} of the two interpolation points and $\vec{b}$ the vector from P_{start} to the vehicle's GNSS position V. Then the vector $\vec{a}_b$ can be calculated by:

$$\vec{a}_b = \frac{\vec{a} \cdot \vec{b}}{\vec{a} \cdot \vec{a}} \cdot \vec{a} \tag{9.2}$$

This enables the projection of the vehicle's current GNSS position onto the map data. In Algorithm 2 then the projected position is used instead of $VBOX_{lat}$ and $VBOX_{lon}$.

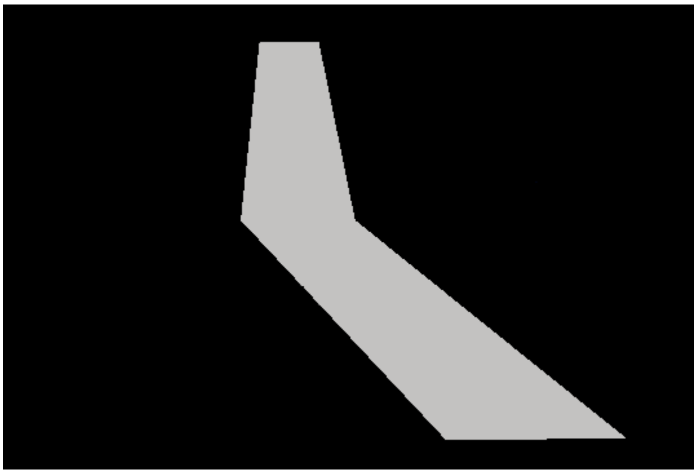

Figure 9.7.: Distortion of the vehicle's GNSS position leads to an inclined first element of the line segment [Ruf16]

9.3.8. Visualisation Results

The afore mentioned steps result in an HUD visualisation that runs smoothly and dynamically, showing the future road course. The visualisation taken during a live test in the vehicle prototype can be seen in Figure 9.8. As the visualisation was developed according to the ADASIS standard it can be easily adapted for series production. To be able to display this visualisation in conjunction with others, such as the Sails visualisation from [WK17], the 3D-FRC was put towards the back of the scene, to not occlude other visualisations, and scaled to be still visible. Screenshots of the result can be seen in Figure 9.9.

Figure 9.8.: Screenshots of our application during a live test in a prototype vehicle.

Figure 9.9.: Screenshots of our application in combination with Sails visualisation. [Ruf16]

9.4. User study

The user study investigated the effect of the visualisation on the drivers. As the visualisation's main intention was to inform the driver about the future road course, we investigated the effect of the visualisation in situations with bad visibility conditions. As the weather effects cannot be controlled in a real environment setup, we investigated the effect of the visualisation around sharp bends of the route. We expected the visualisation to have an effect on the acceleration behaviour of participants. Furthermore we investigated the subjective perception of the visualisation by the users. We also investigated the gaze behaviour introduced by the visualisation. As this was an extensive analysis and resulted in a follow-up study, this part of the user study is described in Chapter 10. The participants were driving two different routes, each about 20 minutes driving time and consisting of several sharp bends. The first time they drove without the HUD enabled and the second time they drove the route with the HUD enabled. The route was shown in both cases on the built-in car navigation system with route pointing in driving direction. The navigation was shown with a static zoom factor, as the 3D-FRC visualisation was also shown with a constant zoom factor. The voice guidance was disabled, as the main objective of the user study was to investigate the perception of the visualisation. The participants were not helped in following the route, only when they left the route, were they supported in getting back onto the route.

9.4.1. Setup of the User Study

As we wanted to investigate the effect of the visualisation in situations with bad visibility conditions, we selected two routes consisting of several sharp bends. The two routes can be seen in Figure 9.10. The participants drove the first time with the HUD visualisation disabled and the second time with the HUD visualisation enabled. To counteract characteristics of the routes the order of the routes was equally distributed amongst the participants. Each route took about 20 minutes driving time.

9.4.2. Participants

We conducted the user study with 20 participants, none of whom had used the system before. We employed 13 male and 7 female test subjects with an average age of 40.7 years

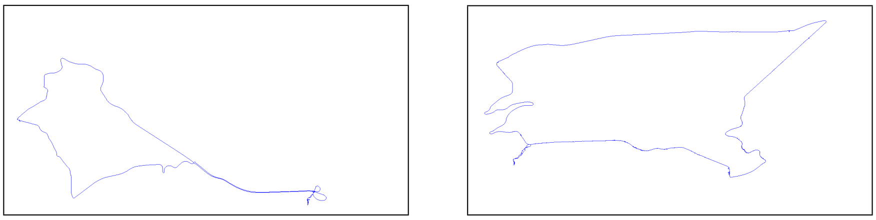

(a) First route of the measurement. (b) Second route of the measurement.

Figure 9.10.: The two routes used in the user study. Images were created by extracting the GNSS coordinates from the test drives.

(standard deviation 8.65). The age stretched from 25 to 58 years. Only one test subject drove less than 10,000 kilometres per year, 10 participants drove between 10,000 and 20,000 kilometres per year and the remaining 9 participants drove above 20,000 kilometres per year. Nine of the participants had driven a car equipped with a HUD before. The test subjects received a 25 Euro voucher for their participation.

9.4.3. Test Procedure

The test subjects drove the first test drive with the visualisation disabled to give them the opportunity to familiarise themselves with the vehicle. In the second test drive the HUD visualisation was enabled. After both test drives, we conducted a qualitative interview to record feedback by the test subjects.

9.4.4. Independent Variables

The independent variable in this within-subjects design was the visualisation. All participants drove the route without and with the visualisation enabled. The first test drive was driven with the visualisation disabled and the second test drive with the visualisation enabled. The order of the routes was equally distributed amongst the test subjects, to counteract for special characteristics of the routes.

9.4.5. Dependent Variables

To investigate the effect of the visualisation on the test subjects around sharp corners we recorded the acceleration of the vehicle throughout the test drives via a CAN trace by recording the vehicle's integrated components. The routes were entered into the built-in navigation system and the voice guidance was disabled in order to investigate the effect of the visualisations. In the second test drive the route was also entered into the electronic horizon provider to supply the HUD visualisation with the route. The test subjects had no information about the course of the route in advance.

9.4.6. Subjective Variables

In order to investigate the subjective perception of the visualisation the participants were asked the following questions after the two test drives:

- ”Does the HUD visualisation increase your feeling of safety in situations with bad visibility?”
- ”Do you find it easier with the HUD visualisation to find the right way?”

In the interview participants were also asked how they rated the maturity of the visualisation and how the visualisation could be improved.

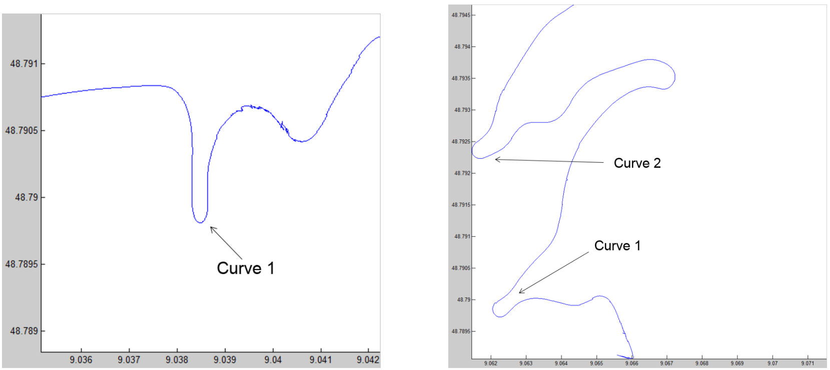

(a) The sharp bend in the first route.

(b) The two sharp bends in the second route.

Figure 9.11.: The sharp bends of the two routes.

9.4.7. Results

This section presents the results of the user study. The mean number of navigational errors decreased from 0.421 to 0.21 by the HUD visualisation. This shows a slight tendency of the users finding it easier finding the right way with the visualisation.

The acceleration behaviour was investigated around sharp bends. The sharp bends in the two routes can be seen in figure 9.11. Initially we selected other curves as well which had to be excluded, one because of a steep slope and another one because of a speed limit just before the curve. For each participant we determined the minimum of the longitudinal acceleration around these sharp bends. The result is presented in Figure 9.12.

As strong braking is not desired in public traffic, we determined for each test drive the maximum deceleration around the sharp curves. We then averaged for each participant the deceleration of the second route and compared this to the deceleration of the sharp curve in the first route. For each participant we then calculated the difference between the two routes, depending on which was driven without or with the HUD visualisation. The mean difference showed an improvement by 0.25 m/s^2 (standard deviation 0.80) with the HUD visualisation enabled. Naturally there can be many factors affecting the acceleration behaviour in public traffic, but an overall tendency is visible.

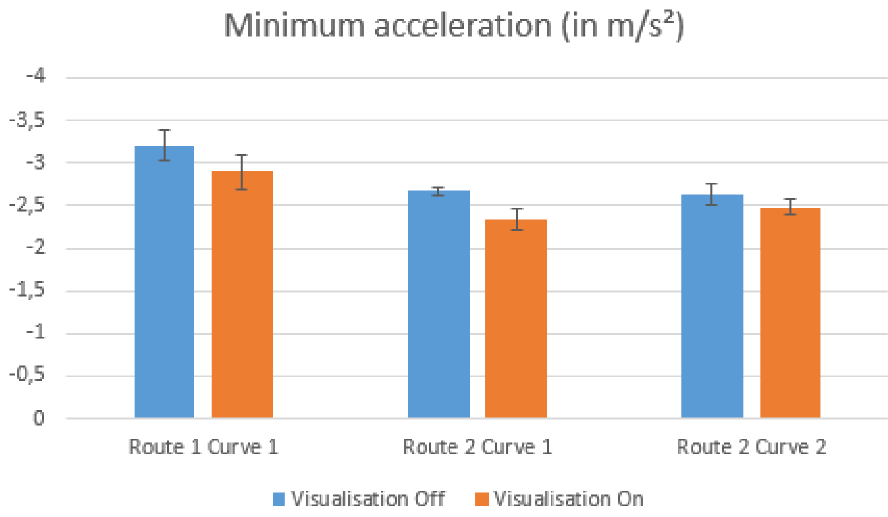

Figure 9.12.: Comparison of minimum acceleration at sharp bends of the road (given in m/s^2) between without and with visualisation.

Subjective results

The results of the subjective questions are presented in Figure 9.13. The visualisation was found helpful by 75% of the users during situations with bad visibility conditions. 80% of the users found it easier with the visualisation to find the right way. An overall tendency of users finding the visualisation helpful is visible.

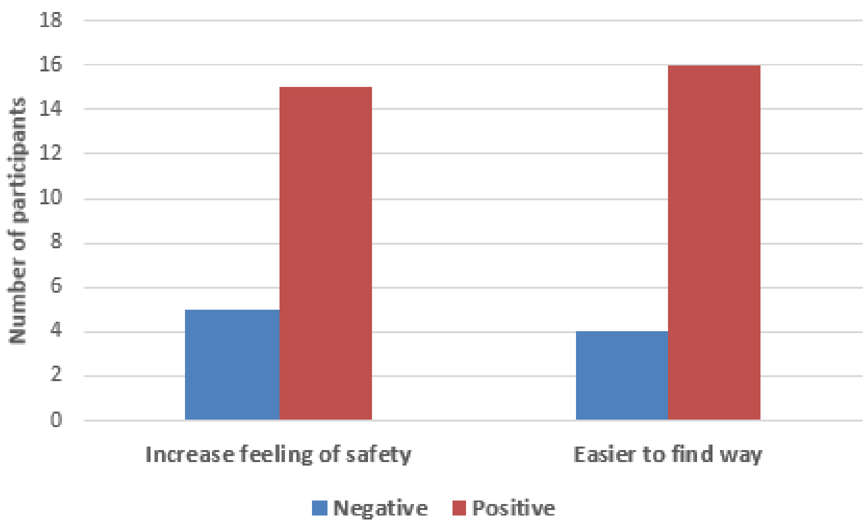

Figure 9.13.: Results of the subjective questions. Overall test subjects had an increased feeling of safety with the HUD visualisation found it easier to find the right way compared to the head unit.

9.5. Summary

This chapter presented the development of a visualisation of the future road course in the HUD. The visualisation is based on the electronic horizon which is already available in modern cars. The visualisation shows the future road course in a three-dimensional manner and was therefore named 3D-FRC (three-dimensional future road course). By complying to the automotive ADASIS standard our visualisation could be easily adapted for series production.

After an extensive description of the development of the visualisation, part of the user study to investigate the perception of the visualisation by users was presented. The visualisation was designed in an *AR-like* manner to be able to integrate the visualisation in a current prototype and observe the effect by such visualisations in real traffic situations, before full AR visualisations become available in the vehicle. The implementation in a real car prototype gives us the opportunity to investigate the effects, which are too complex to be analysed in a simulator. The results show an improvement in the braking behaviour and was well received by the participants. During the course of this user study we employed an eye tracking system, the details of which are presented in the next chapter and resulted in a follow-up study.

Part IV.

Change in gaze behaviour by the HUD visualisations

This part investigates the effect of the HUD visualisations on drivers. Emphasis is placed on the investigation of the gaze behaviour. For this purpose the drivers were equipped with eye tracking glasses to record what the driver was seeing and where he/she was gazing towards. The initial user study showed a significant decrease in gazes towards the instruments in the vehicle, but also showed an increase in the glance duration in the area of the HUD when the visualisation was enabled. To investigate this change in gaze behaviour further, we conducted another experiment to determine the effect by conventional, static representations and how much this effect is increased by dynamic visualisations in the HUD. In order to investigate the cognitive workload in the second experiment, we used another measure, the Detection Response Task (DRT), according to ISO 17488.

10. Initial assessment of the gaze behaviour introduced by HUD visualisations

The previous chapter described parts of the results of the user study conducted in [WRSK17]. It was conducted in a real car prototype under real traffic situations. This also implies that driving speeds cannot be totally attributed to the HUD visualisation, but also depend on confounding factors related to varying traffic conditions. For this reason we also incorporated eye tracking glasses in the evaluation to asses a change in gaze behaviour, which we hypothesized to be one of the strongest behavioural changes. We mainly focused on the reduction of gazes towards the instruments inside the vehicle, namely the head unit and the instrument cluster, by an HUD visualisation. The HUD visualisation in this case was the 3D-FRC visualisation presented in Chapter 9. Furthermore, we investigated if the HUD visualisation led to an increased visual attention in the area of the HUD.

10.1. Introduction

The investigation of the gaze behaviour was conducted in a real car prototype under real traffic conditions to record all the effects which are too complex to be modelled by a car simulator. The test subjects were for this purpose equipped with eye tracking glasses which record what the user was seeing and where the user was gazing towards. This was one of the only options to analyse the gaze behaviour in a meaningful way, as the HUD visualisation is only visible from the driver's point of view. We expected that the gazes towards the head unit and the instrument cluster (IC) to decrease. Additionally, we investigated if the 3D-FRC visualisation led to a new kind of cognitive capture, in which the focus of the drivers on the visualisation and their attentiveness to the ongoing traffic are affected.

10.2. Theory

As we sought to increase the driving safety through our HUD visualisation of the future road course in the HUD, instead of displaying it in the head unit, we investigated the following hypotheses:

H1: 3D-FRC reduces number of glances towards the head unit

H2: 3D-FRC decreases the total glance time towards the head unit

H3: 3D-FRC reduces number of glances towards the IC

H4: 3D-FRC decreases the total glance time towards the IC

H5: 3D-FRC does not increase the mean glance duration in the area of the HUD

We chose to investigate these hypotheses in order to assess the opportunity of decreasing the time for drivers taking their eyes off the road while not leading to a tunnel view.

10.3. Setup of the User Study

We employed several test subjects to investigate the effect of a HUD visualisation on drivers. In addition to the 3D-FRC visualisation, the current speed was also displayed in the HUD. The user study was performed using two different routes. To investigate the change in gaze behaviour participants were equipped with the Dikablis Professional eye tracking glasses by Ergoneers [Erg] while driving the two routes. The first route was driven with the HUD visualisation disabled and the remaining route with the HUD visualisation enabled. To counteract characteristics of the routes, the order of the routes was equally distributed amongst the test subjects. In both cases the route was also entered into the built-in car navigation system, so users could choose between the identification of the route from the HUD visualisation, the car navigation system or a combination of the two. The visualisation in the head unit was showing the route pointing in driving direction and a static zoom factor, as the HUD visualisation also had a constant zoom factor. Because the main objective of the study was to investigate the change in gaze behaviour by the visualisation, voice guidance was disabled. Hints about the route were not given, unless participants left the route, in which case they were supported in getting back on the route again.

As described in Chapter 9 the test subjects drove two routes with a common start and end point, enabling us to distribute the order of the routes amongst the test subjects. Participants drove one route with the HUD disabled and the other one with the HUD enabled. Each route took about 20 minutes driving time, resulting in approximately 800 minutes of recordings.

10.4. Participants

The participants were the same as in Chapter 9. We employed 20 test subjects, none of whom had used the system before. 13 of the test subjects were male and 7 of them were female with an average ago of 40.7 years (std. dev. 8.65 years). The age stretched from 25 to 58 years. One of the test subjects drove less than 10.000 kilometres per year, 10 of the participants drove between 10.000 and 20.000 kilometres per year and the remaining 9 participants drove more than 20.000 kilometres per year. Nine of the participants had driven a car equipped with a HUD before and the test subjects received a 25 Euro voucher for their participation.

10.5. Test Procedure

Prior to the test drives test subjects were equipped with the Dikablis Professional eye tracking glasses to record the field of view of drivers and their gaze direction. The eye tracking system has one RGB scene camera and two infrared cameras, observing the user's eyes to estimate the gaze direction. The system enables the definition of Areas of Interest

(AOIs) relative to optical markers registered by the scene camera. We defined the AOIs as presented in Figure 10.1 for each participant.

Before starting the test drive the eye tracking system needed to be calibrated. We used the image of the HUD for this purpose. The test subjects were asked to orient their head such that the centre of the scene camera is in the middle of the HUD image. We then asked participants to focus on each of the four corners while keeping their head steady. The eye tracking system was calibrated accordingly and the test subjects drove the first test drive with the HUD visualisation disabled while their gaze behaviour was being recorded.

The recording was stopped after the end of the first test drive and the second route was entered into the navigation system as well as into the ADAS Horizon provider for the 3D-FRC visualisation. After calibrating the eye tracking system once again by the same procedure, the second test drive with the HUD visualisation enabled, began.

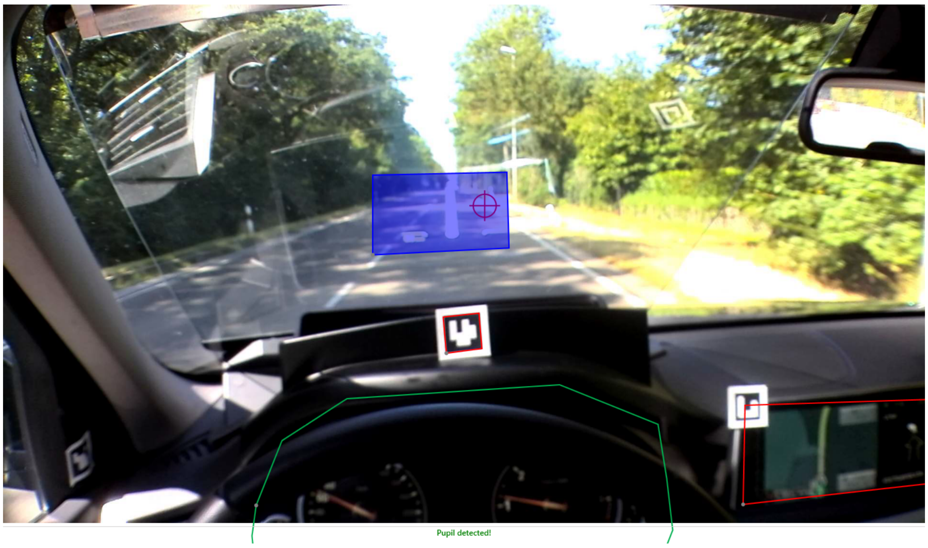

Figure 10.1.: Sample picture of the AOI definition with the eye tracking system. The blue area denotes the area of the HUD, the green area the instrument cluster and the red area the head unit.

10.6. Independent Variables

The independent variable in this within-subjects design was the HUD visualisation. All participants drove without the HUD visualisation and with the HUD visualisation enabled. The first test drive the test subjects drove with the HUD visualisation disabled and the second test drive with the HUD visualisation enabled. To counteract special characteristics of the routes, the order of the routes was equally distributed amongst the test subjects.

10.7. Dependent Variables

We recorded the gaze behaviour during the test drives. We aimed to investigate the effect of the HUD visualisation on the gaze behaviour. By defining the AOIs for the head unit,

Table 10.1.: Mean, standard deviation and median of the eye tracking statistics

Measure	Without visualisation			With visualisation		
	Mean	Std.dev.	Median	Mean	Std.dev.	Median
No. of glances head unit	65.15	39.13979	67	9.26316	9.60233	5
No. of glances IC	40.25	31.94218	37	17.42105	13.04827	20
Total glance time head unit (in seconds)	37.38115	30.46091	32.5915	4.28221	6.02999	2.131
Total glance time IC (in seconds)	19.1083	20.69118	13.578	8.40168	9.12759	6.148
Mean glance duration HUD (in seconds)	0.37945	0.08654	0.3635	0.42589	0.08592	0.401

IC and the area of the HUD we were able to assess the change in gaze behaviour.

10.8. Results and Discussion

We investigated the hypotheses, as formulated in the Theory section. The eye tracking recordings were analysed by using the D-LAB software. The software offers to recalibrate the tracking if the eye tracking glasses were moved during the measurement. We checked for all recordings at least in the middle of the recordings, if the tracking was off and recalibrated accordingly. The calibration was corrected at least at these positions, while with some test subjects even further recalibrations were performed. For each participants we defined, using the D-LAB software, AOIs for the area of the HUD, the IC and the head unit. The software offers an automatic calculation of the number of gazes, the mean glance duration as well as the overall glance duration towards the AOIs. The results of the eye tracking statistics can be seen in Table 10.1.

The measurements were checked for normal distribution by investigating their histograms. The data regarding the hypotheses (H1, H2, H3, H4, H5) was found not to be normally distributed. Therefore, we used the Wilcoxon signed-rank test to check for statistical significance.

We observed that the mean number of gazes towards the head unit (Figure 10.2) were strongly significantly reduced when the visualisation was enabled ($p\text{-}value = 0.0000714$). Additionally, the total glance time spent in the AOI of the head unit (Figure 10.3) was similarly highly significantly ($p\text{-}value = 0.000001907$) reduced when the visualisation was shown.

The effect regarding the IC were similar. The mean number of glances (Figure 10.4) and the total glance time (Figure 10.5) towards the IC were significantly reduced by the HUD visualisation ($p\text{-}value = 0.00111$, 0.001959 respectively). Therefore we can accept hypotheses H1-H4. These results show that the *AR-like* 3D-FRC visualisation has the potential that drivers spend more time observing the road as opposed to reading the information from the instruments inside the vehicle.

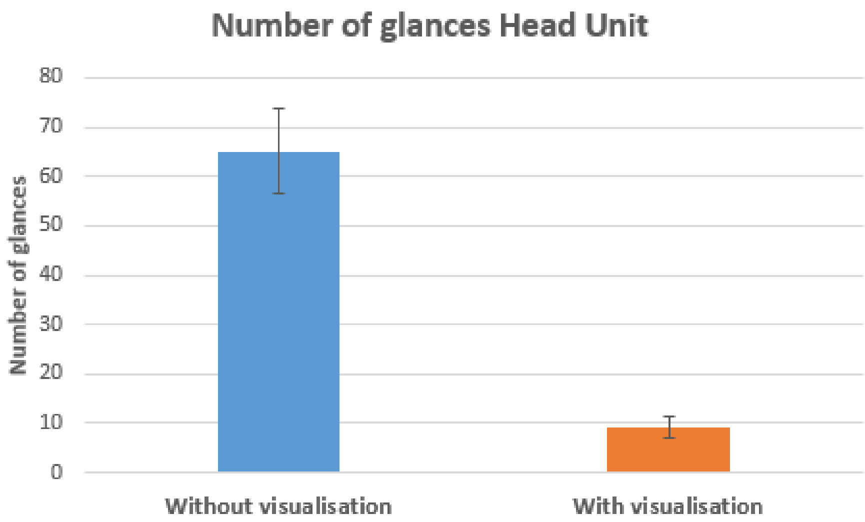

Figure 10.2.: Mean number of glances on the head unit with standard error of means. Left without visualisation (blue) and right with visualisation (orange).

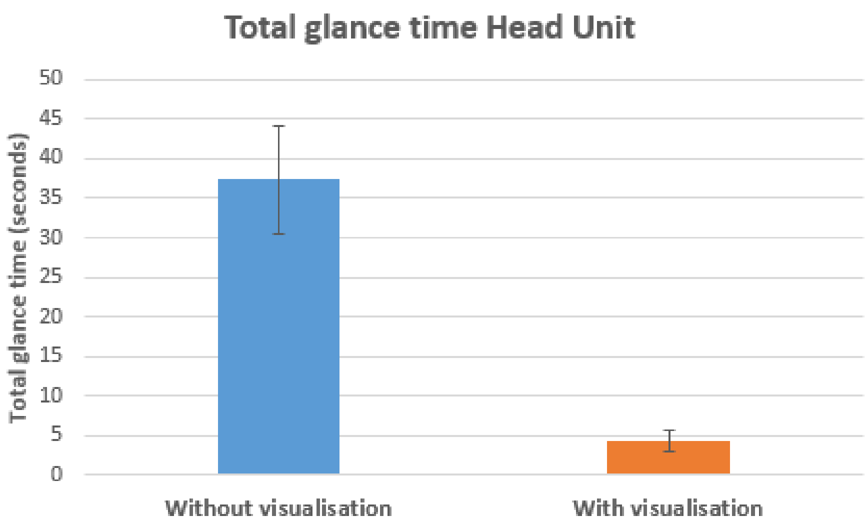

Figure 10.3.: Mean total glance time on the head unit with standard error of means. Left without visualisation (blue) and right with visualisation (orange).

On the other hand, analysis of hypothesis H5 showed that the mean glance duration in the area of the HUD was also significantly increased (p-$value$ = 0.004726) showing that the test subject spent a significant amount of time looking towards the HUD area. It has to be taken into account that the eye tracking system cannot distinguish between glances on the HUD visualisation itself and objects in the area of the HUD, which could bias the results. Further investigation of this point is necessary. All in all the HUD visualisation led to a higher amount of glance time in the area of the windshield. However, one needs to be

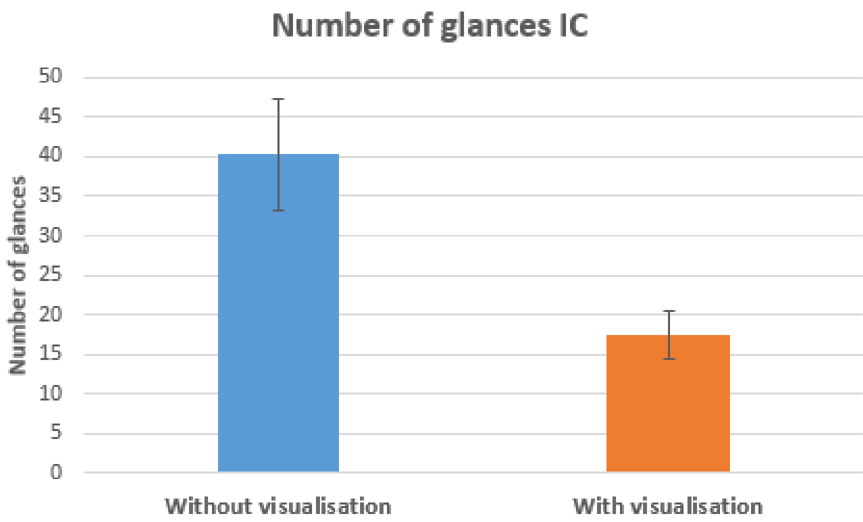

Figure 10.4.: Mean number of glances on the instrument cluster with standard error of means. Left without visualisation (blue) and right with visualisation (orange).

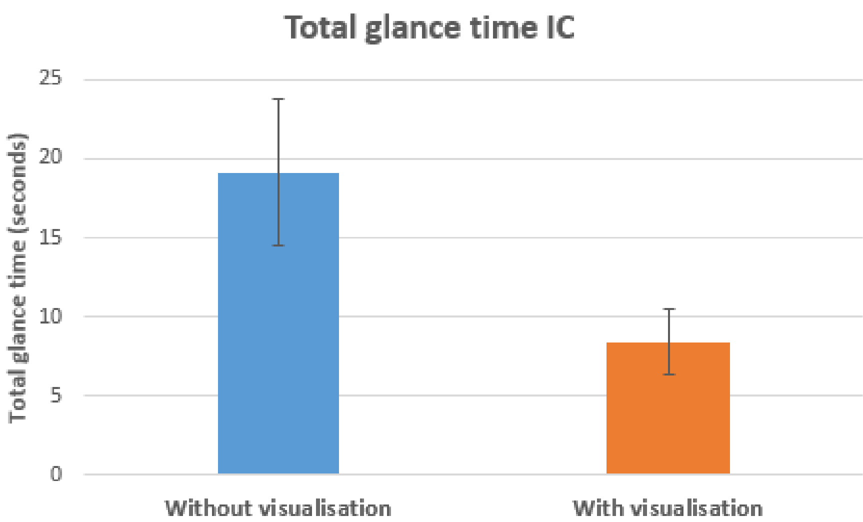

Figure 10.5.: Mean total glance time on the instrument cluster with standard error of means. Left without visualisation (blue) and right with visualisation (orange).

cautious to not create another distraction inside the vehicle by displaying excessively in the HUD.

10.9. Summary

This chapter described the initial assessment of the change in gaze behaviour by an HUD visualisation. We used the *AR-like* 3D-FRC visualisation for this purpose. To do so the

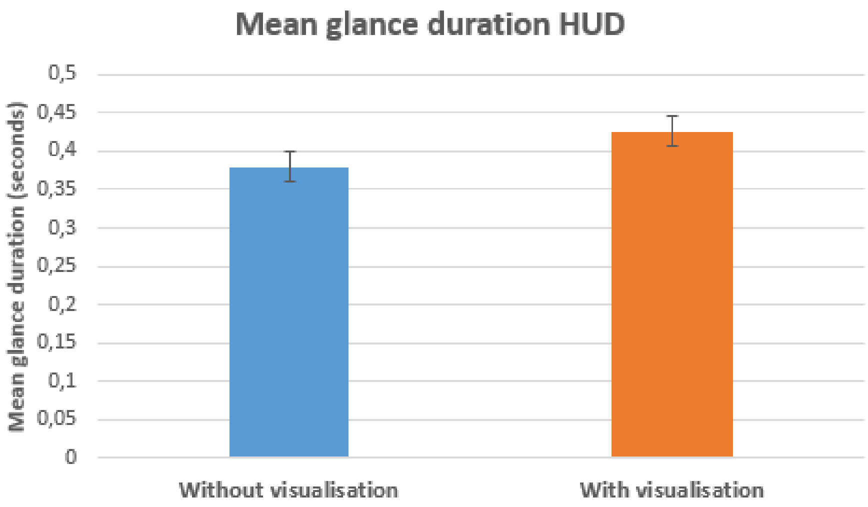

Figure 10.6.: Mean glance duration in the area of the HUD with standard error of means. Left without visualisation (blue) and right with visualisation (orange).

visualisation was implemented in a real car prototype and the study was executed under real traffic situations, as these situations are too complex to be simulated. The investigation of the gaze behaviour showed a significant decrease of the number of gazes as well as the total glance time towards the AOIs of the head unit and the IC. At the same time the mean glance duration inside the HUD area also increased significantly. This indicates the precaution that one needs to take of not displaying too much in the HUD, such that the HUD visualisation does not create another distraction inside the vehicle. It needs to be pointed out that the eye tracking system is not capable of distinguishing between glances towards the visualisation itself and if the user was glancing towards an object behind it. As shown by Cutting and Vishton [CV95] the convergence and accommodation of the eye change only below ten metres. Therefore, further investigation of the cognitive distraction was necessary, which we assessed in the next experiment, described in Chapter 11. In essence the experiment will investigate how much the gaze behaviour is already affected by the introduction of static HUD representations and how much further this change is affected by the introduction of dynamic *AR-like* representations. The initial results already show that *AR-like* visualisations have the potential of decreasing the time drivers spent not looking on the road and can therefore increase driving safety.

11. Refined investigation of distraction introduced by *AR-like* visualisations

In the previous chapter we showed that HUD visualisations have the potential to decrease the time drivers spend looking away from the road by reducing the time drivers spend looking towards the instruments inside the vehicle, such as the head unit and the instrument cluster. Furthermore, we showed that simultaneously the mean glance duration inside the area of the HUD increased significantly. We therefore wanted to investigate further how much the gaze behaviour is already affected by conventional HUD representations, such as the current speed, the speed limit and static navigational hints. And how the glance behaviour is further affected by dynamic *AR-like* representations, such as the Sails or 3D-FRC visualisations presented in [WK17] and [WRSK17]. Therefore, we set up a second experiment delving into the effect of *AR-like* visualisations on the drivers' cognitive load. As we already had insights into the gaze behaviour of drivers using the 3D-FRC visualisation in the previous study, we chose to investigate the effect of the Sails visualisation on the attention level of the drivers.

11.1. Introduction

The goal of this user study was to measure the effect of different HUD visualisations on the cognitive workload of users. In particular we were interested in the change in gaze behaviour. One finding of the previous study was that the glance duration in the area of the HUD increased significantly when the HUD was enabled. This increase seems only natural, as a visualisation is shown in an area which previously lacked one, hence drawing the attention of the drivers. Therefore, we wanted to determine how much the gaze behaviour is already affected by the introduction of an HUD itself and how much further it is affected by the introduction of *AR-like* visualisations in the HUD, as proposed in [WRSK17].

An additional factor to take into account is the inability of the eye tracking system to distinguish between the glances towards the HUD visualisation and the glances of the user while focusing on an object behind it. One possibility to distinguish between the two could be to observe the accommodation and the convergence of the eyes by observing the size of the pupils and the direction of the gaze vector of each eye with the other. Unfortunately accommodation and convergence only change below 10 metres according to [CV95] and the visualisation's virtual image distance (VID) is above this threshold.

We therefore employed another measure to investigate the cognitive effort required to understand the shown HUD visualisations. For this, we attached an LED to the frame of the eye tracking glasses, according to the standard ISO/DIS 17488. The standard describes the Detection Response Task (DRT), which is another measure to assess the cognitive load of the test subjects. The eye tracking glasses with the attached LED can be seen in Figure 11.1. During the test drives the experimenter was able to trigger the LED several times.

The time between switching the LED on by the experimenter and switching it off by the test subject was recorded by a Rasperry Pi controller, which was responsible for triggering the LED and recording the timestamps. The timestamps were then analysed in terms of the hit rate and the mean response time according to the standard ISO/DIS 17488. The DRT is explained further in section 11.5.

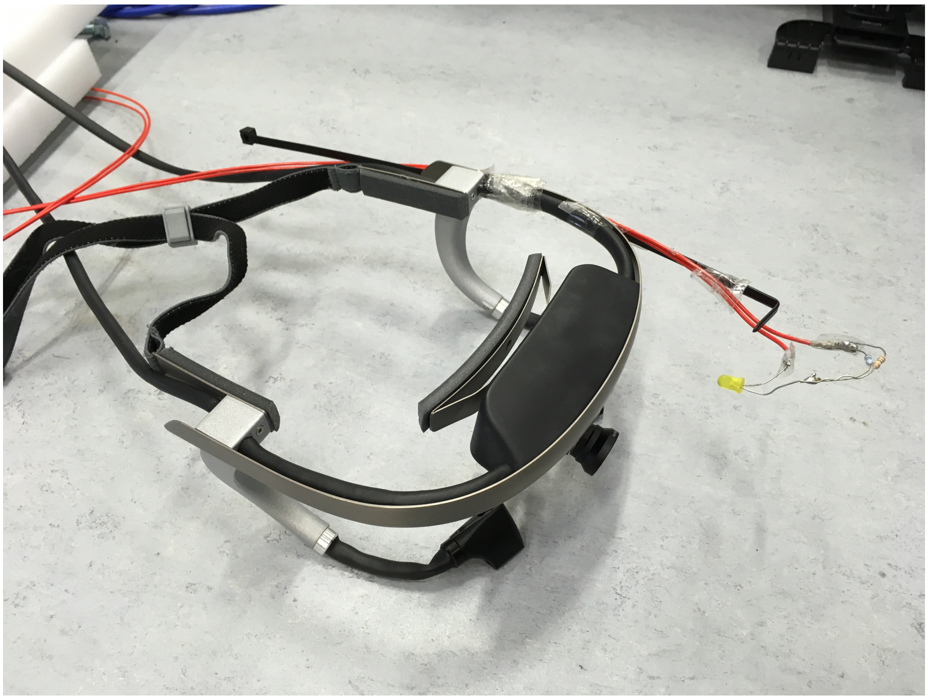

Figure 11.1.: The Dikablis Professional eye tracking glasses with the LED attached according to the standard ISO/DIS 17488.

11.2. Theory

We incorporated eye tracking as well as the DRT into our design of the user study. We therefore had numerous hypotheses concerning the eye tracking and the DRT. Concerning the eye tracking we investigated the following hypotheses regarding the Area of Interest (AOI) of the HUD:

ET1: The number of glances towards the AOI of the HUD increase when the HUD is shown

ET2: The number of glances towards the AOI of the HUD do not increase further when the Sails visualisation is shown

ET3: The total glance time towards the AOI of the HUD increases when the HUD is shown

ET4: The total glance time towards the AOI of the HUD does not increase further when the Sails visualisation is shown

ET5: The mean glance duration towards the AOI of the HUD increases when the HUD is shown

ET6: The mean glance duration towards the AOI of the HUD does not increase further when the Sails visualisation is shown

The hypotheses concerning the DRT were:

DRT1: The introduction of an HUD increases the mean response time of the DRT

DRT2: The introduction of an HUD decreases the hit rate of the DRT

DRT3: The introduction of dynamic representations does not increase the mean response time of the DRT

DRT4: The introduction of dynamic representations does not decrease the hit rate of the DRT

By analysing these hypotheses we aimed to investigate how much the attention of the driver is diverted by the introduction of a conventional HUD compared to *AR-like* representations in the HUD. We suspected that the introduction of an HUD alone increases the visual attention towards that area, but wanted to assess if and how much this is further increased by *AR-like* visualisations.

11.3. Setup of the User Study

The participants drove three different routes, each took about 10 minutes driving time. While driving the first route the HUD was disabled, serving as the baseline condition. After driving the first route the participants experienced either the conventional HUD visualisation, showing the current speed, the speed limit and static navigational hints, or additionally the Sails visualisation (which was presented in Chapter 8). The two visualisations can be seen in Figure 11.2. The third and last test drive was driven with the remaining visualisation.

Figure 11.2.: The conventional and the Sails HUD visualisation.

11.4. Participants

We conducted the user study with 15 participants, 12 of them being male and three being female. The average age was 36.9 years (std.dev. 9.2 years) with the youngest test subject being 22 and the oldest being 52 years old. Regarding the eye tracking statistics we had to exclude five test subjects from the analysis. During the test drives of two different test subjects the D-LAB eye tracking software crashed, resulting in incomplete measurements for these two test subjects. During the test drive of one test subject the lighting conditions became considerably worse, resulting in the lack of detection of the optical markers for the AOI definition. Furthermore, with two test subjects the eye tracking was not working reliably, such that even after extensive manual recalibration, the variance between the test drives was too high and the data of these two test subjects had to be excluded as well. During the DRT measurements the connection to the switch attached to the steering wheel broke during one test drive, so the results of this test subject were excluded in the analysis of the DRT.

11.5. Test Procedure

The test subjects were equipped with the Dikablis Professional eye tracking glasses with an LED attached to it. The LED in turn was controlled by a Raspberry Pi 1 Model B. To the Raspberry Pi, a switch was connected, through which it was possible for the experimenter to trigger the LED at different points in time. The test subjects were instructed to press another switch, which was attached to the steering wheel, when they recognised the LED being lit up. The Raspberry Pi then recorded the time difference between the trigger and the response. The test subjects were instructed to concentrate mainly on driving safely and only concentrate on responding to the trigger at a secondary priority as described in the ISO standard. Prior to the test drives the test subjects were familiarised with the procedure of lighting up the LED and responding to it while the vehicle was at standstill.

For analysing the gaze behaviour, we similarly to the setup in Chapter 10 defined the AOI for the area of the HUD, as this was the main AOI we aimed to investigate further. For calibrating the eye tracking system, we again used the HUD image described previously. The eye tracking system was recalibrated prior to each test drive. We recorded the gaze behaviour during all three test drives.

11.6. Independent Variables

The independent variable in this within-subjects design was the HUD visualisation. All participants experienced the conditions: without HUD, conventional HUD visualisation, HUD with *AR-like* Sails visualisation. The participants drove the first route with the HUD disabled, the second route with either the conventional HUD visualisation or the Sails visualisation and the third route with the remaining visualisation enabled. The order of the second visualisation was equally distributed amongst the test subjects.

Table 11.1.: Mean and standard deviation of the eye tracking statistics

Measure	No. of glances HUD		Total glance time HUD		Mean glance duration HUD	
	Mean	Std.dev.	Mean	Std.dev.	Mean	Std.dev.
Without visualisation	84.2	62.50653	45.0798	36.82396	0.4769	0.2651463
Conventional HUD	228.5	80.00174	139.2339	62.96284	0.5841	0.154244
Sails HUD	169.7	96.61844	89.9314	54.22316	0.4963	0.1492292

11.7. Dependent Variables

We assessed the gaze behaviour during all three test drives. As the main objective of this investigation was to analyse the gaze behaviour towards the HUD, an AOI was defined for the HUD. Furthermore, we recorded the trigger times of the DRT throughout the test drives by determining the time when the LED was lit up and when the participant responded by pressing the button attached to the steering wheel.

11.8. Evaluation of the user study

To evaluate the statistics of the eye tracking, we defined, using the D-LAB software, the AOI for the HUD in relation to four optical markers. Whenever at least one marker is visible, the AOI of the HUD is calculated by the software. If more than one marker is visible, the location of the AOI is interpolated between these markers resulting in a higher accuracy of the AOI location. After definition of the AOIs the software offers an automatic calculation over the measurement. The metrics were evaluated in respect to number of glances, the total glance time and the mean glance duration towards the AOI of the HUD.

For the trigger logs of the DRT, we calculated the hit rate according to ISO 17488. The hit rate is defined as the valid responses divided by the total number of responses. A response is considered valid when the response time is between 100 and 2500 ms. Of the valid responses, the mean response time was calculated by evaluating the arithmetic mean.

11.9. Results and Discussion

We analysed the results for the eye tracking as well as for the DRT.

11.9.1. Eye Tracking Results

Similarly to the previous study, we investigated the eye tracking results with the D-LAB software. The software enables recalibration of the eye tracking and automatic calculation of statistics towards AOIs. We defined an AOI for the area of the HUD, and evaluated it according to number of glances, total glance time as well as the mean glance duration. The overall statistic of means and standard deviation can be seen in Table 11.1.

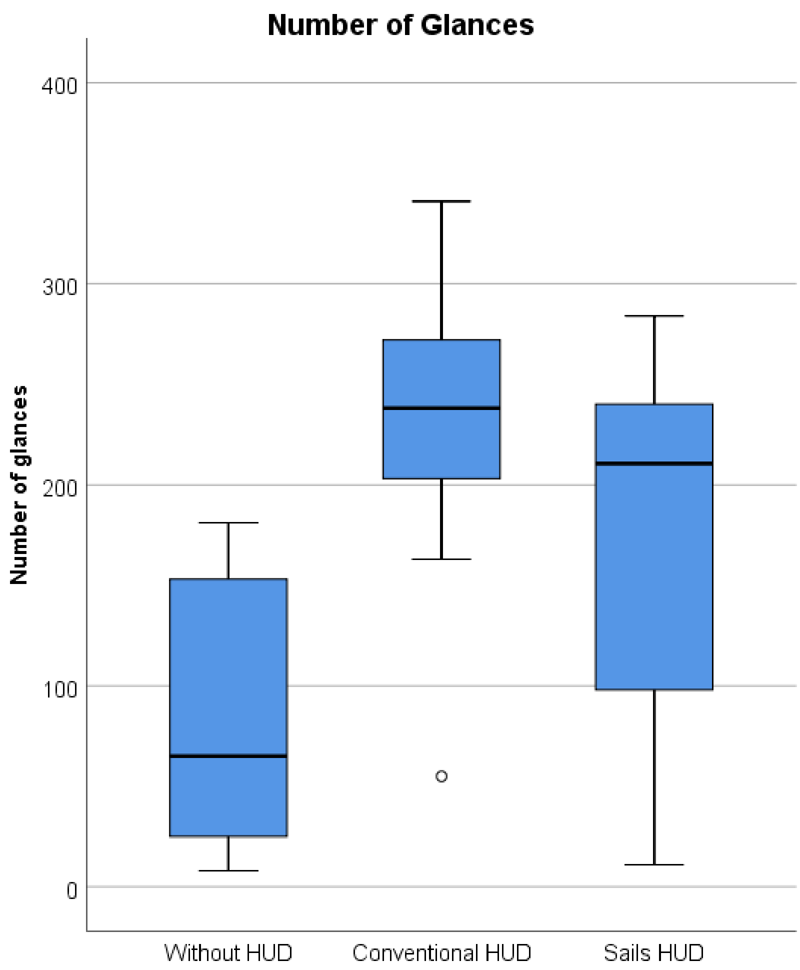

Figure 11.3.: Number of glances towards the HUD.

We investigated the data concerning our hypotheses (ET1, ET2, ET3, ET4, ET5, ET6) for normal distribution by investigating their histograms. The data was not found to be normally distributed. Therefore, we analysed the results with the Friedman's test. The Friedman's test found significant differences for all measures.

- Number of glances ($\chi^2 = 11.4$, p-value $= 0.003346$), see the box plot in Figure 11.3
- Total glance time ($\chi^2 = 9.8$, p-value $= 0.007447$), see the box plot in Figure 11.4
- Mean glance duration ($\chi^2 = 6.2$, p-value $= 0.04505$), see the box plot in Figure 11.5

To find the significant differences between the conditions, we performed a post hoc analysis with a Wilcoxon paired signed-rank test with a Bonferroni correction applied. The Bonferroni correction accounts for the multiple comparisons, and divides the significance

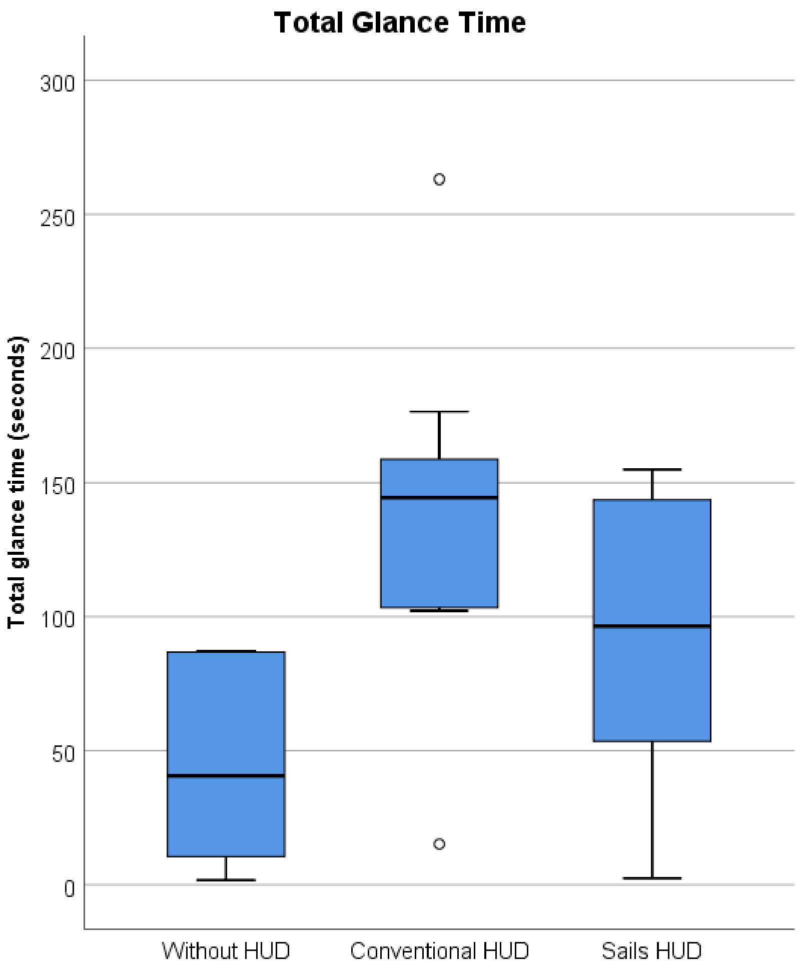

Figure 11.4.: Total glance time towards the HUD.

level by the number of combinations, which is in this case 3 different combinations. Hence the p-values are compared to the new significance value of $\alpha = 0.05/3 = 0.01667$. We found the following significant differences:

- **Number of glances:** Without and Conventional: V = 55, p-value = 0.0009766
- **Total glance time:** Without and Conventional: V = 55, p-value = 0.0009766

All other measures were not found to be significant. Therefore we can accept all our hypotheses ET1-ET6, apart from ET5. The difference in the mean glance duration was found not to be statistically significant (lowest p-value = 0.04199, Without and Conventional). One explanation for this outcome could be that the number of valid test subjects n=10 is too low to show a significant difference. Also, the Bonferroni correction is a conservative method to adjust the p-value for multiple statistical tests [Arm14] and might be too strict in this case.

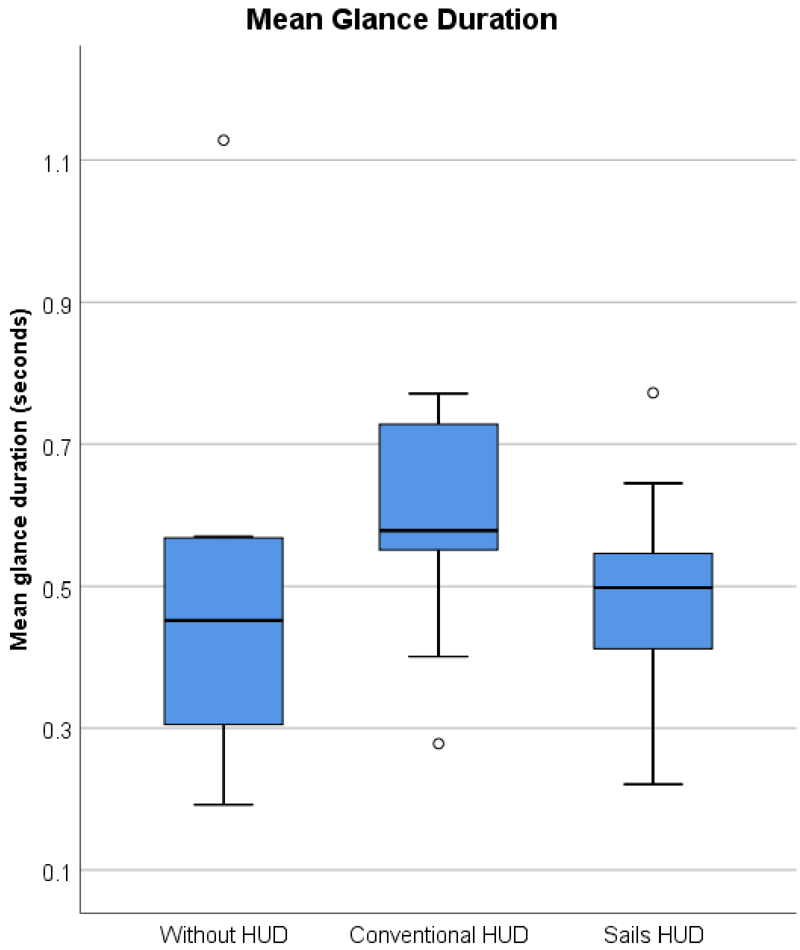

Figure 11.5.: Mean glance duration towards the HUD.

11.9.2. DRT Results

The box plots for the hit rate and mean response time for the DRT over the three conditions can be seen in Figures 11.6 and 11.7.

The mean hit rates for the three conditions are the following: $\bar{h}_{without} = 0.928$ (std.dev. 0.077) $\bar{h}_{conv} = 0.896$ (std.dev. 0.121) $\bar{h}_{sails} = 0.894$ (std.dev. 0.129) The mean response time for the three conditions are the following: $\bar{rt}_{without} = 0.811$ (std.dev. 0.256) $\bar{rt}_{conv} = 0.879$ (std.dev. 0.230) $\bar{rt}_{sails} = 0.889$ (std.dev. 0.276)

We analysed the results with the Friedman's test as the data was not normally distributed. The Friedman test did not find significant differences in either of the two measures. Therefore we need to reject hypotheses DRT1 and DRT2 and can accept hypotheses DRT3 and DRT4. Maybe more test subjects are required to show a significant difference for the DRT task.

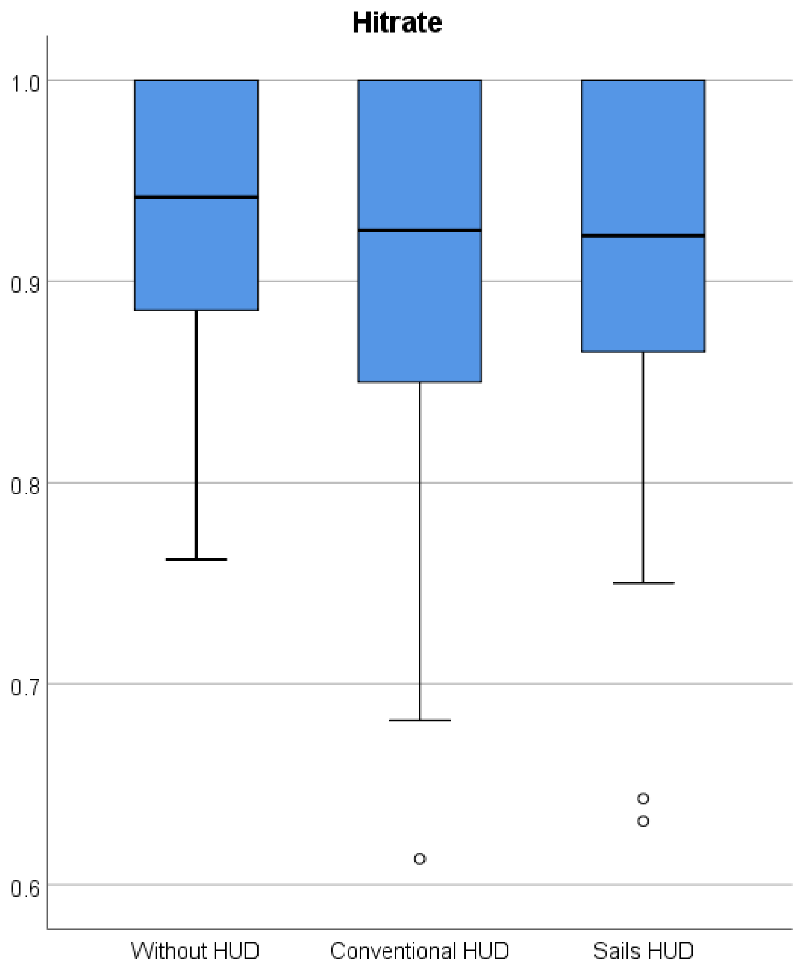

Figure 11.6.: The hit rate regarding the DRT.

11.10. Summary

We analysed the distraction effect of conventional visualisations compared to *AR-like* representations in the HUD. To investigate these effects we used eye tracking as well as the Detection Response Task (DRT) according to ISO 17488. While the DRT did not show significant differences between the test runs, the eye tracking revealed significant differences between test drives without an HUD and conventional HUD representations enabled. The analysis did not find any further increase with the Sails visualisation enabled. In fact, the investigated statistics (number of glances, total glance time and mean glance duration) show an overall lower mean in all three measures, but the differences being non-significant. It could be possible that more test subjects would show a significance in these measures. Compared to the previous study in [WRSK17], we investigated the Sails instead of the 3D-FRC visualisation. One explanation for the non-significant difference between no HUD and the Sails HUD could be that the Sails representation is visually less distracting than the 3D-FRC visualisation. Another explanation could also be that with 10 valid test subjects

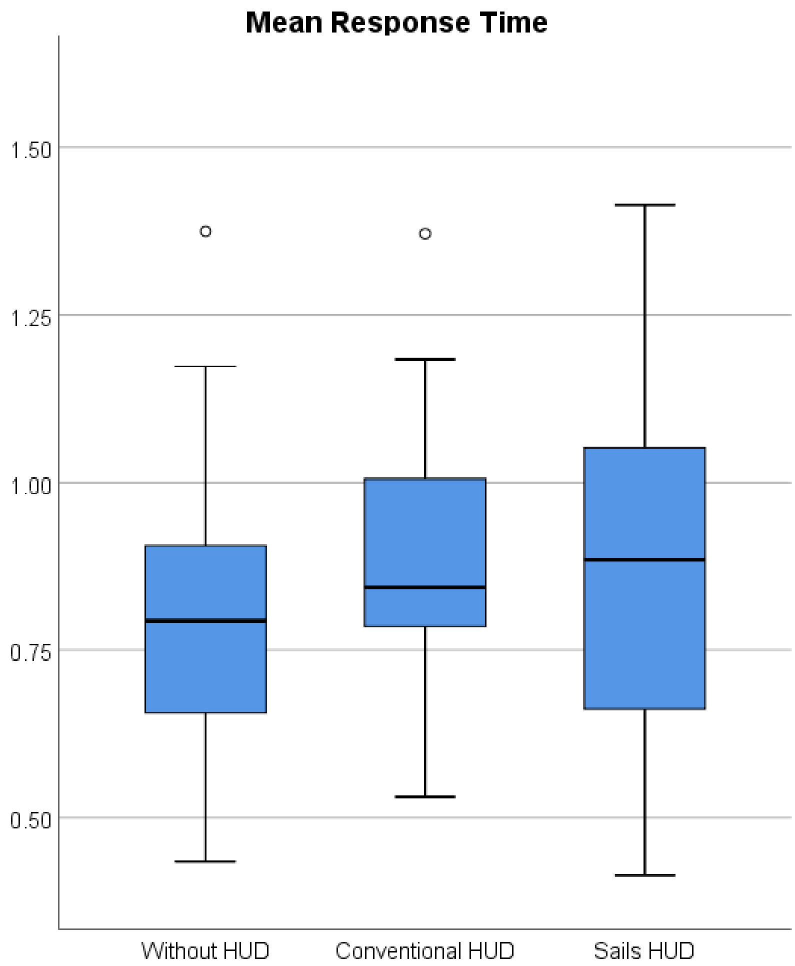

Figure 11.7.: The mean response time regarding the DRT.

the number of participants is too low to show a significant difference. What the study did show was that already a conventional HUD increases the visual attention towards that area significantly and *AR-like* visualisations do not increase this visual attention further, as can be seen from the non-significant lower means of the eye tracking statistics. Taken together with the results of the previous study, HUDs already decrease the gazes towards the instruments inside the vehicle, such as the head unit and the instrument cluster strongly significantly and can therefore potentially increase the time the driver is looking towards the road. Full-AR or our already implemented *AR-like* visualisations can decrease the visual attention towards the area of the HUD. These studies showed that by already implementing *AR-like* visualisations inside a real vehicle, make it possible to investigate the strong potential of AR visualisation in the vehicle.

Part V.

Assessment of the temporal aspect of HUD applications

Latency is a crucial aspect of AR visualisations and this part investigates the temporal aspect of the HUD visualisations. The first part describes the measurement of the end-to-end latency of our HUD prototype in order to get an initial impression of the current latency. After measuring the overall latency, we conducted a user study in order to determine which level of latency is actually perceivable. Through these steps we determined the overall end-to-end latency of the HUD system and to what threshold the latency is perceivable.

12. Measurement of the end-to-end latency of the Head-Up-Display

Latency has always been an important factor when developing AR applications. [ZWL+14] In the case of optical see-through AR (OST-AR) it becomes even more important, as it is obviously not possible to delay the real world which is a possibility in the case of video see-through AR (VST-AR). As the HUD is an optical see-through device, measuring and reducing the latency of the visualisations is a vital part during its development.

In this chapter we describe the measurement of the end-to-end latency of the current HUD prototype. The measurement was conducted in the master's thesis by Shreyas Dhone [Dho17]. The first measurement determines the latency induced by the HUD unit alone. As the overall system latency is larger due to the fact that also sensors, e.g. the camera, are involved in the complete system setup, we measured the end-to-end system latency. For the measurement of the overall end-to-end latency we adapted the design of Billeter et al. [BRW+16] to fit the needs of our application and developed a timestamp generator for this purpose, which is described in the second measurement. The actual latency experienced in the vehicle is different, as the sensors (e.g. camera) involved in the process of generating the visualisation, are already performing some sort of prediction which reduces the overall latency. Therefore, lastly we determined the latency of a particular application.

12.1. Introduction

While developing AR visualisations the latency plays a crucial part. It becomes even more important in the case of optical see-through AR (OST-AR), compared to VST-AR. As a first step we investigated what level of latency is actually generated in the HUD system. For this, we first describe the measurement of the HUD unit alone and then describe two different ways of determining the overall end-to-end latency. One way to determine the end-to-end latency inside the test vehicle was by having a timestamp generator visible in the real environment and showing the processed image within the HUD image. With an external camera observing the timestamp generator and the HUD image, it is possible to determine the end-to-end latency by subtracting the processed timestamp from the current timestamp which is visible on the timestamp generator.

12.2. Related Work

Billeter et al. presented in [BRW+16] a LED-based device for measuring the end-to-end latency. As they point out the use of prediction and extrapolation methods are necessary to decrease the lag between the real scenery and the augmented visualisations, especially in the case of OST-AR. To make use of these methods, accurate knowledge of the overall end-to-end latency is crucial. Therefore they developed a generic end-to-end latency measuring device for camera-based AR systems, with an accuracy of below 1 ms. As some functionalities of our AR HUD system are also camera-based, we adapted their setup to measure the end-to-end latency of our AR HUD.

Apart from determining the end-to-end latency of the system, it is necessary to determine the threshold where users are able to perceive latency. The study of Ng et al. [NLW$^+$12] shows, using touchscreen applications, that users are able to perceive latency well below 10 ms. In their study, they present the development of a high performance touch system which is capable of displaying latencies between 1 ms and several 100 ms with an accuracy of 1 ms. With this prototype they determined the perceivable latency of touchscreen applications through a user study. They determined the Just-Noticeable Difference (JND) with an average of around 6 ms, and a range from 2.38 ms to 11.36 ms (std.dev. 4.33 ms). However, the results of touchscreen applications cannot be directly applied to OST-AR applications, such as the AR HUD. Therefore, we analysed the perceptual effects of AR HUD applications in Chapter 13.

The work by Sielhorst et al. in [SSK$^+$07] analyses the measurement of the absolute latency of VST-AR systems. Their method, in essence, is to encode the current time into the image and decode it after camera feedback. They point out different sources of latency, which are: Exposure; read out and transfer to memory; tracking of objects; visualisation generation; and the latency of displaying the resulting image. Their method achieves a precision well below 1 ms. As they point out, it is of advantage to analyse the histograms of the recorded latency, which they consider more expressive than the average and the standard deviation. Itoh et al. [IOH$^+$16] describe the design of a low latency OST Head-Mounted Display (OST-HMD) which is achievable with off the shelf hardware. They achieve a mean temporal error of below 1 ms, with a median spatial error of below 0.3°in the viewing angle with a maximum error of 1.0°. The estimated time delay is below 1 ms and the average below 0.5 ms.

12.3. Measurement of the HUD alone

To measure the latency induced by the HUD unit, we used a lab-box of our HUD prototype. With the lab-box, it is possible to display the HUD visualisation on a small TFT display, which is the same in the actual HUD prototype. Also the physical interfaces, such as CAN and Ethernet, are the same as in the vehicle.

We needed a trigger to start and to stop the latency measurement of the HUD unit. Because of its low latency, we chose to perform the measurements using a phototransistor. The phototransistor used in our measurement was an OSRAM BPX43. It was connected to a microcontroller of the type CD4066B by Texas Instruments, which made it possible to adjust the threshold of the phototransistor, to create a binary signal. According to the data sheets the rise and fall time of the phototransistor is between 9 and 18 μs, and the switching characteristics of the microcontroller are below 70 ns. As the latency we seek to measure is in the ms region, these delays are negligible. A complete system overview can be seen in Figure 12.1.

As a trigger for the start and the stop of our measurement, we chose to change the current speed visualisation. The speed value changed between 0 and 200 . The phototransistor was facing the first digits of the display. We chose these speed values to change from one to the other, because this results in the first two digits being lit up and the value of 200 covers a slightly larger area than the value of 100 and is a valid number in most cars, if one considers it to be interpreted as kph. The start trigger of changing the visualisation was sending out the CAN message with the speed value and the stop trigger was the phototransistor changing its state.

We performed two different measurements, one where the response was received via a RS232 connector and in the other case the response was received via a digital I/O port of the CAN hardware.

The measurements were repeated in a time window of several hours to determine the induced latency on average. As the HMI rendering as well as the CAN measurement software were running on a PC, the timings can differ quite strongly, because the computational load of the operating system varies.

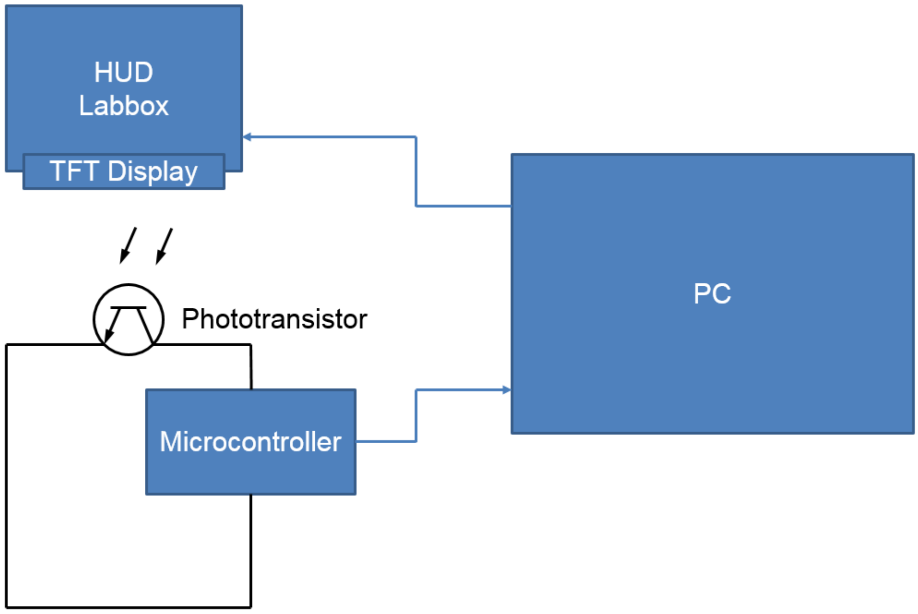

Figure 12.1.: System overview of the latency measurement of the HUD unit.

12.3.1. Results and Discussion

In the case of the RS232 connection, the mean response time was found to be 103.70 ms (std. dev. 15.54 ms) and in the case of the digital I/O port of the CAN hardware we found a mean response time of 124.06 ms (std. dev. 14.67 ms).

We expected the CAN digital I/O results to be lower than the RS232 ones, as the trigger and response was taking place in the same hardware, but this expectation was not confirmed by our results. There are several possible explanations for the results. The results hint that the I/O port of the CAN hardware performs a processing of the signal, leading to the higher response time. Moreover, one needs to take into account that these were two different measurements running on a PC. Therefore, depending on the computational load of the computer, the results can differ. Also both results show a standard deviation of around 15 ms, so the actual response time might very well be around 115 ms.

12.4. Measurements of the End-to-End Latency

After the measurement of the latency of the HUD unit alone, we sought to determine the latency of the complete system inside the vehicle. For this purpose we developed a timestamp generator, to be able to observe the difference between the current timestamp in the environment and the processed timestamp displayed on the HUD. This measurement is described in the first part of this section.

In case of a particular application, the observed latency would be different than the timestamp difference, as the involved sensors often perform prediction techniques. Therefore, we also measured the observed latency of a particular application and present the results in the last part of this section.

12.4.1. Measurement by a timestamp generator

This measurement is inspired by the work of Billeter et al. in [BRW+16]. In their work, they describe the development of a timestamp generator and the automatic calculation of the end-to-end latency. The timestamp generator encodes the current time in Gray code, and toggles an array of 16 LEDs accordingly. One of the main reasons for encoding the time in Gray code was that two consecutive timestamps differ only by one bit.

We developed a similar timestamp generator based on a Raspberry Pi. The time is shown on a seven-segment display, in a human readable format, to easily determine the induced latency. The timestamp generator is then placed in front of the vehicle's camera and recorded. The area of the image observing the current timestamp is then extracted and this part of the image is displayed on the HUD.

An external camera (Sony RX100) is observing the scene, recording the current timestamp in the world and the processed timestamp on the HUD. By subtracting the processed timestamp from the current timestamp one can calculate the resulting end-to-end latency. An overview can be seen in Figure 12.2.

Results and Discussion

For determining the end-to-end latency, we manually observed the current timestamp shown on the timestamp generator, and subtracted the processed timestamp shown on the HUD. The external camera had a refresh rate of 100 Hz, while the vehicle's camera had an update rate of 30 Hz. Therefore, we naturally found a difference between investigating the first image and the last image of the processed timestamp compared to the current timestamp on the timestamp generator.

We revealed the average latency from investigation of the first image to be 142.05 ms (std.dev. 15.92 ms) and the average latency in the last image to be 159.47 ms (std.dev. 17.08 ms). The global average of the two is therefore 150.76 ms.

In conclusion this means by subtracting the found latency of the HUD unit alone of 115 ms an extra of about 35 ms is induced by the processing of the vehicle's camera and the communication between the HMI application and the vehicle's camera.

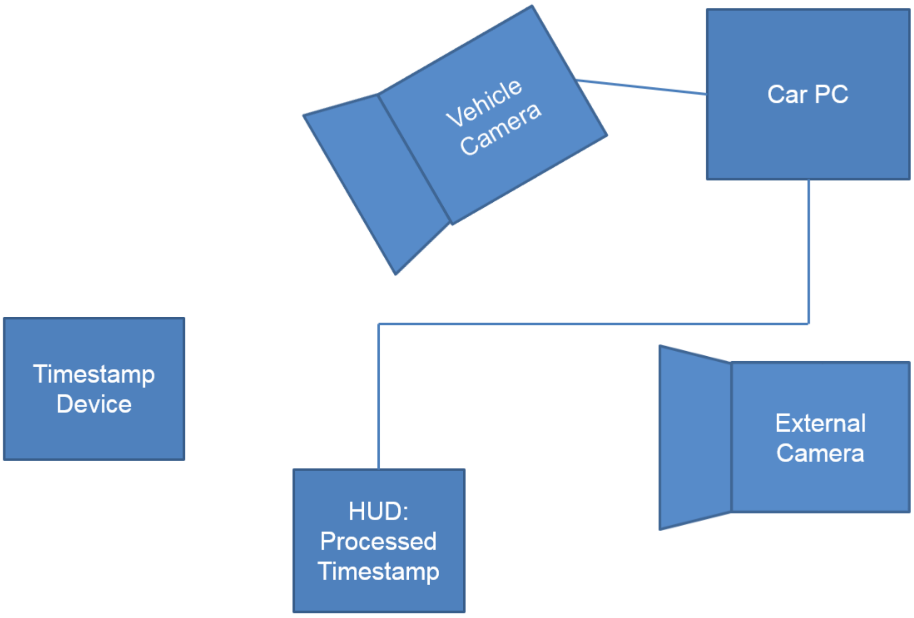

Figure 12.2.: System overview of the latency measurement of the end-to-end latency.

12.4.2. Latency measurement while analysing a specific application

In the previous section, we described one way of determining the end-to-end latency. In an actual application, the involved sensors often perform some prediction and the observed latency can therefore be lower. For this reason we chose a particular application to measure the actual end-to-end latency.

The application chosen was supporting the Adaptive Cruise Control (ACC). The ACC system keeps a defined distance (measured in seconds at the current travel speed) to the vehicle in front, and keeps the speed within a fixed limit. Especially when drivers experience the system the first few times, they naturally distrust the application to recognise the vehicle in front. In such cases, the HUD is ideal for reassuring the user that the system is recognising the vehicle in front, by highlighting the respective vehicle.

In our test vehicle, the leading vehicle was captured using a stereo video camera, designed for automative applications. The stereo video camera extracts the position of the object in front and classifies it, e.g. a car or a motorcycle. This information is then sent to our HMI software via Ethernet. The camera itself performs some pre-filtering for determining the position and classifying the object. If the object is classified as a car our HMI software then renders a virtual depiction of a clamp just below that object.

In order to perform the measurement we drove the leading vehicle along a certain route and came to a halt. This situation was simultaneously recorded by a video camera, which observes the leading vehicle and the trailing virtual clamp. The resulting overall latency was then manually determined, by investigating the video to count the frames between when the leading vehicle reaches a certain position and how long it takes for the virtual clamp to reach the same position. In Figure 12.3 this situation is visualised.

Another measurement investigated the initial time which was required to first recognise the leading vehicle. For this purpose, we obstructed the view of the camera with a piece

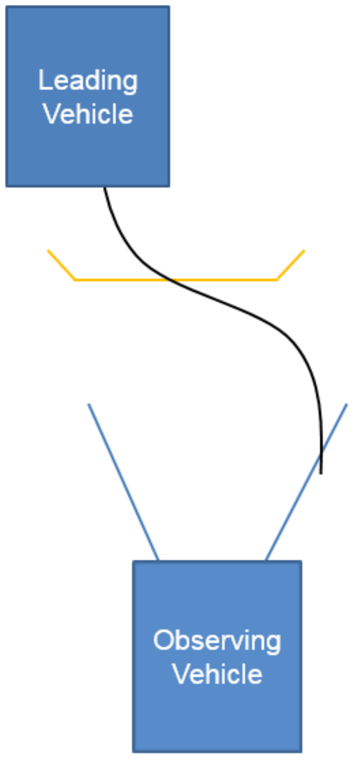

Figure 12.3.: The leading vehicle drives the given path and comes to a halt. The trailing virtual clamp is following the vehicle. The observing vehicle is capturing the scene and when the leading vehicle came to a stop, the time is measured until the clamp reaches this position.

of cardboard and measured the time the leading vehicle was visible until the virtual clamp was being displayed and remained steadily at its position.

Results and Discussion

We found the latency of the virtual clamp following an already tracked vehicle in front to be 73.33 ms (std.dev. 23.09 ms). For the initial detection of the vehicle in front, its correct classification as as car, and displaying the ACC clamp took approximately 1.3 s (std.dev. $\approx$ 0.05 s).

These results show that an actual AR HUD application has a lower latency than the complete end-to-end latency as the sensor is performing a prediction of the objects in front. Hence, using prediction algorithms are a necessity for displaying AR HUD applications smoothly.

12.5. Summary

This section gave a general overview of the latency in the current system. We started by measuring the latency of the HUD unit alone and then went on to measure the end-to-end latency. To measure the end-to-end latency we used two different techniques. One measured the raw end-to-end latency by passing through the image of a recorded timestamp by the vehicle's camera and displaying the image on the HUD. The complete latency was then determined by substracting the processed timestamp from the current timestamp on the timestamp device.

In an actual AR HUD application the data would be prefiltered and some kind of prediction performed. We therefore conducted a second measurement, in which we investigated the occuring latency of the ACC visualisation.

These measurements gave a first impression of what amount of latency is induced in the current prototype of an AR HUD application. To determine the threshold until which the latency is actually perceivable, we conducted user studies which are presented in the following chapter.

13. Perceivable latency of AR applications in the HUD

In the previous chapter, we measured the latency induced by the AR HUD prototype. One important factor during the development of an AR system is to determine which threshold of latency is actually perceivable, i.e. to what extent needs the system to be improved to meet the criteria of creating a seamless AR experience. For touchscreen applications, this threshold has already been determined [NLW$^+$12]. However, it is unclear if the results of the touchscreen applications can directly be applied to the case of AR applications. We therefore conducted two experiments to investigate whether we are able to reproduce the touchscreen results for AR applications as well, or if they differ.

The first experiment investigates the perceivable latency of a virtual object following a real object, the depiction of a car in our case. The setup of the experiment is similar to the setup of the user study in [NLW$^+$12]. In essence, we were able to reproduce the results of their touchscreen user study for AR applications. One major drawback of the experiment is that, users merely determine the spatial distance between the car depiction and the following object, than judging the temporal offset, as pointed out in [JNDW13]. Therefore, we conducted a second experiment, investigating the time difference of the appearing car depiction, and the appearing virtual object with a certain delay. These results are presented in the second part of this chapter.

13.1. Introduction

Besides measuring the actual latency of the AR HUD prototype as described in Chapter 12, it is important to determine the level of latency which is actually perceivable. This threshold is vital, as it marks the threshold to which the engineering of the system is necessary. While this has been determined for touchscreen applications [NLW$^+$12] [JNDW13], in the case of AR applications, this threshold is still unknown. Therefore, we used a similar experimental setup, as the touchscreen applications, for a proposed ACC application, in which a virtual clamp is following a real car.

As we were required to control the temporal error, we needed to determine this error in a virtual setting. Therefore, the car the virtual clamp is following after is a virtual depiction. In the virtual setting we were able to control the temporal delay of the virtual clamp towards the depiction of the car. The first part investigates the following of the clamp to the moving car. The second part investigates the delay between the appearance of the car depiction and the appearance of the virtual clamp.

13.2. Perceivable Latency of a virtual Clamp following a Car

The goal of this user study was to measure the perceivable latency of AR HUD applications. This user study concentrated on the use case of ACC, where a virtual clamp is following the real car. Because it was necessary to control the overall latency of the system, we conducted the user study in a virtual setting. In this setting the car is depicted as a virtual depiction of it and the delay of the virtual clamp can be set arbitrarily. The setup does not take into account the effects of real world lighting and other factors, which may affect the perception of the latency. It is however important to control the temporal delay, to be able to determine the threshold of perceivable latency over a realistic setting. This was one of the main reasons for choosing this minimal setup in order to determine the level of perceivable latency.

13.2.1. Setup of the User Study

To determine the perceivable latency threshold, we chose to design our user study similar to the one described by Ng et al. [NLW+12]. In their work, they describe how they determined the Just-Noticeable Difference (JND) of a dragging task on touchscreens. The design of their user study was a two alternative forced-choice setup. Particpants were asked to move their finger from left to right and right to left on the touchscreen and then had to choose which of the two was "faster". We adapted this design to our setup, where the test subjects were shown two sequences, each 5 seconds long, and with no possibility of seeing any of the sequences again. The sequences showed the depiction of the car with a trailing virtual clamp. One sequence had a delay of 1 ms between the car depiction and the virtual clamp, and was termed the *reference*. In the other sequence the delay was between 2 ms and 200 ms, and was termed the *probe*. The order of the two sequences was randomised and after each two sequences were shown, the participants had to choose the sequence which they thought had the lower delay.

The JND is described as the level where participants are able to identify the reference case in 75% of the cases, also referred to as X_{75}. The setup is based on the weighted-up down method described by Kaernbach [Kae91]. In this method correct responses result in a decrease of the latency, i.e. the stimulus perceived, by the base step size (initially 16 ms). Incorrect responses lead to an increase of the latency by three times the base step size. A change from a correct response to an incorrect response, and vice versa, are termed as a *reversal*. These reversals converge towards the value of X_{75}.

The up-down methods are also often referred to as staircase methods because of the resulting shape of a staircase [Kae91]. We ran two staircases interleaved so the participant is unable to keep track of the performance history. One staircase started at the upper boundary of 200 ms, and the other staircase started at the lower boundary of 2 ms. After 10 reversals occurred in both staircases, the experiment was stopped. For each staircase, the values of the first 10 reversals were recorded. At every reversal the base step size was halved until a minimum of 1 ms.

13.2.2. Participants

We conducted the experiment with 10 participants, 9 of them being male and one being female. The average age was 32.8 years (std.dev. 8.5) with the youngest test subject being 22 and the oldest being 47 years old.

13.2.3. Evaluation of the User Study

For each participant we calculated the arithmetic mean of the last 5 reversals. As described earlier, a reversal describes the change from a correct response to an incorrect one, and vice versa. As the step size was halved at every reversal, with an initial step size at 16 ms, the base step size in the last 5 of a total of 10 reversals is at 1 ms. This is the accuracy we desired for the perceivable latency.

13.2.4. Independent Variables

The independent variable in this user study was the ACC visualisation in which a virtual clamp is following the depiction of a car.

13.2.5. Dependent Variables

The dependent variable in this user study was the calculation of Just-Noticeable Difference (JND), calculated by the average of the last 5 reversals. The JND was calculated for each participant.

13.2.6. Results and Discussion

The calculated JND levels and the respective standard deviation for each participant can be seen in Figure 13.1.

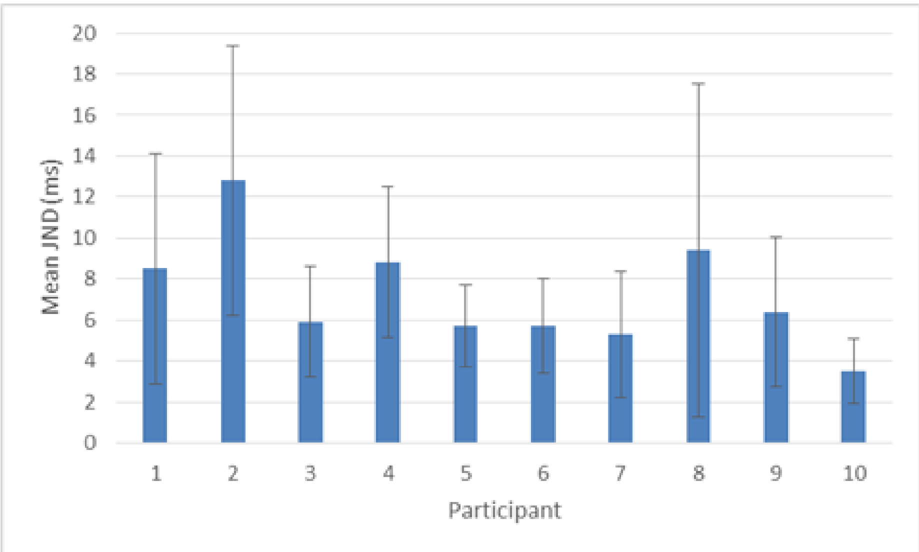

Figure 13.1.: Overview of JND levels of each participant.

The minimum of the JND levels over the last 5 reversals was 3.5 ms and the maximum was 12.8 ms. The overall average was found to be 7.2 ms (std.dev. 2.68 ms). These results are in accordance with the results of [NLW+12]. For the case of the virtual clamp following the car the results argue for a correlation in the perception of latency for AR applications and touchscreen applications. However, one possibility can be that the participants are mostly focusing on the spatial distance between the virtual clamp the car, rather than concentrating on the temporal delay. Therefore we conducted a second experiment to investigate the temporal perception of AR applications further.

13.3. Latency Perception of appearing Objects in an AR application

After determining the level of perceivable latency in the previous study, and finding the JND levels similar to the ones perceivable in the touchscreen cases, we set up a further user study to determine the level of perceivable latency in the case of appearing objects. As Jota et al. [JNDW13] pointed out, the perceived latency in touchscreen applications is considerably higher in the case of tapping compared to dragging on the touchscreen. In the case of dragging, the user has in addition to the temporal feedback, also a spatial feedback from the visualisation following the finger. In their work, they found the JND in the tapping case a lot higher (64 ms compared to 6 ms). Therefore we investigated whether these results can be reproduced for AR applications.

13.3.1. Setup of the User Study

The setup of the user study was similar to the previous study (13.2.1) and based on Ng et al. [NLW+12] and Jota et al. [JNDW13]. Again the design was a two alternative, forced-choice setup. To describe the setup in short, participants were shown two sequences. In one of the sequences the virtual clamp was appearing with the minimum latency of 1 ms, termed the reference, and in the other sequence the clamp appeared with a delay between 2 and 200 ms, termed the probe. Unlike the previous experiment, here, after seeing both sequences, the participant had the possibility to either see the first sequence again or decide right away. Again, the setup was based on a weighted up-down method described by Kaernbach [Kae91]. If the participant identified the reference correctly, the latency of the probe was decreased by the base step size. If the participant selected the wrong sequence, the latency of the probe was increased by three times the base step size. The base step size was initially 16 ms. A change from a correct response to an incorrect and vice versa was termed a reversal. After each reversal the step size was halved, until a step size of 1 ms. The up-down methods are also referred to as staircase methods because of the resulting shape of a staircase [Lev71]. Similar to the previous experiment, one staircase started at the maximum latency of 200 ms, and another staircase started at the minimum latency of 2 ms. The staircases were run interleaved, so that the participant is not biased by the preceding stimuli. The experiment ended after 10 reversals occurred in both staircases. The JND was calculated by calculating the arithmetic mean of the last 6 reversals of each staircase.

13.3.2. Participants

We conducted the experiment with 20 participants, 11 of them being male and 9 being female. The average age was 35.85 years (std. dev. 9.41 years) with the youngest test subject being 22 and the oldest being 59 years old.

13.3.3. Independent Variables

The independent variable in this user study was the visualisation shown to the users, displaying the depiction of a car first and a virtual clamp appearing with a given delay.

13.3.4. Dependent Variables

The dependent variable in this user study was the delay between the appearance of the car depiction and the virtual clamp.

13.3.5. Results and Discussion

The calculated results and the respective standard deviation for each participant are shown in Figure 13.2.

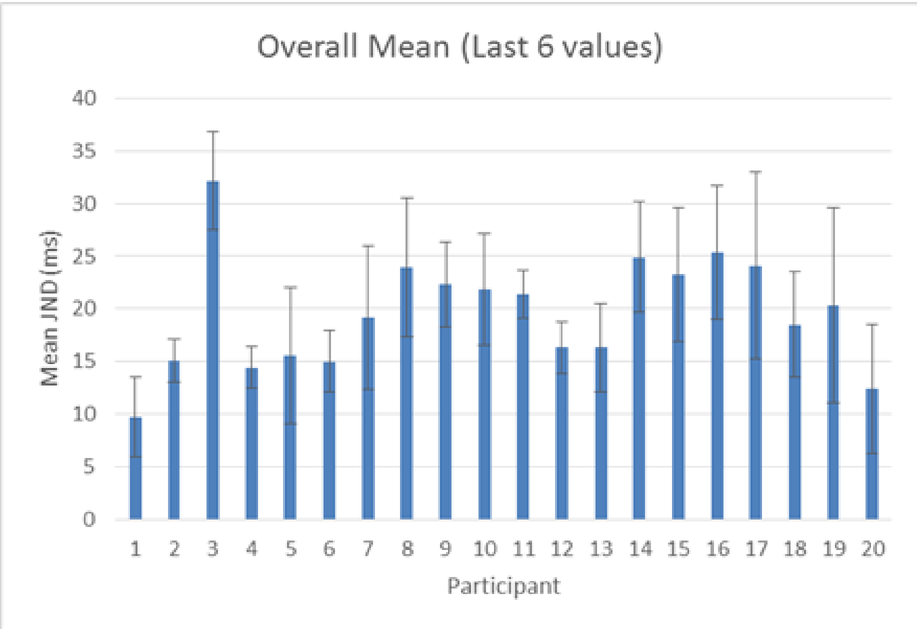

Figure 13.2.: Overview of JND levels of each participant.

The mean JND, over all participants, was found to be 19.60 ms with a std. dev. of 5.35 ms. The minimum JND was 9.75 ms and the highest was 32.17 ms. These results are considerably lower than the results for touchscreen applications as found by Jota et al. [JNDW13]. They determined JND levels for touchscreen applications of 64 ms with a std. dev. of 24 ms. The perceived latency in the case of touchscreen applications and the appearance of AR visualisations differs. In the case of the dragging task on touchscreens and the following of a virtual clamp the results are comparable. In terms of tapping on a touchscreen, and the appearance of an object, and the visualisation attached to it the results differ. One possible explanation for this outcome is the fact that in the case of touchscreens, the appearance of the square visualisation is triggered by a haptic trigger and the response is visual. In the case of AR visualisations, such as the one of the ACC visualisation we investigated, the trigger and the response are both visual. As the trigger and response of touchscreens are in two different channels, haptic and visual, and in the case of AR visualisations trigger and response are in the same channel, visual, this could well explain the lower threshold in the perceived latency of the latter.

The results show that in the case of appearance the perceived latency is considerably lower in AR applications than touchscreen applications. The perception of the latency level can be even lower, as the user study was performed on display with a refresh rate of 60 Hz, therefore a new image is presented every $16.\bar{6}ms$. We did not expect the perceivable latency to be this low, therefore future work can investigate if the perception threshold is even lower.

13.4. Summary

Our results show that the perception of latency in the dragging use case of touchscreen applications [NLW$^+$12] is comparable to the trailing ACC clamp in AR applications. We postulate that in both cases the participants are mostly basing their decision on the spatial distance between the finger and the visualisation in the case of touchscreen applications, or the distance between the car and the virtual clamp in the AR case. [JNDW13] already pointed this out and performed a follow-up study on the tapping use case of touchscreen applications. Following a similar logic, we performed a follow-up user study to investigate if the perception of latency is also comparable in the appearance of objects in the AR case.

With these two user studies, we were able to show that the perception of latency is similar in the dragging case of touchscreen applications to the following of virtual objects after a reference object, which can often be experienced in the case of OST-AR applications. The same perception of latency does not hold in the case of tapping a touchscreen and a visualisation appearing compared to a reference object appearing and a virtual object afterwards. Examples of these can include the sudden appearance of a real object (e.g. the view was obstructed) or turning on a display device. One can argue that the user is able to distinguish between initialisation of the system or the initial registration of an object compared to tracking and following an object and the tolerance for the initial phase is much higher than in the active phase. Nevertheless, this work shows that the perception thresholds are quite lower than they have been estimated so far. Assessing the perception levels of latency in AR applications is an important topic in order to create a seamless AR experience.

Part VI.

Conclusion and Future Work

This part gives an overall conclusion of the work conducted in this thesis. It summarises the main achievements of this thesis and points out future research areas.

14. Conclusion

This work investigated various areas critical for developing robust AR HUD applications. As the AR HUD is a component in which several disciplines, such as engineering, computer science and psychology meet, the investigation of these areas can be vast. We therefore concentrated on a few specific aspects, such as the spatial accuracy, the limited display area and the temporal aspect, always keeping the user in mind by performing user studies with the potential AR HUD applications.

We began with an analysis of the accuracy of current GNSS systems used in series vehicles. The results show that the current positional error of approximately 6 metres is still too high to steadily position an AR visualisation at a given location. Therefore we developed a novel AR visualisation not requiring a high positional accuracy, being solely dependent on the distance to the turn. Through its design and animation it is conveying a spatial appearance and creating a dynamic behaviour. This visualisation was well perceived by the participants of a comparative user study.

Furthermore we developed a visualisation of the future road course in the HUD, the 3D-FRC. The 3D-FRC's only data source is the electronic horizon which is already available in current premium series vehicles. Therefore it is relatively easy to integrate this visualisation into any given vehicle. The 3D-FRC was designed in an *AR-like* manner to be able to display a dynamic visualisation in the given field of view of the AR HUD prototype. The conducted user study showed an overall improvement of the braking behaviour by participants around sharp corners.

Another area which was investigated by us in depth was the change of the gaze behaviour by HUD and AR HUD applications. The initial assessment of the gaze behaviour showed a significant decrease of the visual attention towards the Head Unit and the Instrument Cluster when the AR HUD was enabled. Furthermore it revealed a significant increase of the visual attention towards the area of the HUD. We therefore investigated, in another experiment, how much the visual attention is already distorted by the introduction of conventional HUD representations and if this distortion is further increased by the introduction of *AR-like* representations. The results showed that already conventional HUD representations increase the visual attention towards the area of the HUD and that this attention is not increased further by the introduction of *AR-like* representations. Taken together, these results show that AR HUD applications have the potential to increase the time drivers spend looking towards the road and also show that one needs to display the HUD information in a subtle way to not create another distraction in the vehicle.

The last area presented in this thesis is the assessment of the temporal aspect of HUD applications. Along with the measurement of the end-to-end latency of the current HUD prototype, the perception of AR applications was also investigated. The end-to-end latency of the complete system was found to be around 150 ms, which is quite high in order to create a seamless AR visualisation. The perceivable latency of the trailing virtual ACC clamp showed, with a just-noticeable difference (JND) of 7 ms, results comparable to the dragging case of touchscreen applications. On the other hand, in the case of appearing objects, the JND level was, at an average of 19.6 ms, considerably lower than the tapping use case of touchscreen applications. The results show that the current prototype is taking too long to display AR visualisations smoothly, but give an impression of what latency levels are actually perceivable.

Taken together this work investigated the areas of sensors (GNSS), the development of AR visualisation concepts (Sails and 3D-FRC), their evaluation through user studies (subjective and objective measures, such as gaze behaviour) and the assessment of the temporal aspect of HUD applications. These areas are essential for the development of AR HUD applications. In all areas results are presented and provide a deeper understanding of what is currently available and what is necessary to successfully develop AR HUD applications. This thesis is another step towards the introduction of Augmented Reality into vehicles.

15. Future work

There are various areas one can investigate further to introduce AR HUD applications into series vehicles. This thesis is another stepping stone towards their introduction into the vehicle. As mentioned in the conclusion, due to its complexity, through the combination of several different research areas, the possibilities of future work regarding the AR HUD are plenty.

The first area of investigating the accuracy of current GNSS sensors shows that the sensors in series vehicles are currently not adequately accurate enough to steadily display full-AR visualisations in the environment. Therefore GNSS sensors with a higher accuracy for AR applications are needed. These high-precision GNSS sensors might also be introduced in the context of automated driving, as the accurate positioning of the vehicle is necessary for these applications as well.

The development of the Sails visualisation and the 3D-FRC show that it is necessary to also take inevitable errors, such as the location inaccuracy or the limited field of view of the HUD, into account. During system development, it is essential that one asserts which parameters can currently not be changed, and accounts for these in the design of the visualisation concept.

Furthermore, it is also of high importance that the potential AR HUD applications are developed in an actual prototype. The complete perception of the visualisations changes in the actual prototype when compared to simulators. Therefore we argue for user studies in real prototypes. As pointed out several times, it is important to frequently assess the quality of the AR HUD applications in real traffic situations.

Another area which was investigated by this work in detail was the change in gaze behaviour by the introduction of HUDs. This topic can be investigated further by assessing the role of optical focus in HUD applications. Further studies could devise experiments to investigate this part in detail, by either employing sophisticated hardware offering the opportunity of determining the focus or developing own techniques for determining a change in focus of users.

In order to create a seamless AR experience it is also necessary to create an HUD system with a considerably low latency. Current prototypes are still not fast enough to achieve this. The first results of the user studies regarding the perception of latency can be investigated further by the use of a latency simulator with a higher sophistication.

This thesis is another step on the path of introducing Augmented Reality into the vehicle. It builds upon the previous work, mostly conducted in simulators and takes the important step of investigating the applications inside an actual vehicle. Naturally there are inevitable side effects which can be introduced by actual traffic situations. However, it is crucial to investigate these applications in the real environment, as their perception changes dramatically in a real prototype. This work is a first step towards this investigation.

A. Appendix

A.1. Questionnaire of the User Study investigating Navigational AR Visualisations (SUS and TLX)

Fragebogen zur AR-HUD Navigation

Benutzer-zentrierte Entwicklung und Evaluation
von Visualisierungen für Navigationsvarianten im AR-HUD

	trifft nicht zu 1	2	3	4	trifft zu 5
1. Ich denke, dass ich die Visualisierung häufig verwenden werde	☐	☐	☐	☐	☐
2. Ich habe die Visualisierung unnötig komplex gefunden	☐	☐	☐	☐	☐
3. Ich denke, dass die Visualisierung einfach zu verstehen war	☐	☐	☐	☐	☐
4. Ich denke, dass ich technische Betreuung benötige, um mit der Visualisierung zurecht zu kommen	☐	☐	☐	☐	☐
5. Ich finde, dass die verschiedenen Funktionen gut in die Visualisierung integriert worden sind	☐	☐	☐	☐	☐
6. Ich denke, dass die Visualisierung zu inkonsistent war	☐	☐	☐	☐	☐
7. Ich kann mir vorstellen, dass die meisten Leute den Umgang mit der Visualisierung sehr schnell lernen	☐	☐	☐	☐	☐
8. Ich finde, dass der Umgang mit der Visualisierung sehr mühselig war	☐	☐	☐	☐	☐
9. Ich habe mich sehr sicher im Umgang mit der Visualisierung gefühlt	☐	☐	☐	☐	☐
10. Ich muss eine Menge über die Visualisierung lernen, bevor ich diese effektiv nutzen kann	☐	☐	☐	☐	☐

NASA-TLX

Benutzer-zentrierte Entwicklung und Evaluation
von Visualisierungen für Navigationsvarianten im AR-HUD

Information: Geben Sie bitte an, wie hoch die Beanspruchung in den einzelnen Dimensionen war. Markieren Sie dazu auf den folgenden Skalen bitte, in welchem Maße Sie sich in den sechs genannten Dimensionen von den Aufgaben beansprucht oder gefordert gesehen haben.

1 Geistige Anforderung

Wie viel geistige Anstrengung war bei der Informationsaufnahme und bei der Informationsverarbeitung erforderlich (z.b. Denken, Entscheiden, Rechnen, Erinnern, Hinsehen, Suchen...)? War die Aufgabe leicht oder anspruchsvoll, einfach oder komplex, erfordert sie hohe Genauigkeit oder ist sie fehlertolerant?

gering hoch

2 Körperliche Anforderungen

Wie viel körperliche Aktivität war erforderlich (z.B. ziehen, drücken, drehen, steuern, aktivieren...)? War die Aufgabe leicht oder schwer, einfach oder anstrengend, erholsam oder mühselig?

gering hoch

3 Zeitliche Anforderungen

Wie viel Zeitdruck empfanden Sie hinsichtlich der Häufigkeit oder dem Takt mit dem Aufgaben oder Aufgabenelemente auftraten? War die Abfolge langsam und geruhsam oder schnell und hektisch?

gering hoch

4 Ausführung der Aufgaben

Wie erfolgreich haben Sie ihrer Meinung nach die vom Versuchsleiter (oder von Ihnen selbst) gesetzten Ziele erreicht? Wie zufrieden waren Sie mit ihrer Leistung bei der Verfolgung dieser Ziele?

gut schlecht

5 Anstrengung

Wie hart mussten Sie arbeiten, um Ihren Grad an Aufgabenerfüllung zu erreichen?

gering hoch

6 Frustration

Wie unsicher, entmutigt, irritiert, gestresst und verärgert (versus sicher, bestätigt, zufrieden, entspannt und zufrieden mit sich selbst) fühlten Sie sich während der Aufgabe?

gering hoch

Kontrollieren Sie bitte, ob Sie zu allen Fragen Angaben gemacht haben. Bei Unklarheiten, wenden Sie sich bitte an die anwesenden Versuchsleiter.

NASA-TLX

Benutzer-zentrierte Entwicklung und Evaluation
von Visualisierungen für Navigationsvarianten im AR-HUD

Im Folgenden werden jeweils zwei der sechs Beanspruchungsdimensionen in verschiedenen Kombinationen gegenübergestellt. Geben Sie jeweils an, welche Beanspruchungsdimension für die Gesamtbeanspruchung, die Sie empfunden haben, bedeutsamer war. Es geht also zunächst nicht darum, wie hoch die Beanspruchung in den einzelnen Dimensionen war, sondern wie wichtig die jeweilige Dimension für das Gesamtempfinden war.

Körperliche Anforderungen	☐	☐	Zeitliche Anforderungen
Anstrengung	☐	☐	Geistige Anforderungen
Frustration	☐	☐	Körperliche Anforderungen
Anstrengung	☐	☐	Frustration
Geistige Anforderungen	☐	☐	Zeitliche Anforderungen
Körperliche Anforderungen	☐	☐	Anstrengung
Zeitliche Anforderungen	☐	☐	Ausführung der Aufgaben

Frustration	☐	☐	**Geistige Anforderungen**
Zeitliche Anforderungen	☐	☐	**Frustration**
Ausführung der Aufgaben	☐	☐	**Anstrengung**
Anstrengung	☐	☐	**Zeitliche Anforderungen**
Frustration	☐	☐	**Ausführung der Aufgaben**
Ausführung der Aufgaben	☐	☐	**Körperliche Anforderungen**
Geistige Anforderungen	☐	☐	**Ausführung der Aufgaben**
Geistige Anforderungen	☐	☐	**Körperliche Anforderungen**

Kontrollieren Sie bitte, ob Sie kein Vergleichspaar vergessen haben.

A.2. Questionnaire of the User Study investigating Navigational AR Visualisations (AttrakDiff)

Segel/Pfeil

Proband Nr.:

Nachfolgend finden Sie Wortpaare, mit deren Hilfe Sie die Beurteilung vornehmen können. Sie stellen jeweils extreme Gegensätze dar, zwischen denen eine Abstufung möglich ist.

Ein Beispiel:

unsympathisch ○ ○ ○ ○ ◉ ○ ○ sympathisch

Diese Bewertung bedeutet, dass das Produkt eher sympathisch, aber noch verbesserungsbedürftig ist.

Denken Sie nicht lange über die Wortpaare nach, sondern geben Sie bitte die Einschätzung ab, die Ihnen spontan in den Sinn kommt. Vielleicht passen einige Wortpaare nicht so gut auf das Produkt, kreuzen Sie aber trotzdem bitte immer eine Antwort an. Denken Sie daran, dass es keine "richtigen" oder "falschen" Antworten gibt - nur Ihre persönliche Meinung zählt!

menschlich	○	○	○	○	○	○	○	technisch
isolierend	○	○	○	○	○	○	○	verbindend
angenehm	○	○	○	○	○	○	○	unangenehm
originell	○	○	○	○	○	○	○	konventionell
einfach	○	○	○	○	○	○	○	kompliziert
fachmännisch	○	○	○	○	○	○	○	laienhaft
hässlich	○	○	○	○	○	○	○	schön
praktisch	○	○	○	○	○	○	○	unpraktisch
sympathisch	○	○	○	○	○	○	○	unsympathisch
umständlich	○	○	○	○	○	○	○	direkt

stilvoll	○	○	○	○	○	○	○	stillos
voraussagbar	○	○	○	○	○	○	○	unberechenbar
minderwertig	○	○	○	○	○	○	○	wertvoll
ausgrenzend	○	○	○	○	○	○	○	einbeziehend
bringt mich den Leuten näher	○	○	○	○	○	○	○	trennt mich von Leuten
nicht vorzeigbar	○	○	○	○	○	○	○	vorzeigbar
zurückweisend	○	○	○	○	○	○	○	einladend
phantasielos	○	○	○	○	○	○	○	kreativ
gut	○	○	○	○	○	○	○	schlecht

verwirrend	○	○	○	○	○	○	○	übersichtlich
abstoßend	○	○	○	○	○	○	○	anziehend
mutig	○	○	○	○	○	○	○	vorsichtig
innovativ	○	○	○	○	○	○	○	konservativ
lahm	○	○	○	○	○	○	○	fesselnd
harmlos	○	○	○	○	○	○	○	herausfordernd
motivierend	○	○	○	○	○	○	○	entmutigend
neuartig	○	○	○	○	○	○	○	herkömmlich
widerspenstig	○	○	○	○	○	○	○	handhabbar

Alter:

Geschlecht:

List of Figures

List of Tables

Bibliography

[ADA] ADASIS. `http://ertico.com/projects/adasisforum/`.

[Arm14] Richard A. Armstrong. When to use the Bonferroni correction. *Ophthalmic and Physiological Optics*, 34(5):502–508, sep 2014.

[Azu97] Ronald Azuma. A survey of augmented reality. *Presence: Teleoperators and Virtual Environments*, 6(4):355–385, 1997.

[BBL15] Adam Bolton, Gary Burnett, and David R Large. An investigation of augmented reality presentations of landmark-based navigation using a head-up display. In *Proceedings of the 7th International Conference on Automotive User Interfaces and Interactive Vehicular Applications - AutomotiveUI '15*, pages 56–63, New York, New York, USA, 2015. ACM Press.

[BKM09] Aaron Bangor, Philip Kortum, and James Miller. Determining what individual sus scores mean: Adding an adjective rating scale. *J. Usability Studies*, 4(3):114–123, May 2009.

[Bro96] John Brooke. SUS - A quick and dirty usability scale. *Usability evaluation in industry*, 189(194):4–7, 1996.

[Bro16] Nora Broy. *Stereoscopic 3D user interfaces: exploring the potentials and risks of 3D displays in cars.* PhD thesis, Universität Stuttgart, 2016.

[BRW+16] Markus Billeter, Gerhard Rothlin, Jan Wezel, Daisuke Iwai, and Anselm Grundhofer. A LED-Based IR/RGB End-to-End Latency Measurement Device. In *2016 IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct)*, pages 184–188. IEEE, sep 2016.

[Com15] Wikimedia Commons. File:lens3b.svg — wikimedia commons, the free media repository, 2015. [Online; accessed 11-May-2018].

[CPMA11] Vassilis Charissis, Stylianos Papanastasiou, Lewis Mackenzie, and Sachi Arafat. Evaluation of collision avoidance prototype head-up display interface for older drivers. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 6763 LNCS(PART 3):367–375, 2011.

[CV95] James E. Cutting and Peter M. Vishton. Perceiving Layout and Knowing Distances: The integration, relative potency, and contextual use of different information about depth. In William Epstein and Sheena Robers, editors, *Handbook of Perception and Cognition: Perception of Space and Motion*, volume 5, pages 69–117. Academic Press, 1995.

[DCT09] Anup Doshi, Shinko Yuanhsien Cheng, and Mohan Manubhai Trivedi. A novel active heads-up display for driver assistance. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(1):85–93, 2009.

[Dho17] Shreyas Dhone. Latency Analysis of Head Up Displays. M.Sc. Thesis, Technische Universität Chemnitz, 2017.

[Dus15] Markus Duschl. *Critical Assessment of Data Visualization from the Electronic Horizon in the Context of a Deceleration Assistance System.* PhD thesis, Technische Universität München, 2015.

[DWSS17] Catherine Diaz, Michael Walker, Danielle Albers Szafir, and Daniel Szafir. Designing for depth perceptions in augmented reality. In *2017 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 111–122. IEEE, 2017.

[EGN] EGNOS. website. `http://www.esa.int/Our_Activities/Navigation/EGNOS/What_is_EGNOS`.

[Erg] Ergoneers. Dikablis professional eye tracking glasses. `http://www.ergoneers.com/en/hardware/dikablis-glasses/`.

[FMHY97] Steven Feiner, Blair Macintyre, Tobias Höllerer, and New York. A Touring Machine : Prototyping 3D Mobile Augmented Reality Systems for Exploring the Urban Environment Columbia University Graduate School of Architecture , Planning and Preservation Columbia University. *Scenario*, 97:74–81, 1997.

[FMS93] Steven Feiner, Blair Macintyre, and Dorée Seligmann. Knowledge-based augmented reality. *Communications of the ACM*, 36(7):53–62, 1993.

[GFK14] Joseph L. Gabbard, Gregory M. Fitch, and Hyungil Kim. Behind the glass: Driver challenges and opportunities for AR automotive applications. *Proceedings of the IEEE*, 102(2):124–136, 2014.

[GS08] Joseph L. Gabbard and J. Edward Swan. Usability engineering for augmented reality: Employing user-based studies to inform design. *IEEE Transactions on Visualization and Computer Graphics*, 14(3):513–525, 2008.

[Has04] Marc Hassenzahl. The Interplay of Beauty, Goodness, and Usability in Interactive Products. *Human Computer Interaction*, 19(4):319–349, 2004.

[Hol] Microsoft HoloLens. `https://www.microsoft.com/en-us/hololens`.

[HS88] Sandra G. Hart and Lowell E. Staveland. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. *Advances in Psychology*, 52(C):139–183, 1988.

[HSK09] Manuel Huber, Michael Schlegel, and Gudrun Klinker. Temporal calibration in multisensor tracking setups. *Science and Technology Proceedings - IEEE 2009 International Symposium on Mixed and Augmented Reality, ISMAR 2009*, pages 195–196, 2009.

[HSK14] Manuel Huber, Michael Schlegel, and Gudrun Klinker. Application of Time-Delay Estimation to Mixed Reality Multisensor Tracking. *Journal on Virtual Reality and Broadcasting (JVRB)*, 11(3), 2014.

[HUD] HUDway. `http://hudwayapp.com/`.

[IDAK15] Yuta Itoh, Maksym Dzitsiuk, Toshiyuki Amano, and Gudrun Klinker. Semi-Parametric Color Reproduction Method for Optical See-Through Head-Mounted Displays. *IEEE Transactions on Visualization and Computer Graphics*, 21(11):1269–1278, nov 2015.

[IOH+16] Yuta Itoh, Jason Orlosky, Manuel Huber, Kiyoshi Kiyokawa, and Gudrun Klinker. OST Rift: Temporally consistent augmented reality with a consumer optical see-through head-mounted display. In *2016 IEEE Virtual Reality (VR)*, pages 189–190. IEEE, mar 2016.

[ISO98] ISO. Ergonomic requirements for office work with visual display terminals (VDTs) – Guidance on usability. Standard 9241-11:1998, International Organization for Standardization, Geneva, CH, March 1998.

[ISO10] ISO. Ergonomics of human-system interaction – Human-centred design for interactive systems. Standard 9241-210:2010, International Organization for Standardization, Geneva, CH, March 2010.

[ISO14] ISO. Road vehicles - measurement of driver visual behaviour with respect to transport information and control systems - part 1: Definitions and parameters. Standard 15007-1:2014, International Organization for Standardization, Geneva, CH, November 2014.

[Isr12] Boris Israel. *Potenziale eines kontaktanalogen Head-up Displays für den Serieneinsatz*. PhD thesis, TU Munich, 2012.

[Jan18] Axel Jansen. *Augmented Reality in der Fahrzeugnavigation*. PhD thesis, Technische Universität München, 2018.

[JNDW13] Ricardo Jota, Albert Ng, Paul Dietz, and Daniel Wigdor. How Fast is Fast Enough? A Study of the Effects of Latency in Direct-Touch Pointing Tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*, number 1, page 2291, New York, New York, USA, 2013. ACM Press.

[Kae91] Christian Kaernbach. Simple adaptive testing with the weighted up-down method. *Perception & Psychophysics*, 49(3):227–229, may 1991.

[KD09] SeungJun Kim and Anind K. Dey. Simulated augmented reality windshield display as a cognitive mapping aid for elder driver navigation. In *Proceedings of the 27th international conference on Human factors in computing systems - CHI 09*, page 133, New York, New York, USA, 2009. ACM Press.

[KKKO13] Takaya Kawamata, Itaru Kitahara, Yoshinari Kameda, and Yuichi Ohta. Poster: Lifted road map view on windshield display. In *2013 IEEE Symposium on 3D User Interfaces (3DUI)*, pages 139–140. IEEE, mar 2013.

[KMW16] Steffen Kutter, Stefan Meier, and Christian Wiesner. Verfahren zum Betreiben eines Head-Up-Displays, Head-Up-Display Einrichtung, 2016. DE102016203080.

[KWGP13] Hyungil Kim, Xuefang Wu, Joseph L Gabbard, and Nicholas F Polys. Exploring head-up augmented reality interfaces for crash warning systems. *Proceedings of the 5th International Conference on Automotive User Interfaces and Interactive Vehicular Applications - AutomotiveUI '13*, pages 224–227, 2013.

[Lev71] H. Levitt. Transformed Up-Down Methods in Psychoacoustics. *The Journal of the Acoustical Society of America*, 49(2B):467–477, feb 1971.

[LSG+03] Mark A. Livingston, J. Edward Swan II, Joseph L. Gabbard, Tobias H. Höllerer, Deborah Hix, Simon J. Julier, Yohan Baillot, and Dennis Brown. Resolving multiple occluded layers in augmented reality. In *The Second IEEE and ACM International Symposium on Mixed and Augmented Reality, 2003. Proceedings.*, pages 56–65. IEEE Comput. Soc, 2003.

[MD16] Tobias Müller and Ralf Dauenhauer. A taxonomy for information linking in augmented reality. In *International Conference on Augmented Reality, Virtual Reality and Computer Graphics*, pages 368–387. Springer, 2016.

[Met] Meta. `https://www.metavision.com/`.

[MKPP11] Zeljko Medenica, Andrew L. Kun, Tim Paek, and Oskar Palinko. Augmented reality vs. street views: a driving simulator study comparing two emerging navigation aids. *Proceedings of the International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '11)*, pages 265–274, 2011.

[MTUK94] Paul Milgram, Haruo Takemura, Akira Utsumi, and Fumio Kishino. Augmented Reality: A class of displays on the reality-vistuality continuum. *Systems Research*, 2351(Telemanipulator and Telepresence Technologies):282–292, dec 1994.

[NDPR09] Simon Nestler, Markus Duschl, Darya Popiv, and Mariana Rakic. Concept for visualizing concealed objects to improve the driver's anticipation. *17th World Congress on Ergonomics (IEA 2009)*, 2009.

[Neu12] Alexander Neumann. *Simulationsbasierte Messtechnik zur Prüfung von Head-up Displays.* PhD thesis, Technische Universität München, 2012.

[Nie12] Jakob Nielsen. Usability 101: Introduction to usability. `https://www.nngroup.com/articles/usability-101-introduction-to-usability/`, 2012.

[NLW+12] Albert Ng, Julian Lepinski, Daniel Wigdor, Steven Sanders, and Paul Dietz. Designing for low-latency direct-touch input. In *Proceedings of the 25th annual ACM symposium on User interface software and technology - UIST '12*, page 453, New York, New York, USA, 2012. ACM Press.

[One] Magic Leap One. `https://www.magicleap.com/magic-leap-one`.

[PWB15] Lisa Pfannmüller, Matthias Walter, and Klaus Bengler. Lead me the right way ?! The impact of position accuracy of augmented reality navigation arrows in a contact analogue head-up display on driving performance, workload, and usability. *Proceedings of the 19th Triennial Congress of the International Ergonomics Association(IEA) 2015*, (August), 2015.

[REKB06] Christian Ress, Aria Etemad, Detlef Kuck, and Marcus Boerger. Electronic Horizon - Supporting ADAS Applications with Predictive Map Data. *ITS European Congress*, 2006.

[Rif] Oculus Rift. `https://www.oculus.com/rift/`.

[Ruf16] Mike Ruf. Entwurf und Umsetzung einer visuellen Fahrerrückmeldung bei eingeschränktem Sichtfeld durch Analyse des elektronischen Horizonts für ein Augmented Reality Head-Up-Display. B.Sc. Thesis, Hochschule Karlsruhe, 2016.

[SC99] Daniel J Simons and Christopher F Chabris. Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception*, 28(9):1059–1074, 1999.

[Sch08] Marcus Schneid. *Entwicklung und Erprobung eines kontaktanalogen Head-up-Displays im Fahrzeug.* PhD thesis, TU Munich, 2008.

[Shr11] Anand Shroff. *An Eye on Numbers: A Ready Reckoner in Ophthalmology*, page 97. 2011.

[SJS14] William Steptoe, Simon Julier, and Anthony Steed. Presence and discernability in conventional and non-photorealistic immersive augmented reality. *ISMAR 2014 - IEEE International Symposium on Mixed and Augmented Reality - Science and Technology 2014, Proceedings*, pages 213–218, 2014.

[SKKO06] Akihiko Sato, Itaru Kitahara, Yoshinari Kameda, and Yuichi Ohta. Visual navigation system on windshield head-up display. *Proc. 13th World . . .*, pages 1–8, 2006.

[SLS+06] J. Edward Swan II, Mark A. Livingston, Harvey S. Smallman, Dennis Brown, Yohan Baillot, Joseph L. Gabbard, and Deborah Hix. A Perceptual Matching Technique for Depth Judgments in Optical, See-Through Augmented Reality. In *IEEE Virtual Reality Conference (VR 2006)*, pages 19–26. IEEE, 2006.

[SP11] Ali Serpengüzel and Andrew W. Poon. *Optical Processes in Microparticles and Nanostructures: A Festschrift Dedicated to Richard Kounai Chang on His Retirement from Yale University.* 2011.

[SSBG15] Missie Smith, Jillian Streeter, Gary Burnett, and Joseph L. Gabbard. Visual search tasks. In *Proceedings of the 7th International Conference on Automotive User Interfaces and Interactive Vehicular Applications - AutomotiveUI '15*, pages 80–87, New York, New York, USA, 2015. ACM Press.

[SSE15] J. Edward Swan, Gurjot Singh, and Stephen R. Ellis. Matching and Reaching Depth Judgments with Real and Augmented Reality Targets. *IEEE Transactions on Visualization and Computer Graphics*, 21(11):1289–1298, nov 2015.

[SSK+07] Tobias Sielhorst, Wu Sa, Ali Khamene, Frank Sauer, and Nassir Navab. Measurement of absolute latency for video see through augmented reality. In *2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality*, pages 1–4. IEEE, nov 2007.

[Stu] CGI Studio. `http://www.fujitsu.com/us/products/devices/semiconductor/gdc/products/cgistudio.html`.

[Sut68] Ivan E. Sutherland. A head-mounted three dimensional display. *Proceedings of the December 9-11, 1968, fall joint computer conference, part I on - AFIPS '68 (Fall, part I)*, page 757, 1968.

[TBK06] Marcus Tönnis, Verena Broy, and Gudrun Klinker. A survey of challenges related to the design of 3D user interfaces for car Drivers. *Proceedings - IEEE Virtual Reality*, 2006:134, 2006.

[TK06] Marcus Tönnis and Gudrun Klinker. Effective control of a car driver's attention for visual and acoustic guidance towards the direction of imminent dangers. In *2006 IEEE/ACM International Symposium on Mixed and Augmented Reality*, pages 13–22. IEEE, oct 2006.

[TKK08] Marcus Tönnis, Leslie Klein, and Gudrun Klinker. Perception thresholds for augmented reality navigation schemes in large distances. *Proceedings - 7th IEEE International Symposium on Mixed and Augmented Reality 2008, ISMAR 2008*, pages 189–190, 2008.

[TKP09] Marcus Tönnis, Gudrun Klinker, and Marina Plavšic. Survey and Classification of Head-Up Display Presentation Principles. *Proceedings of the International Ergonomics Association (IEA)*, 2009.

[TPK13] Marcus Tönnis, David A. Plecher, and Gudrun Klinker. Representing information - Classifying the Augmented Reality presentation space. *Computers and Graphics (Pergamon)*, 37(8):997–1011, 2013.

[WK17] Christian A. Wiesner and Gudrun Klinker. Overcoming location inaccuracies in augmented reality navigation. In *International Conference on Augmented Reality, Virtual Reality and Computer Graphics*, pages 377–388. Springer, 2017.

[WLLGB03] Christopher D Wickens, John Lee, Yili D Liu, and Sallie Gordon-Becker. Introduction to human factors engineering. 2003.

[WRSK16] Christian A. Wiesner, Mike Ruf, Demet Sirim, and Gudrun Klinker. Visualisation of the electronic horizon in Head-Up-Displays. In *2016 IEEE International Symposium on Mixed and Augmented Reality (ISMAR) -Adjunct, 2016*. IEEE, 2016.

[WRSK17] Christian A. Wiesner, Mike Ruf, Demet Sirim, and Gudrun Klinker. 3D-FRC: Depiction of the future road course in the Head-Up-Display. In *2017 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 136–143. IEEE, Oct 2017.

[ZWL+14] Feng Zheng, Turner Whitted, Anselmo Lastra, Peter Lincoln, Andrei State, Andrew Maimone, and Henry Fuchs. Minimizing latency for augmented reality displays: Frames considered harmful. In *2014 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 195–200. IEEE, sep 2014.