Ref. Ares(2018)4033489 - 31/07/2018

Delivery Report for

MeBeSafe

Measures for behaving safely in traffic

Deliverable Title	Vehicle Measures evaluation
Deliverable	D2.1
WP	WP2 In-vehicle nudging solutions
Task	Task 2.1 Sensing driver and vehicle state Task 2.2 Sensing and predicting cyclist intent Task 2.3 Hazard perception and prediction Task 2.4 In-vehicle nudges

This project (MeBeSafe) has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 723430.

Disclaimer

Copyright

Grant Agreement No.	723430
Project Start Date	01/05/2017
Project End Date	31/10/2020
Duration of the Project	42 months
Deliverable Number	D2.1
Deliverable Leader	Olaf Op den Camp, TNO
Dissemination Level (Confidentiality)	Public
Nature	Report
Status	Final
Due Date	M15 (July 2018)
Deliverable Leader/ Responsible Organization	Olaf Op den Camp, TNO
Contributing Author(s)	Bram Bakker, Cygnify Marie-Christin Harre, OFFIS Jeroen Uittenbogaard, TNO Olaf Op den Camp, TNO Mikael Ljung Aust, Volvo Cars
Reviewer(s)	Teri Lillington, Shell
Formal review and upload by Coordinator	Institute for Automotive Engineering - RWTH Aachen University (ika)

Abstract
This report describes different ideas for nudging solutions that can be implemented in vehicles to nudge the driver to: o Make better use of safety functions onboard state-of-the-art vehicles that are equipped with various advanced driver assistance systems. The ideas for making better use of safety functions will be elaborated as part of the MeBeSafe coaching framework in WP4. In this report only an introduction to this type of in-vehicle solutions has been given. o Direct their attention to potential hazards on the road. As a use case for this type of nudging, we focus at the interaction between cyclists and passenger cars on the road; representing a large number of traffic casualties which is difficult to address by current advanced driver assistance systems due to the high maneuvrability of cyclists. To direct the attention to potential hazards, two basic system components are needed: 1. A model to estimate the level and type of hazard and 2. A human-machine-interface to provide appropriate information regarding this hazard to a driver. The report describes the set up of such a hazard prediction model and its components: a static world model referring to road layout and traffic rules, a dynamic world model that considers the actual detections of potential hazards on the road, and a cyclist trajectory prediction model that is intended to predict where a cyclist is going in the upcoming couple of seconds. Al the information from these components is integrated to estimate a hazard level in an approach of a cyclist intersection. Moreover, different options for transferring information regarding the estimated hazard to the driver have been identified. The report shows which design rules and approach need to be followed to develop these options into an effective human-machine-interface. Both in the hazard model as in the HMI-options, there is room for selecting parameter values that influence the effectiveness of the combined nudging-

solution. In a next step in MeBeSafe WP2, tests with simulations and test with simulators will be used to determine the most promising in-vehicle nudging solution that will actually be implemented as a prototype in one FIAT 500X vehicle for testing in WP5.

Version	Date	Comment
1	12.02.2018	Template and outline
2	21.02.2018	Completed outline and introduction
3	14.06.2018	Draft world model and cyclist trajectory prediction
4	26.06.2018	Contributions regarding ACC and HMI integrated
5	28.06.2018	Hazard prediction model integrated
6	02.07.2018	Peer-review ready report prepared (incl. TNO internal review)
7	10.07.2018	Revision Chapter 4, HMI based on peer-review.
8	17.07.2018	Revision Chapter 1, 2 and 3, based on peer-review.
9	24.07.2018	Final version for official review by coordinator

table of document history

Table of Contents

List of Figures

List of Tables

Acronyms

ACC	Adaptive Cruise Control
ADAS	Advanced Driver Assistance System
AEB	Autonomous Emergency Braking
AR	Augmented Reality
CAN	Control Architecture Network
FCA	FIAT Chrysler Automobiles
FCW	Forward Collision Warning
GPS	Global Positioning System
HMI	Human Machine Interface
HTA	Hierarchical Task Analysis
HUD	Head-up Display
LSTM	Long short-term memory
OSM	Open Street Map
TA	Task Analysis
THW	Time Headway
TTC	Time-to-Collision
VCC	Volvo Car Company

1 Introduction

1.1 MeBeSafe and Work Package 2

The aim of the MeBeSafe project is to develop, implement and validate measures that direct road users towards safer behaviour in common traffic situations. MeBeSafe is planning to do this by changing habitual traffic behaviour using 'nudging' and coaching, with the aim of improving driving behaviour. Nudging is a technique that subconsciously stimulates drivers to drive safer, while with coaching, drivers are given feedback on their driving behaviour by a coach in order to learn about their own driving behaviour and enhance driving performance. The main aim of WP2 is the development and implementation of in-vehicle hardware and software solutions to nudge drivers of passenger cars to show safer behaviour. Moreover, an interface to in-vehicle sensor systems, e.g. to provide an off-line coaching scheme with the necessary information, will be implemented in WP2.

MeBeSafe is organised in altogether 6 Work Packages (WPs), as shown in Figure 1.

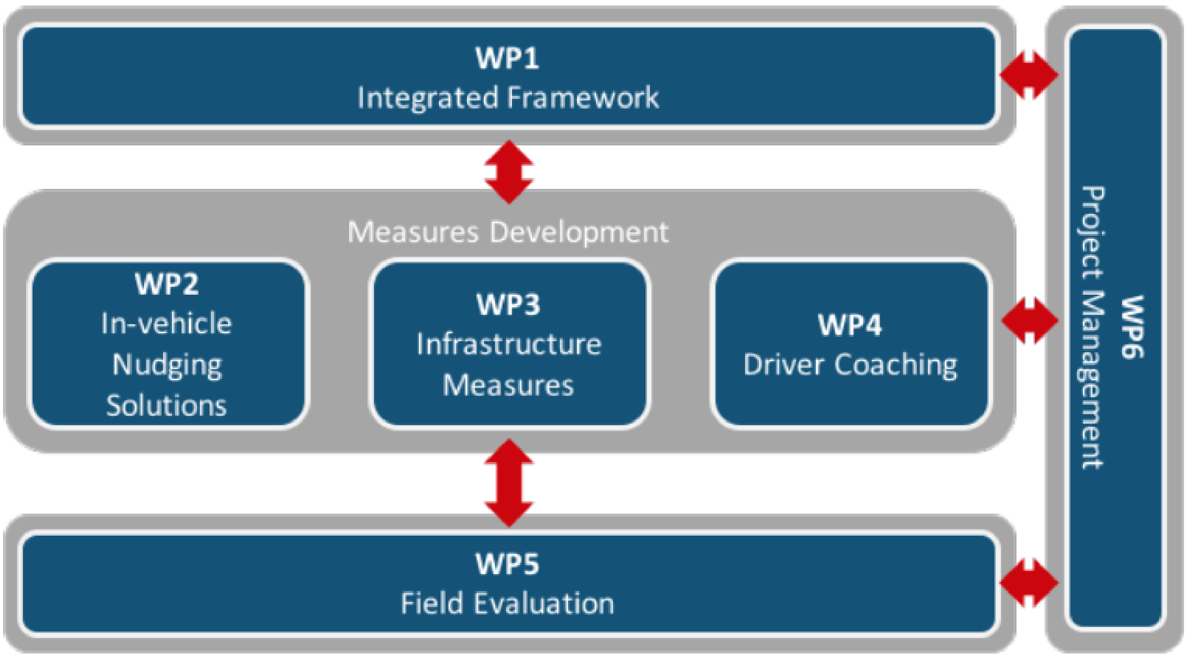

Figure 1 Work Packages in MeBeSafe.

The in-vehicle solutions are mainly dedicated to drivers of passenger cars, with a focus on stimulating the use of safety functions, particularly Adaptive Cruise Control, and directing the driver's attention to potential hazards. In this report (D2.1), different possible solution concepts will be proposed. As development and implementation of a prototype solution into a vehicle (up to a level that the vehicle is allowed to drive on the road) is very costly, an evaluation funnel is developed in T2.5. With the evaluation funnel, consisting of a combination of driving simulator tests and virtual simulation tests, the most promising solution will be selected for implementation.

WP2 targets solutions to meet the following MeBeSafe objectives:

- **O2: Increase the use of Adaptive Cruise Control systems** throughout the journey to prevent close following. Insufficient distance between vehicles in close following are a direct causation in 10% of road accidents [2].
- **O3: Direct the attention of the driver to potential hazards** to increase the timely perception of actual hazards. "Failure to look properly" has been shown to be a major causation factor in 30% of accidents [4].

1.2 Description of Tasks

Task 2.1 – Sensing driver and vehicle state

Interfaces will be defined and implemented with the sensors that provide information on the driver and vehicle state:

- Interface to the driver drowsiness state sensor as implemented in VCC cars including a definition and implementation of data transfer to a VCC coaching app.
- Interface to the driver direction of attention sensor that is used during the tests to evaluate the different solutions to influence the drivers direction of attention.
- Interface to get information regarding the vehicle state such as current speed, heading, and acceleration, and whether ACC is switched on or not.

Task 2.2 – Sensing and predicting cyclist intent

To provide the appropriate nudging information towards the driver in an earlier stage than when a critical situation is imminent, information on the intent of cyclists that might interfere with the path of the host vehicle needs to be available some seconds before the critical situation occurs. This information comes available from the interpretation of the vehicle's sensor system. Based on the view of the bicycle's trajectory over the last few seconds, a prediction is made over the intended trajectory for the coming seconds. The prediction will come with an estimated probability for the cyclist's manoeuvre. To develop such a sensing and prediction system requires the following tasks:

- Develop a probabilistic cyclist's intent prediction model for the most common interaction scenarios between cyclists and passenger cars on a typical intersection.
- Perform an observation study to determine typical bicycle-to-car manoeuvres and to estimate the model parameters.
- Perform a sensor study to determine the accuracy of path prediction of a cyclist, based on in-vehicle sensor observations. In this way, the detections made in the observation study are coupled to the paths as monitored from the car.

Task 2.3 – Hazard perception and prediction

Current Autonomous Emergency Braking systems only brake when a collision with a (cyclist) target is imminent. The AEB decision logic and control law uses the relative position and movement of the target to make this judgement. Nudging responses to the driver need to occur at a larger distance from the (potential) targets, whose positions and intended manoeuvres with respect to the car need to be known (or estimated), considering the local relevant traffic rules and infrastructure layout. The objective of T2.3 is to build a world model based on an available map (public domain), the localization of the host vehicle on this map and the location and intended manoeuvres of surrounding cyclist targets from the vehicle's sensor information.

The world model consists at each point in time of all relevant information for the decision logic and control law to respond with appropriate nudging actions. The development of the world model includes:

- Interfacing to a GPS-based map, using the GPS position of the host vehicle and the corresponding driving direction.
- Sensing possibly hazardous situations from fusion of sensor data (obstructions e.g. by parked cars) and the world map regarding possible traffic crossing the host vehicle path. Relevant information from the UDRIVE project [1] will be utilised.
- Integration with the target intent models to provide a complete picture of potential hazards and probability measures. A simulation application is built as a development tool to carry out stochastic simulations in order to support the development. Its purpose is to determine the difference between actual hazard and perceived hazard.

Task 2.4 – In vehicle nudges

The information on a possibly hazardous situation is input to nudge the driver several seconds before a cyclist is crossing the path of the host vehicle. The nudge is intended to avoid any critical situation. Should a critical situation occur (i.e. a Time-to-Collision of < 2 sec), then the available advanced driver assistance system (ADAS) should be triggered in addition to the nudge. The objective of T2.4 is to design the decision logic and control law of the nudging system according to the framework of WP1. Output from these systems is provided to a human-machine-interface (HMI), which is also developed and implemented in this task.

For the Adaptive Cruise Control (ACC) awareness nudging, the main approach is to present the driver with information on current percentage of ACC usage (by the driver) over a certain prior time period in such a way that usage of ACC is encouraged.

For the intersection conflict nudge, information on possibly hazardous situations will be presented several seconds before they are predicted to escalate (e.g. a cyclist is crossing the path of the host vehicle), in order to nudge the driver towards adaptation of larger safety margins, and hence avoid the potentially critical situations.

The actions to be undertaken within this task are:

- Decision logic and control law development and implementation;
- HMI development for Directing the Driver Attention towards possibly hazardous situations involving crossing bicycles;
- HMI development for ACC Awareness;
- HMI implementation for evaluation in a virtual test environment and for testing in the driving simulators at FIAT Chrysler Automobiles (FCA) and Volvo Car Company (VCC).

Task 2.5 – Solution selection

A task in which different options for in-vehicle nudging solutions as proposed and developed in tasks T2.2, T2.3 and T2.4 will be evaluated in driver simulator tests, and by means of virtual simulations. It is the objective of task T2.5 to select the most promising and feasible option for implementation in the test vehicles for testing in WP5.

Task 2.6 – Implementation of the nudge solution in the test vehicles

Based on the results of task T2.5, the nudging system and corresponding HMI will be implemented into the test vehicles for validation in a field trial (WP5). The vehicles will be prepared to run a field trial. Results out of the field trial will be used to update and optimize the system.

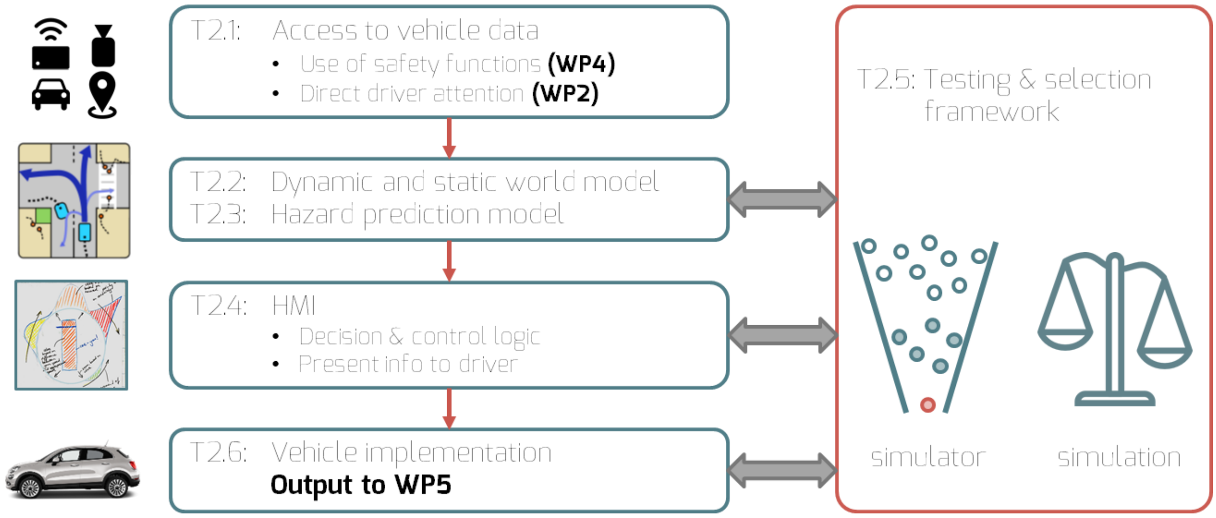

Figure 2 Schematic overview of tasks and their relations in WP2

1.3 Structure of deliverable and contribution by partners

In this deliverable, we focus on concept development of in-vehicle nudging solutions to:

- o Stimulate the use of safety functions onboard vehicles;
- o Direct driver attention to potential hazards.

Chapter 2 provides an introduction to the state-of-the-art of in-vehicle nudging solutions. Background information is provided regarding safety systems that are currently available, to determine the solution space for nudging technology. Basic concepts for nudging are provided, differentiated to the 2 objectives.

VCC focussed on:

- o Increasing the usage of Adaptive Cruise Control (ACC): this requires a signal that shows the current use of ACC in the vehicle. VCC not only makes sure this signal is available from the vehicles that will be used in the tests in WP5. The interaction with VCC's coaching app is elaborated in WP4, and is consequently no part of this document.

- Make drivers take a break when they are really tired: most Volvos are equipped with a driver drowsiness sensor. The data from this sensor are made available and used as input for WP4.

TNO, Cygnify and OFFIS focussed on:

- Solutions that direct driver attention to potential hazards. For such in-vehicle solutions, a distinction is made in the development of a hazard prediction model (Chapter 3) and the development of HMI options to inform the driver on the potential hazards according to the model (Chapter 4).

In Chapter 3, the process towards directing driver attention to potential hazards is further elaborated. As a use case, directing driver attention to potentially hazardous interactions between passenger cars and cyclists in urban areas is considered, from the perspective of the host vehicle (the vehicle in which the in-vehicle solutions will be integrated). A hazard prediction model is proposed and the different flows of information required to come to a hazard prediction are shown. TNO focused mainly on the static world model to describe the conditions that are relevant for hazard prediction (road infrastructure layout, traffic rules, view blocking obstructions). A definition of 'hazard' is given, and it is shown how to determine the hazard based on information regarding the environment (static world model), the host vehicle and the cyclists that possibly interact with the host vehicle.

Cygnify took responsibility to draft a model for predicting the trajectory of cyclists that come into the view of the sensors on board the host vehicle. Based on image processing techniques based on Artificial Intelligence, the manoeuvre of each individual cyclist that has a possible interaction with the host vehicle is predicted a few seconds ahead. This is not a straightforward extrapolation of current manoeuvring. The model includes the use of features (such as hand gestures) to enhance the confidence in the predicted path. TNO and Cygnify integrated the different models into one hazard prediction model that covers both the static as well as the dynamic world model.

In Chapter 4, different options for the human-machine-interface (HMI) are described by OFFIS. Options have been generated in a brainstorming workshop with all partners involved in WP2 following OFFIS's Konect method. Based on a potential hazard level, an HMI is used to nudge the driver in adapting her/his behaviour to appropriately deal with this hazard. The HMI options mainly provide information regarding the hazard level and the direction from which the hazard might be expected in an intuitive way towards the driver. Each of the three main options is elaborated and possible ways for implementation are given.

Chapter 5 of the deliverable provides the conclusions as input to task T2.5.

2 In-vehicle nudging solutions

This chapter provides an introduction to the state-of-the-art of in-vehicle solutions. Background information is presented on different ideas for in-vehicle nudging of drivers that can be implemented in addition to already existing active systems for the enhancement of safety, such as ADAS.

2.1 Stimulating the use of safety functions onboard vehicles

MeBeSafe aims to explore risk management areas in new ways. The "new" part refers to the fact that while most current measures to increase road safety by means of risk management try to make traffic users behave in specific safer ways by explicitly appealing to reason, MeBeSafe is based on the realization that driver behaviour is largely automated and habitual, with reason playing a limited role. The project therefore takes another approach by aiming at changing habitual traffic behaviour using nudging. This is a concept adapted from behavioural economics that relates to (subconsciously) pushing humans in a desired direction without being prohibitive against alternative choices of action.

A particular application related to road safety is the fact that several driver assistance functions, which do have a documented safety impact potential, remain underused by the drivers, thus not realizing their full safety potential. In MeBeSafe, nudging will be applied to stimulate drivers toward higher usage of two such in-vehicle safety systems. The first is to nudge people to use Adaptive Cruise Control (ACC[1]) more, and the second is to nudge people toward higher compliance with Drowsiness Alerts, leveraging the power of incentives.

While nudging holds a lot of promise, it remains a general concept that needs to be made concrete on a use case basis to become meaningful. In the following two

[1] A cruise control function that automatically adapts its speed to the speed of a slower lead vehicle in order to keep a set minimum distance to this vehicle.

sections, concrete traffic safety use cases where nudging through an in-vehicle driver interface is expected to prove successful, will be presented.

2.1.1 Increase the usage of Adaptive Cruise Control

In order to crash into a lead vehicle (have a rear end collision), accident causation research shows that generally two things are required:

- A distracted driver, and…
- … a lead vehicle.

It may seem needless to bring up the second condition, but in fact, this is very important. The reason is that many proposed solutions to this problem only focus on the first issue (the distracted driver), and do not address the second one (close following). However, research on this conflict type clearly shows that the risk of crashing is highly influenced by how far behind you are when the unexpected happens [3]. If you are not following a lead vehicle very closely, you are much more able to resolve the conflict once it arises. Rear end crashes can therefore most likely be addressed just as well by avoiding close following as by avoiding distracted drivers.

Given this assumption, the next question is how we can make drivers avoid close following. Given that driving is largely automatized and habitual, changing how a particular individual is managing his/her distance keeping is indeed challenging. However, there exists a simpler approach, illustrated by Figure 3 Frequency of short time headway events per 100 km of driving below. The figure captures a result from euroFOT [1], where driving with and without ACC active in vehicle following situations was compared (Baseline=manual driving, Treatment=ACC active).

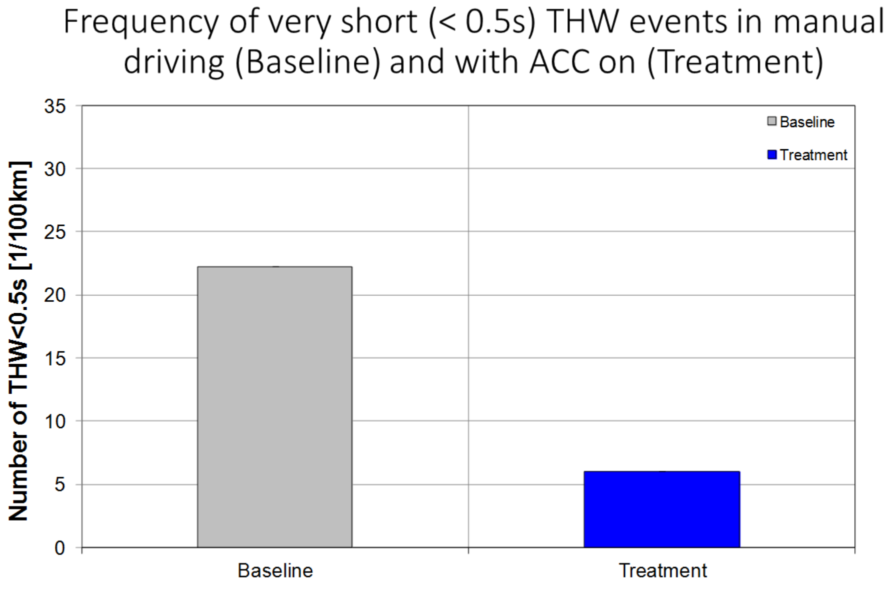

Figure 3 Frequency of short time headway events per 100 km of driving

As is evident from Figure 3, the frequency of close following events, i.e. events that satisfy the risk mechanism for rear end crashes as described above, is much lower in Treatment, i.e. while ACC is active. ACC usage thus lowers the risk of rear end crash involvement.

With the merits of ACC usage illustrated, the next question is: how to make drivers use ACC more in their everyday driving? Several approaches are possible. One could simplify activation (e.g. single button press activation), implement an Opt-out rather than Opt-in activation (e.g. let ACC start automatically in selected lead vehicle following situations), leverage social norms (campaigns in the style of "Don't drink and drive", e.g. "Don't tailgate") and give incentives or rewards, similar to collecting miles with airline reward programs.

In the current approach within MeBeSafe, another mechanism will be employed. This mechanism consists of feedback on ACC usage over time, displayed on an extra display mounted next to the centre console[2]. The display content is modelled on ECO

[2] For ease of implementation in a large number of vehicles in WP5, a separate display is proposed, to avoid large efforts in implementation in the centre console or instrument panel of the vehicle.

driving coaching[3]. In its simplest form, the display would indicate what the driver's current ACC usage score is, and how much more ACC time is required to get in the "green zone", where a certain level of ACC usage is deemed appropriate, given the context and driving patterns. The basic idea is that everything else being equal, drivers generally prefer to stay in the "green" zone if that can be achieved with limited effort, and this mechanism is what the in-vehicle display is designed to leverage for nudging.

2.1.2 Make drivers take a break when they are really tired

Drowsiness is a large traffic safety problem (4). It does not receive quite the same level of attention as other crash contributing factors, probably because it is hard to detect and assert in retrospect, which is when most traffic accident reporting is done. However, it is safe to say that if we could address drowsy driving properly, traffic safety would be significantly increased.

What is interesting about drowsiness is that when it comes to detecting that a driver is drowsy, this problem has been technically solved to a large extent (at least for manual driving). For example, VCCs' Driver Alert system has close to 100% detection rate of when drivers are about to go into a micro sleep that is long enough to leave the lane completely.

The drowsiness problem is thus not primarily a technical problem. Rather it is a behavioural one; few drivers actually take the break needed when drowsiness is detected. From VCC interviews with tired drivers, it is clear that the reasons for not taking that break are many and varied; you want to get home, there's no good place to stop, you only have a few minutes left to drive, etc. However, it is also clear that on a meta level, all these reasons have one thing in common, and that is an unwillingness to change the current course of action before it is completed. In other words, when you are really tired, you really do not want to exchange your current

[3] VCC's own coaching app.

plan of action for something else. Drowsy people simply put are neither flexible nor willing to re-prioritize.

Given that this analysis is correct, the nudging problem for drowsy drivers becomes clear. For cars with Driver Alert, the default setting is that it is always on and hard, if not impossible, to switch off. The driver thus does not have to do anything to get the feedback from the system on whether s/he is drowsy or not. All that is left is nudging the driver to take that break or switch drivers when the system indicates drowsiness, instead of alerting her/him in the traditional way with audible and/or visual warnings.

So, what can be done? Again, as shown in the table below, there are several ways to nudge people into taking that break they really need:

1	Simplify usage	Automatic guidance to rest spot / gas station, etc.
2	Increased performance awareness I	Continuous feedback on drowsiness level rather than on/off state
3	Increased performance awareness II	Also display current level of lane keeping/positioning performance over time. This will co-vary with the drowsiness level, hence making the impact on driving of drowsiness more apparent.
4	Leverage social norms	Paraphrase "don't drink and drive" → "Don't sleep and drive"
5	Incentives/Rewards	If you take a break you get a free coffee. This offer is valid within the next 10 minutes.
6	Gamification	Every time you stop after a Driver Alert warning, you get one lottery ticket. People with lottery tickets are then eligible for the quarterly raffle where the winner gets a brand-new car.

Table 1 Possible approaches to make drivers take a break when drowsy

Of these, no. 5 is the most promising candidate for MeBeSafe, and thus the one which will be implemented in the VVCs test fleet. The challenge here is that the incentive itself has to trigger some other experience and/or expectation in the driver then that of just receiving an economic reward (most people who can afford to drive also can afford a cup of coffee). By offering something for free, and within a limited time, one introduces two additional dimensions. Getting something for free usually draws a lot of interest, and setting a time limit on the offer introduces an element of competition as well. A series of pilot tests will determine if the incentive needs to be implemented as 'immediately available', or if a delayed reward also would be acceptable in terms of impact on behaviour.

The options regarding the use of safety functions onboard vehicles, incl. the function that indicates a driver to take a break when the system senses a certain level of drowsiness of the driver, are further elaborated in WP4 (Driver Coaching).

2.2 Directing driver attention to potential hazards

In urban traffic, hazards can come from every possible direction, since many road users, including cyclists and pedestrians, exhibit independent manoeuvrability resulting in trajectories that potentially interact with approaching cars. Drivers are having difficulty in predicting episodes of increased risk; they are especially not used to paying attention to potential hazards in pre-conflict situations. The nudging-coaching framework [6] has been used to map out the characteristics of the addressed problem (see Figure 4). It was concluded that:

- Drivers' inattention to possible risks is an unintentional error and that drivers generally are willing to direct attention to possible risks in order to avoid traffic accidents.
- However, the driving context does not support drivers to direct attention to possible risks; mostly because possible risks even more than actual risks are difficult to recognize.

This combination of factors shows according to the nudging-coaching framework that nudging is expected to be a successful intervention strategy to direct the attention of drivers towards potential hazards. Tentative insights from WP1 suggest that the design of the nudge should be beneficial to all. WP1 recommends to design the nudging intervention in such a way, that the drivers cognitive load is considered, especially in demanding situations.

The current state-of-the-art of active safety systems, Forward Collision Warning (FCW) and Autonomous Emergency Braking (AEB), have been widely introduced since 2014. A car equipped with FCW/AEB makes use of on-board sensors such as camera and radar, to track and trace traffic participants that possibly interfere with the trajectory of the host car. This information is used to warn the driver in case of a possibly critical situation and/or to brake in case the driver does not respond and the risk of collision does not decrease. Such FCW/AEB systems aim to support drivers that either are not aware of or do not recognize a potential hazard. One or more on-board sensors continuously scan the area in front of the car, to identify the type of traffic participant (as car, truck, pedestrian, cyclists, motor cyclist) and to determine the speed and heading of each potential target object that is in the field-of-view of the sensors. Based on these measurements, considering the vehicle's own speed and heading, a Time-to-Collision (TTC) is calculated for each object.

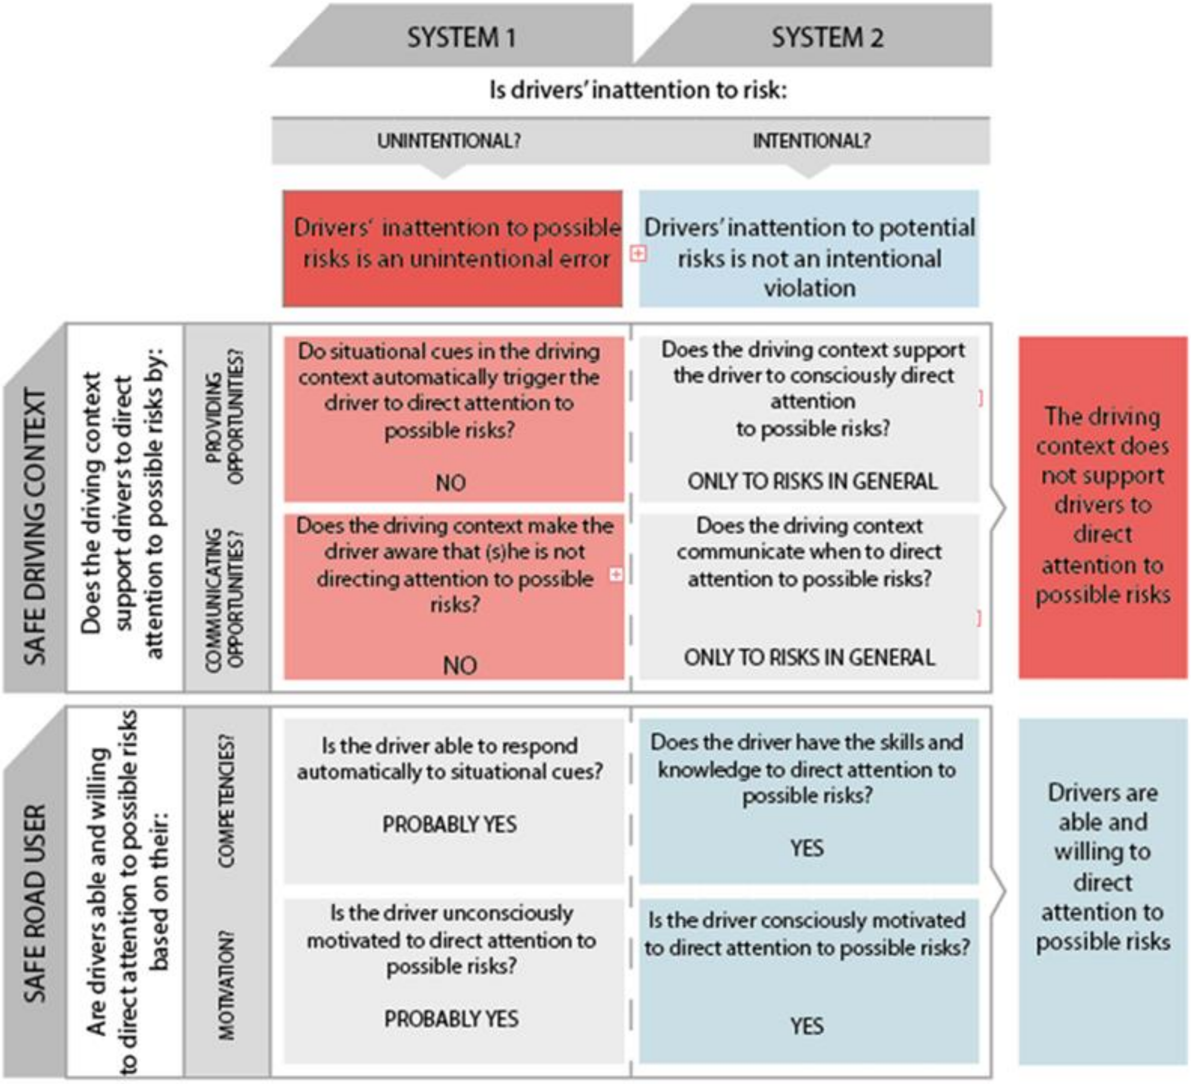

Figure 4 Overview of design considerations and tentative conclusions drawn in Stage 1 addressing drivers' inattention to risks

The TTC is defined such, that a collision with the object would occur as soon as the TTC equals zero. Most AEB systems will warn the driver, in case the TTC for any traffic participant is approximately between 1 and 2-3 seconds. Such a warning can be an audio warning (strong beep or sequence of beeps), a visual warning (a flashing red or orange symbol on the dashboard or head-up display), a haptic warning (vibration of the steering wheel or gas pedal, or a very short brake pulse of the car) or a combination of those. The issued warning intends to provoke a driver response to avoid the collision by braking or steering. In case the TTC continues to decrease and the driver fails to respond, the AEB system will initiate full braking of the car at a TTC of between 0.5 and 1 second before impact. Forceful braking is applied in an attempt to avoid the collision, or at least to mitigate the collision by reducing impact speed. The typical timeline sketched here, shows that the AEB response is usually limited to a maximum of 2-3 seconds.

Cyclists pose a particular large problem in traffic due to their high speed compared to pedestrians and their large manoeuverability compared to motorized vehicles. The study into a test protocol for an AEB-system that is capable to avoid or mitigate a collision with cyclists, revealed limitations to system performance due to the specific cyclist behaviour related to their speed and manoeuverability [27]. To develop this test protocol [8], typical scenarios of car-to-cyclist accidents as collected in the German In-Depth Accident Database GIDAS [7] have been studied. A typical situation where drivers easily overlook the potential hazard is given by a near-side cyclist crossing with a view-blocking obstruction. Cyclist crossing scenarios represent more than 50% of all car-to-cyclist fatalities and seriously injured, whereas a view-blocking obstruction (building, fence, vegetation, parked cars) is reported for between 25% and 50% of these cases [21]. A driver might easily overlook the potential hazard in these cases as:

- The cyclist coming from the right (EU-mainland driving directions) is not seen because of the presence of the view-blocking obstruction, but often has right of way;
- The situation is not recognized as hazardous as the priority road from the right is not easily visible to the driver.

Current state-of-the-art AEB systems have difficulty in supporting the driver in such cases, as the cyclist comes from behind the view-blocking obstruction into the sensor field-of-view at a short time before a possible collision.

Another reason for a late (close to TTC=0) AEB response is the result of the fact that the system needs to make a prediction whether or not a collision will take place. Consider for instance a situation where a car and bicyclist cross paths and where a collision will take place within 2 seconds in case the heading and speed of the collision partners does not change in these 2 last seconds before impact. Not only the driver might control the car to avoid a collision, also the bicyclist might make a turn, steer away or brake. Cyclists have many options to avoid a collision, especially since they are able to brake up to 7 m/s^2, which allows them to come to a standstill down from 20 km/h within less than a second and little over 2 meters displacement. To avoid that the AEB system is fully braking in a case where the cyclist is able to avoid the collision, usually AEB systems will wait to issue the braking action until it is very certain that the cyclist is no longer able to avoid the collision. Based on the above, a braking action to avoid a collision with a cyclist will often be issued less than 1 second before the impact would happen. This limits the speed reduction that the car is able to achieve upon the time of impact. Current state-of-the-art systems (developed to be introduced already in 2018) will consequently show a significant speed reduction strongly mitigating a collision, but full collision avoidance will be difficult to achieve for higher car speeds (50 – 60 km/h).

Simply issuing warning and braking activation earlier at higher TTC is no solution, as the Cyclist-AEB system would then often warn and brake unnecessarily in densely populated areas with many cyclists (easily more than 10 false positive responses per hour). This would incontestably lead to devaluation of the warning signal and to driver annoyance. Ultimately, the driver might not accept the ADAS functionality and switch off the AEB system, in which case no support is provided even in cases where it is necessary.

The typical operation of FCW/AEB systems is given in Figure 5. It shows the escalation level of the system as function of TTC. In MeBeSafe, we will add a nudging system in addition to available FCW/AEB systems, without the intention to alter such FCW/AEB system. Nudging will be provided as a continuous low-level input (depending on the prediction of the level and direction of the possible hazard) to the driver and it is important that this input is consistent with the escalation of the FCW/AEB system in case an actual hazard is perceived.

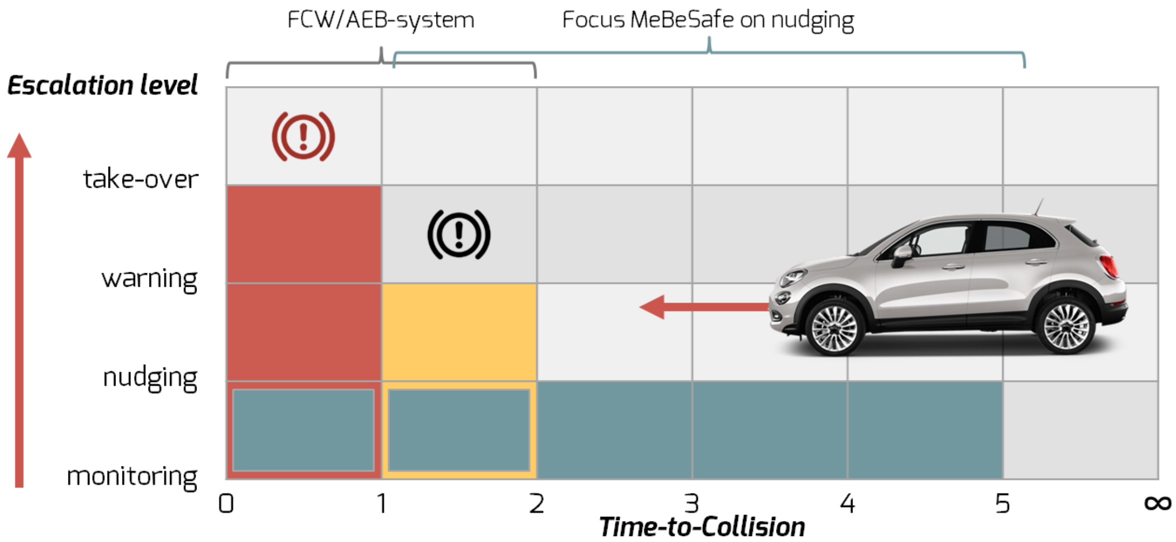

Figure 5 Schematic view on the escalation level of a driver assistance system based on the Time-to-Collision.

The nudging measures developed in MeBeSafe intend to direct driver attention towards the potentially hazardous situation, much more than current indistinctive beeps or warning lights provided by FCW/AEB systems. Options that are considered include:

- An augmented reality projection of a safe zone for driving as overlay in the head-up display (HUD). In colours the preferred driving zone and the potentially hazardous areas are indicated.
- Providing moving icons of bicycles in the HUD at the location from which cyclists might be expected. Actually present cyclists might be highlighted with a colour depending on the potential level of hazard that they represent.
- A projection of a circle in the HUD in which deviations in shape and colour indicate the direction and level of the potential hazard in an intuitive way.
- Each option comes with a way of nudging the driver towards a car speed that is considered appropriate for the situation.

In T2.6, the most feasible and effective option, selected by the procedure developed in T2.5, will be implemented into one FIAT 500X. In WP5, this vehicle is used in a field-operational-test for a final evaluation and estimation of the effectiveness of the selected solution in increasing road safety.

3 World modelling and hazard prediction

According to ISO 45001 on occupational health and safety [7], a *hazard is a source with a potential to cause injury*. In the context of this paper, the hazard is a possible collision between a passenger car and another traffic participant such as a cyclist. The risk is defined as the probability that the hazard actually causes harm. Where a hazard refers to an existing situation (a car and a cyclist approach each other at an intersection), the risk refers to an anticipated situation with a certain likelihood of happening (the probability that the car and the cyclist actually collide with each other).

AEB systems aim to eliminate hazards or mitigate the harm caused in highly risky situations, i.e. critical pre-crash situations in which the probability of collision is very high. Nudging systems can provide interventions even before situations become critical; they provide information to the driver, also in cases where the probability of a critical situation evolving is not close to 100%. In such cases, the intervention (information) can be provided much earlier in the process, over a time horizon of typically 5 seconds (the timing of the interventions needs to be tuned based on simulation and/or field studies), under the condition that the driver can easily digest and use the nudging intervention without being forced to act or respond to it. Should the situation become critical anyway, then the present FCW/AEB system will come into action and escalation will follow the control and decision logic of this system.

Nudging has the advantage that information can be provided even in cases with a lower probability of a hazard actually evolving into a critical situation or for lower probability levels of a hazard actually being present. The reason for this is that nudging is less likely to anoy the driver, resulting in switching-off the ADAS. There are many factors that influence hazards and the exposure to hazards. The following figure provides different examples of similar situations, with different influence factors that determine the level of perceived, potential and actual hazard and the associated risk:

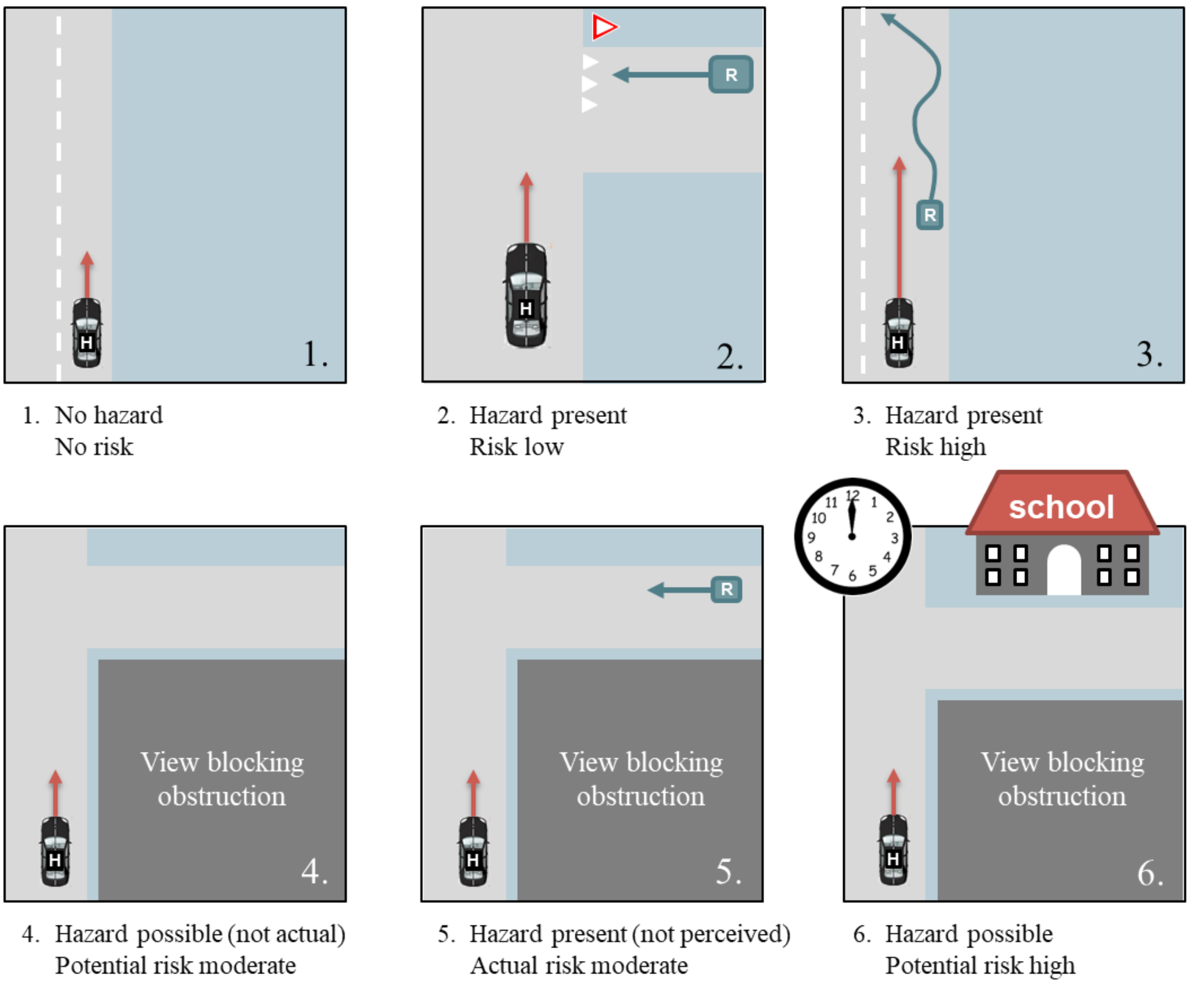

Figure 6 Different situations with an indication of the actual and potential hazard and the actual and potential risk. In this case the hazard is a collision with another road user from the perspective of the host vehicle (H).

Figure 6 shows examples of factors and circumstances that influence the perception of hazard and the associated risk from the perspective of the host vehicle:

- The actual presence of road users that exhibit manoeuvres that potentially conflict with the manoeuvre of the host vehicle. The risk of collision depends not only on the manoeuvres itself (existing situation), but also on the anticipation of the road users towards the hazard (e.g. a swerving bicyclist) and the locally applicable traffic rules;
- View-blocking obstructions (e.g. building blocks, hedges, and parked cars or trucks) that limit or fully block the view on possibly approaching road users which hampers anticipation and increases the potential risk;

- o Similarly, the presence of a school increases the number of potential hazards depending on the time of day – e.g. just before classes start or when classes are adjourned.

To provide a prediction of the potential hazard (and the associated risk) to the nudging HMI (HUD), the following scheme is followed:

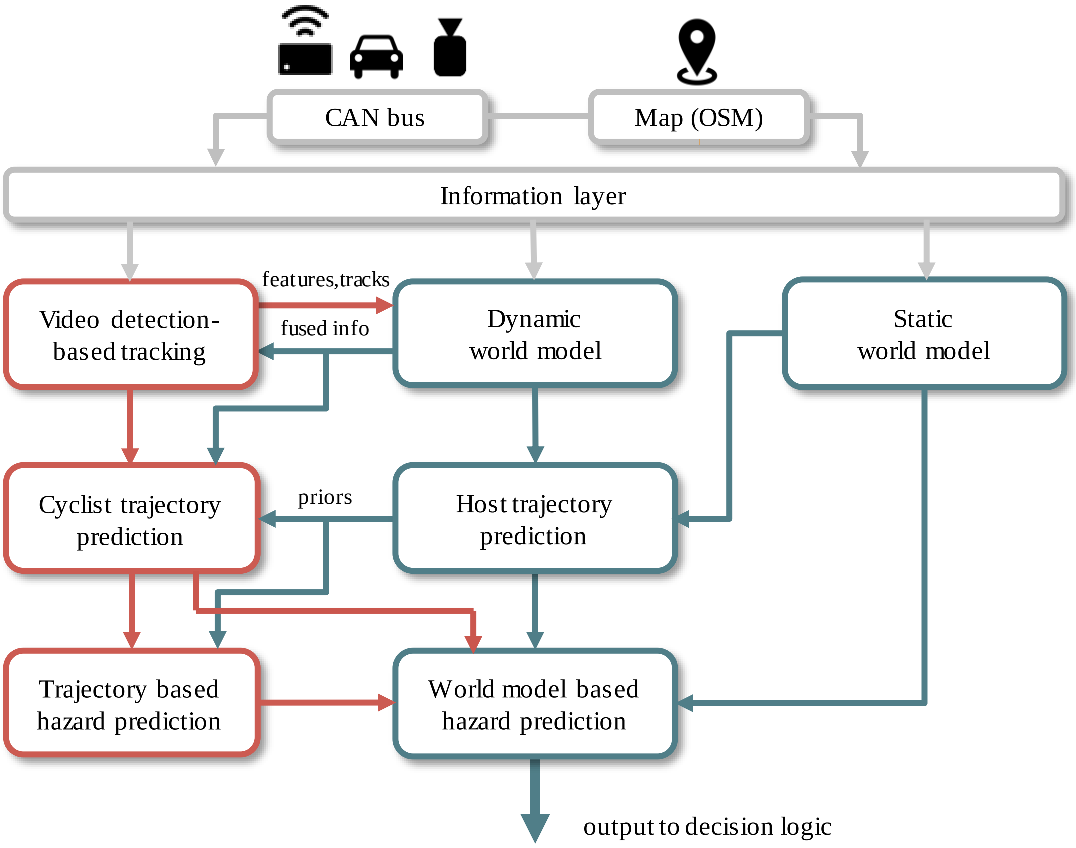

Figure 7 Schematic of information flow from the vehicle sensor system to the predicted hazard as input to the decision logic of the nudging system and HMI

The vehicle sensors (including at least forward-looking radar and video camera) provide input on the current situation with state variables (position, heading, speed) regarding the most important objects surrounding the host vehicle and the host vehicle itself. A connection is made to open street map (OSM). Based on the GPS position, the location of the host is indicated on the map. The map is assumed to provide information on the applicable traffic rules and the presence of the permanent view-blocking obstructions.

The hazard prediction model provides an indication of the most important detected and potential hazards, with a prediction of the risk associated with each of these distinct hazards.

3.1 Dynamic and Static World Model

Within the dynamic world model, the object level data provided by each of the sensors is fused into a congruent world image for the host vehicle. This is a standard component in current ADAS functions as well. The dynamic world model is based on the road users that are actually within the field-of-view of the vehicle's sensor system and that are detected. Object detections by the individual sensors are combined to provide the state of those surrounding road users that are most important for the host vehicle. The state includes the type of road user, the relative distance, the speed and heading of the road user. Based on variables such as TTC and lane occupancy (in case the sensors also provide lane and line marker information), the importance of each object for the host vehicle is rated.

The hazard prediction model proposed in this paper goes beyond the limitation of the dynamic world model: the latter estimates only the actually detected hazards, but not the potential hazards. The hazard prediction model also comprehends a static world model: this static world model provides additional information that is used by the hazard prediction model to determine a qualitative indication of the expected risk. The static world model considers the local road layout (e.g. the presence of view-blocking obstructions), combined with the observed traffic flows depending on the time of day and day of the week at the intersection. An example of such a cyclist accidents heat map for Budapest (2011-2013) is given in [8]. While the available information for local road locations may differ from place to place and may not always be so extensive, the idea is that *where* such data is available, it can and will be used and help to improve the static world model.

The outputs of the dynamic world model (states of the surrounding traffic participants) and the static world model (priors of the static environment that affect the level of hazards and potential risks) are combined with a prediction of the host trajectory (expected manoeuvre over the intersection by the host vehicle) to determine a qualitative prediction of the hazards in crossing the intersection in the *world model based hazard prediction*-module. These predictions indicate the relative position of different hazards and the associated risk with each of the hazards.

3.2 Cyclist Trajectory Prediction

The manoeuvrability of cyclists is particularly high, due to their typical speeds (between 15 and 20 km/h nominal speed [9]) and their capability of sharp turning (changing direction) and heavy braking [10][11]. Simulation studies have shown that cyclists travelling at 15 km/h crossing the path of an approaching passenger car can avoid a collision at a TTC as low as 0.8 sec by an extreme braking action of up to 7 m/s^2 [6].

Because of this high manoeuvrability, the current state of cyclists according to the host vehicle's dynamic world model, as well as direct linear extrapolation of its state (e.g. derived from Kalman filters), is a poor prediction for the state that evolves over the next 2 to 5 seconds, the time interval that was taken as typical for nudging. Consequently, the risk evolution predicted from a straightforward dynamic world model that does not take these peculiarities into account has limited value for hazards that involve cyclists. It is for this reason that models focussing in particular on the prediction of the cyclist trajectory are developed.

The input for such a model consists of the observation of the trajectory of the cyclist (position, speed, heading) over the last seconds, enriched with features that provide indications of the cyclist's intent: hand gestures, turning of the head, roll movement of the bicycle, pedalling behaviour of the cyclists, combined with an approach inspired on "social forces" modelling [12],[20], which models the interaction between the individual cyclist of interest to other road users and to road infrastructure elements.

The whole prediction method consists of a stepwise process which makes use of modern deep learning-based Artificial Intelligence (AI) approaches for detection, tracking and prediction.

Data sources: As the most important data sources video (taken from a forward-looking camera mounted at windshield level) is used as input for trajectory prediction, next to radar data; and map data are coupled and exploited as well to get information on the road infrastructure and its constraints (see Figure 7). Data on bicyclists and other road users derived from video is fused with radar data for better position and velocity estimates of the road users.

Detection: The stepwise process starts with a deep learning-based computer vision algorithm for object *detection*, so we can confidently distinguish between cars, trucks, buses, motorcycles, pedestrians, and bicycles. See Figure 8 for an example. We use both YOLOv2 [13] and Faster-RCNN [14] for this purpose. Note that we need to identify, track and perform elementary prediction of road users other than bicyclists as well (i.e. cars, trucks, buses, motorcycles, and pedestrians) in order to, first of all, distinguish the bicyclists from other road users and, second, to model their interactions with each other.

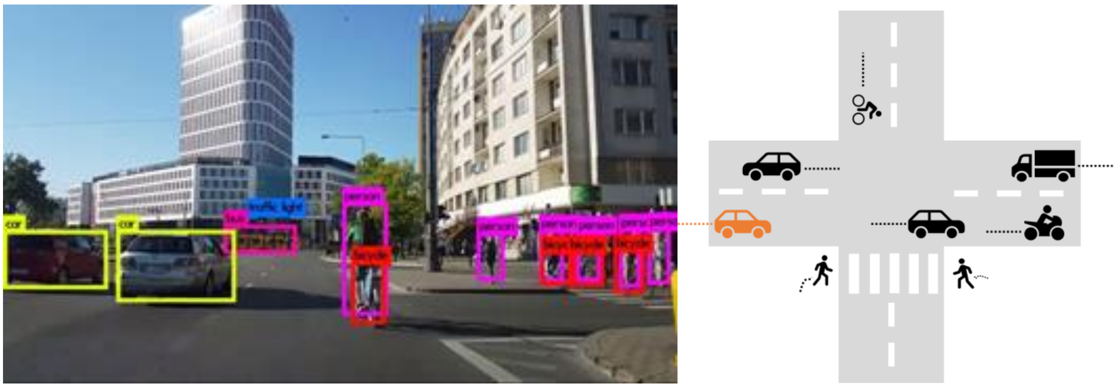

Figure 8 A (left): Deep learning-based road user detection in an image. Note the accurate identification of road user class (such as car vs. bicycle vs. pedestrian), and accurate localization in the image. B (right): Cartesian space trajectories of the host vehicle (orange) and other, detected and tracked road users (black), relative to the road infrastructure.

Tracking: Tracking is the process of following the individual detected road users over time (their trajectories); ideally even when there are large appearance changes and occlusions which make such tracking difficult. In particular, we use *multi-object* (or multi-target) tracking (e.g. [15][16]). This process is done such that a spatio-temporal representation can be obtained of the trajectories of the multiple other road users around our driver – which describes the traffic situation around our driver and which gives us the recent trajectory of the road users over the last seconds, a necessary ingredient for the later step of predicting for the next few seconds. The tracking method relies heavily on accurate detections from the object detection algorithm and on so-called "appearance feature vectors" derived from those detections [17], which act as a kind of "fingerprint". They remain relatively constant for subsequent detections of the same object over time in different image frames (e.g. the same male bicyclist with green shirt and black backpack) while distinguishing clearly between different instances of the same object class (e.g. a male bicyclist with green shirt and a female bicyclist with a black coat). In this way, by storing past appearance vectors for different "active" tracks, and comparing newly detected appearance vectors to stored vectors, we can solve the so-called association problem and assign new detections to existing active (or new) tracks. With new image frames and detections, the storage of appearance vectors associated with tracks is updated as well, to account for appearance changes over time.

Transform to cartesian space: The tracked trajectory in the image space (i.e. sequences of bounding box coordinates relative to image pixels) is translated to the cartesian real-world space (i.e. X,Y coordinates with respect to the road infrastructure). In this step, we also perform data fusion with radar data that provides detections of (many of) the same objects. A process which after data association gives better relative position and velocity estimates for our detected and tracked road users. Combined with position, velocity, and acceleration data from the host vehicle's CAN and GPS systems, this gives us the information to position the trajectories of our host vehicle and the most important traffic participants around the host into a local fine-grained map, i.e. relative to the particular road infrastructure at that point. Figure 8 illustrates schematically the type of cartesian space representation that we obtain of the trajectories of the host vehicle (in orange) and other road users (in black) relative to the infrastructure.

The main reason to transform to the cartesian space and relate the coordinates directly to a fine-grained map is that we can subsequently, for prediction, exploit knowledge about traffic rules and typical traffic patterns which are coupled to the map (e.g. to the particular intersection), as described above in the discussion of the static and dynamic world model and hazard prediction model.

Prediction: A common approach to predict road user trajectories is to consider road users as individual, monolithic objects or points, and base the prediction on their recent trajectories, e.g. based on individual Kalman filters that extrapolate trajectories of each road user into the near future, basically by assuming linear extrapolations of the observed recent acceleration, velocity, and positions. Often, some type of possible collision course estimation is used to deal with possible interactions between road users.

To enhance the prediction of road user trajectories we enrich the object information by adding additional appearance information derived from video. For bicyclists, this is *head and body pose information* which can be derived from image patches corresponding to detected bicyclist objects with remarkable accuracy using modern deep learning-based methods [18]. From that we derive estimates of where the road user is looking, whether the legs are moving or not, and whether the bicyclist sticks out his arm to indicate intended direction. See Figure 9 for an illustration of this object information enriching approach, on consecutive image frames of bicycle video data from the Netherlands, with our pose estimates applied to the images.

We subsequently exploit knowledge about *traffic rules* and *typical traffic patterns* coupled to the map (e.g. to the particular intersection). Statistics about the most common bicyclist traffic patterns in combination with the locally applicable traffic rules, that we get from the static world model, have been shown to improve forecasting [19].

Figure 9 Consecutive image frames (left to right, with some frames omitted) of bicycle video data from Utrecht, The Netherlands (video source: bicycledutch.wordpress.com), with our pose estimates superimposed on the images.

We take into account the notion that road users are constantly aware of each other and take each other (and static obstacles) into account. In particular, we use a *social forces* approach [14], which is based on the idea that social agents (in our case, road users) naturally "repel" each other in a way that can be modelled as forces. This idea has recently been modernized using deep learning, by means of the social Long short-term memory (LSTM) approach [20]. In addition to the standard social LSTM approach, we include the appearance and head/body pose information as described above, and by including infrastructure constraints (as in [19]) and infrastructure statistics (see section 3.1).

Related to the distinctive elements in our approach described above, and in contrast to most trajectory prediction approaches, we use a *multi-level* approach to predict trajectories. That is, we consider the problem of predicting the "strategic high-level" trajectory choice that a bicyclist makes (such as either going straight, or turning left, or turning right on an intersection) as separate from predicting the "operational low-level" trajectory. The low-level trajectory corresponds to individual high-frequency (e.g. 10 Hz) cartesian coordinates (and velocity and acceleration) predictions for the next several seconds *given a particular destination choice*. Figure 10 shows an example. Specifically, the high-level destination choice prediction is based in part on the map infrastructure constraints and statistics but not on social forces. On the other hand, the low-level position prediction is based in particular on the social forces of other road users. In other words, other road users do not determine where a particular road user will go to in the end (its destination), but they do influence his or her particular low-level position trajectory to get there. The appearance and pose information derived from videos is used in both the high-level and the low-level trajectory prediction for that particular road user.

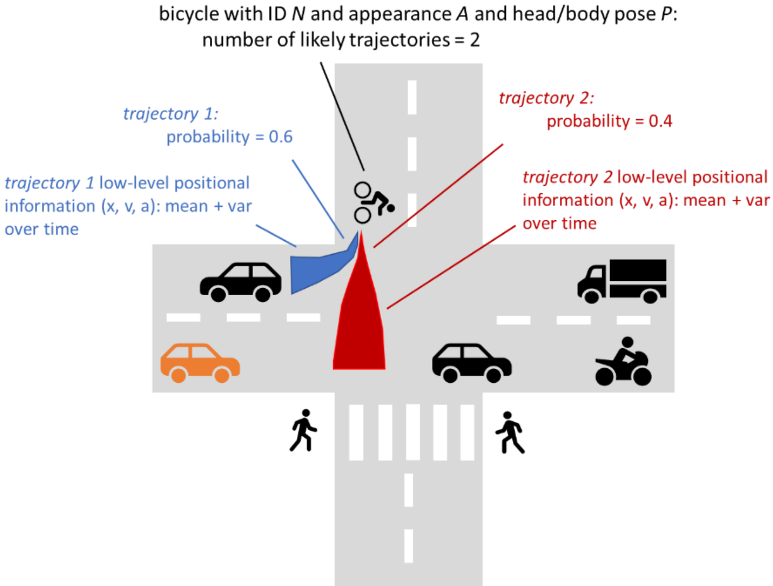

Figure 10 Example of how we represent and predict trajectories at both a high level (the "destination") and a low level position and velocity corridor for each destination.

In this example, we predict that the bicyclist turns right with a probability of 0.6, and goes straight with a probability of 0.4. For each of these options, we predict high-frequency position estimates for the next few seconds, each with a mean and a variance. Typically, over time (further into the future) the variance increases, as shown by widening of the red cone shape in Fig. 3.5 of the low-level trajectory prediction estimates.

The resulting prediction system produces high-level probabilistic predictions of cyclists' destination choices, translated into low level trajectories in a cartesian coordinate system, with a variance associated with trajectories due to model-derived uncertainties. The resulting cyclist trajectory prediction model constitutes a near-real time prediction system: 5 to 10 Hz prediction updates with a latency of 200 to 400 ms appear feasible. With each new data point the prediction is continually updated to reflect the latest sensor information.

3.3 Observation study to provide input for hazard prediction

The in-vehicle nudging system that is developed in the MeBeSafe project will be implemented in a production vehicle (FIAT) that is equipped with a commercial AEB system incl. a standard state-of-the-art sensor set. This sensor set is capable to identify and track passenger cars, trucks, motor cycles, pedestrians and cyclists, when these dynamic objects come into the field-of-view of the sensor system. Consequently, a dynamic world model is already on-board the vehicle for sensor fusion; the outputs of the different sensors are merged into one congruent view on the host vehicle's surroundings by the sensor fusion algorithm. This means that the dynamic world model does not need to be developed from scratch in MeBeSafe. The output of the dynamic world model (the object classification IDs, relative positions, headings and speeds) is assumed to be available for any of the modules developed within the project.

To build up a static world model, an observation study has been performed on one specific urban intersection in the city of Eindhoven at the intersection of Sint Trudostraat/Karolingersweg and Hastelweg. This is a busy uncontrolled 4-armed intersection.

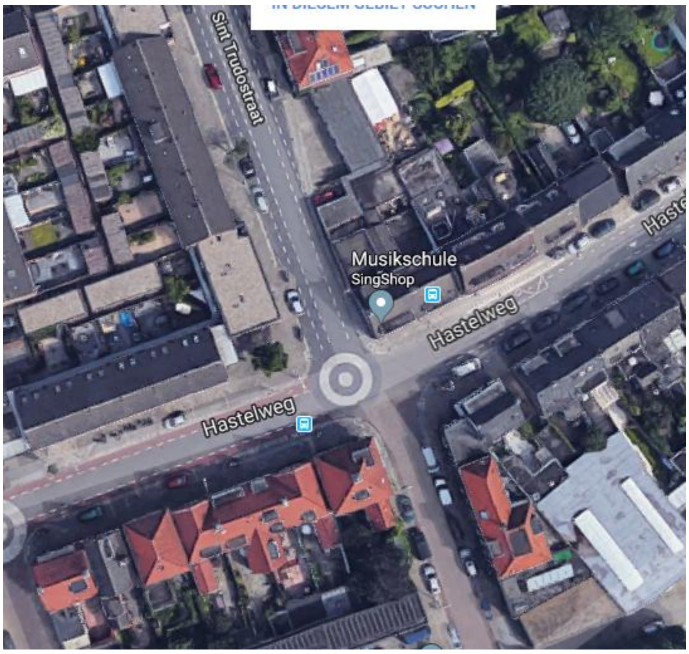

Figure 11 Google earth view on the selected intersection

At two corners, the view is permanently obstructed by a house, only separated from the road by a pedestrian side walk. This view blocking obstruction is rather challenging as parked cars (during a large period during the day) force cyclists to drive closer to the middle of the road. The legal speed limit is 30 km/h, however, previous studies [21] on the same intersection revealed that most vehicles drive (slightly) faster than that. A very shallow speed bump is found at the centre of the intersection. The road markings clearly indicate a crossing of traffic, but the geometry of the speed bump does not challenge the speed of an approaching car.

The objective of this first observation study was to determine the cyclist flows over time at the intersection during a complete week (24/7 including weekend). To this end, a high-definition camera was mounted high on a lamppost to provide a good overview of the complete intersection. From the video footage, all cyclist movements at the intersection were manually annotated. A total of 11868 cyclist manoeuvres have been counted, distinguished over 12 possible manoeuvres (3 directions -turn left, go straight, turn right- for each of the 4 starting points at the intersection). Moreover, it was counted how many cyclists passed the intersection per hour, providing a flow in a number of cyclists/minute. The results are collected in Figure 12.

From the cyclist flow and the layout of the intersection, a relatively detailed static world model can be constructed for this specific intersection. For each of 4 possible approaches of the intersection by a host vehicle, the static world model can provide a prediction of the presence of a cyclist and its possible manoeuvre over the intersection depending on the day of the week and the hour of the day. Based on the host manoeuvre, an estimate can be made for the potential hazard that such a manoeuvre poses onto the host, based on possibility for interaction between host manoeuvre and cyclist manoeuvre. Such a model will provide a hazard estimation when the cyclist cannot be seen from the perspective of the host vehicle, e.g. due to the view-blocking obstruction. In case a cyclist is present in the field-of-view of the host vehicle, the dynamic world model for this specific cyclist computes its hazard separately, making use of the model described in Section 0, next to the hazard prediction of the static model.

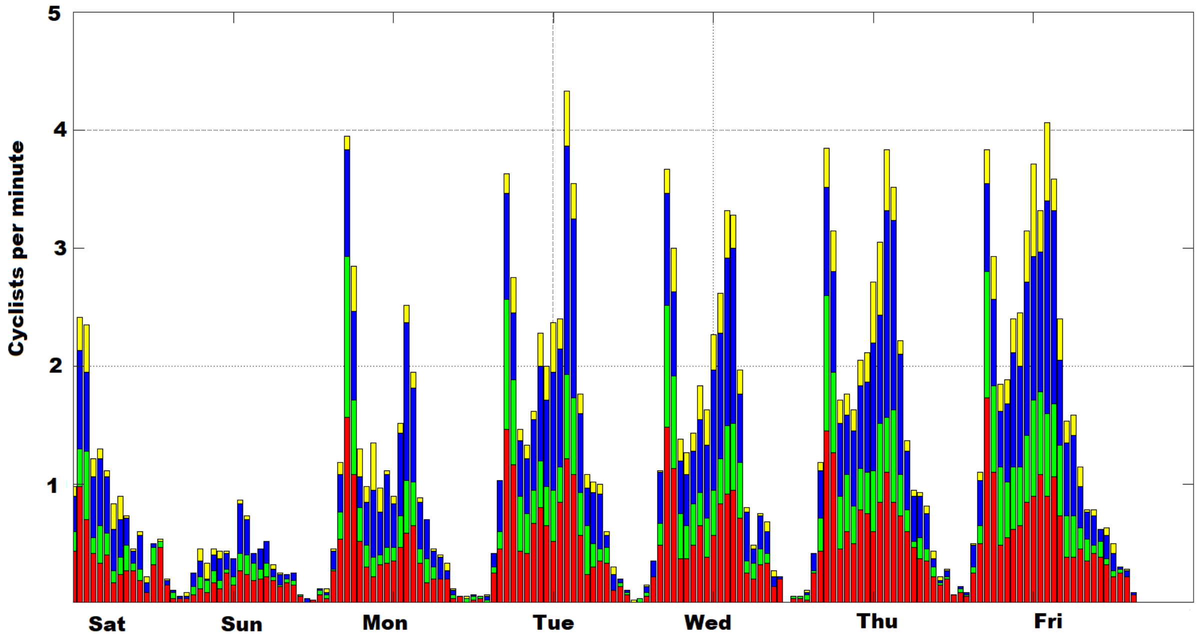

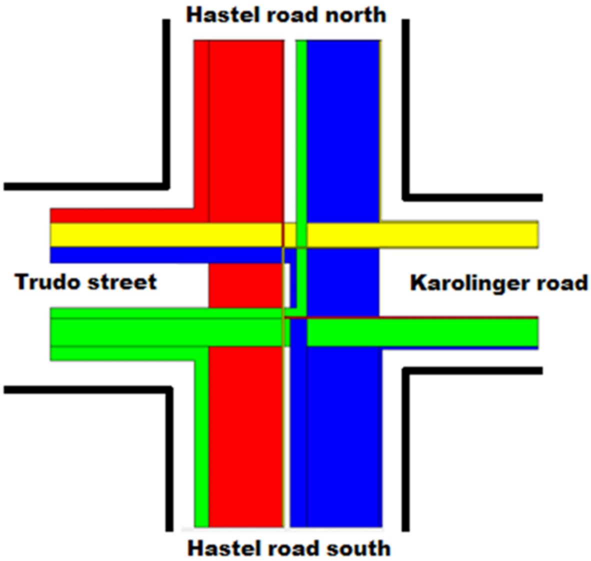

Figure 12 The top figure shows the time and day dependent cyclist flow over the intersection (in cyclists per minute) for every hour for a complete week divided (stacked) in the direction which they enter the intersection corresponding to bottom figure. The bottom figure shows the mean Sankey diagram of the cyclist manoeuvres over the 4-armed intersection: the width of the bar indicates the number of cyclists following one of 12 possible manoeuvres. The colours indicate from which direction the cyclists approach the intersection.

3.4 Hazard prediction model

Hazard prediction is separated into two parts. The first part, which is based on the static world model, will provide a hazard based on the cyclists that cannot be observed (due to a view-blocking obstruction) and the second part, based on the dynamic world model, will provide a hazard on the cyclist that is actually observed from the perspective of the approaching car. These hazards can be combined into one hazard prediction or presented separately, based on the HMI development (Chapter 4) towards a successful nudging strategy. Furthermore, this strategy depends on the dynamic world model to include all cyclists that can be observed. Both parts are discussed separately below.

3.4.1 Static hazard algorithm

The static hazard algorithm is based on the general idea is that a driver needs a certain amount of time to detect, classify and brake for a crossing cyclist to avoid a possible collision. If there is less time available and there is a probability a safety critical cyclist appearing from behind the view-blocking obstruction, a hazard can be issued. A situation sketch is shown in Figure 13. Each needed part is explained below.

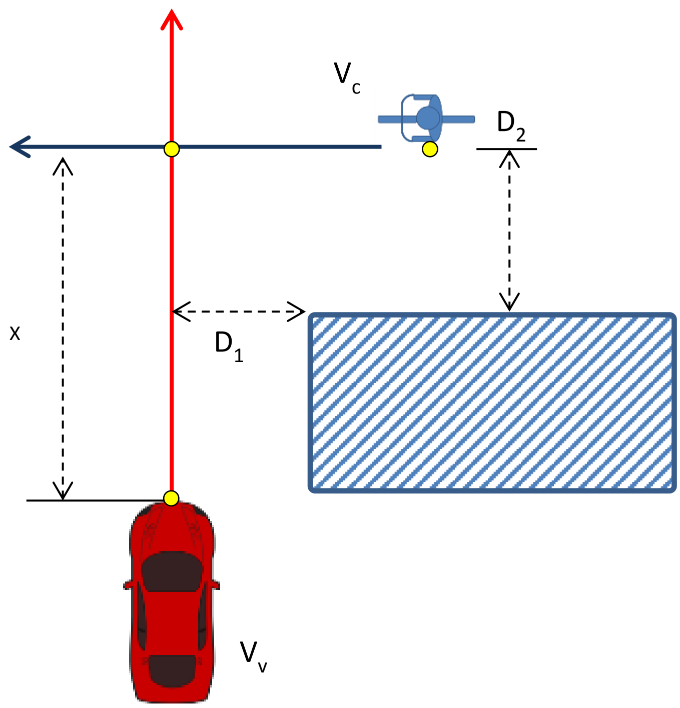

Figure 13 Sketch of a situation to explain the static hazard model. A safety critical cyclist may appear from behind a view-blocking obstruction (shaded box).

The available time (t_a) for the driver to react to prevent a collision can be simplified by assuming a constant characteristic deceleration (a). Doing so, this time (t_a) can be computed by subtracting the time needed to come to a standstill (t_s) from the time to reach the impact point (TTC):

(1) $$TTC = \frac{X}{V_v} \qquad [s]$$

(2) $$t_s = \frac{V_v}{2a} \qquad [s]$$

(3) $$t_a = TTC - t_s \qquad [s]$$

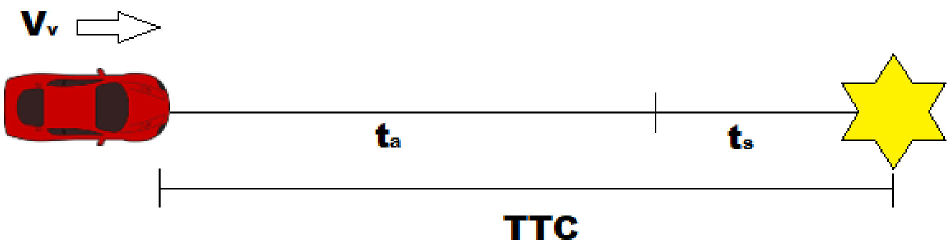

Figure 14 Schematic overview of the different times defined. TTC: Time-to-Collision when velocity remains constant, t_s: amount of time needed to come to a standstill, t_a available time to detect, classify and apply brakes. Note that when the velocity of the vehicle changes the TTC becomes smaller and t_s larger.

This means that, next to the speed of the host (V_v), the distance to the impact point (X) is also required (in real-time). The next step is to define a characteristic time (CT) that is preferred for collision avoidance. A simple linear interpolation is used to determine how much of this characteristic time is used (R) as a fraction between 0 and 1:

(4) $$R = \frac{CT - t_a}{CT} \qquad [-]$$

When a vehicle is entering the intersection with a constant speed V_v, R will start at 0 and will linearly rise over time to 1 when the last moment for braking has been reached at CT = t_a.

The hazard level does not only depend on the capability to brake up to the potential collision point. It needs to be combined with the probability of encountering a safety critical cyclist from behind the view-blocking obstruction. A cyclist is defined as safety critical when the cyclist itself including a safety distance (S) is within the width of the vehicle (W) extrapolated to a Time-to-Collision (TTC) of 0 sec.

Would a cyclist appear from behind the view-blocking obstruction, the following equation is used to compute the cyclist velocity range which is needed for the cyclist to become safety critical:

$$V_c = \frac{V_v D_1}{V_v TTC - D_2} \pm \frac{(S+W)}{TTC} \text{ [m/s]} \quad (5)$$

In other words, if the cyclist appears from behind the obstruction with a velocity $v=V_c$, and the vehicle would not change its velocity, then a collision will result.

To compute V_c, information regarding the location of the view-blocking obstructions of the intersection (see Figure 13) is required. Here, D_1 is the lateral and D_2 the longitudinal distance from the impact point to the path of the vehicle.

To include the layout of the intersection into a hazard model, the intersection is simplified as presented in Figure 15. The locations of the view-blocking obstructions relative to the paths of the vehicles and cyclists from all four directions are shown. Using this simplified model of the intersection, only possible cyclists from the near and far side are relevant for the static hazard model. Cyclists from the oncoming direction are assumed to be clearly in the sensor's field of view of the approaching vehicle and these cyclists are then covered by the dynamic world model and the dynamic hazard algorithm. Furthermore, it can be seen that the near and far side are not symmetrical (in terms of both the impact point and the location of the view-blocking obstruction), meaning that a separate hazard needs to be issued for both directions.

The probability for a cyclist to have a velocity within the cyclist velocity range (p) computed in equation (5) is derived from a naturalistic cyclist velocity probability density function [6]. This function is modelled by a normal distribution with a mean of 15.1 km/h and a standard deviation of 2.6 km/h.

Combining this probability (p) with the ratio of available time used from equation (4) will provide an instantaneous hazard level (Hi) between 0 and 1:

(6) $$Hi = R * p \qquad [-]$$

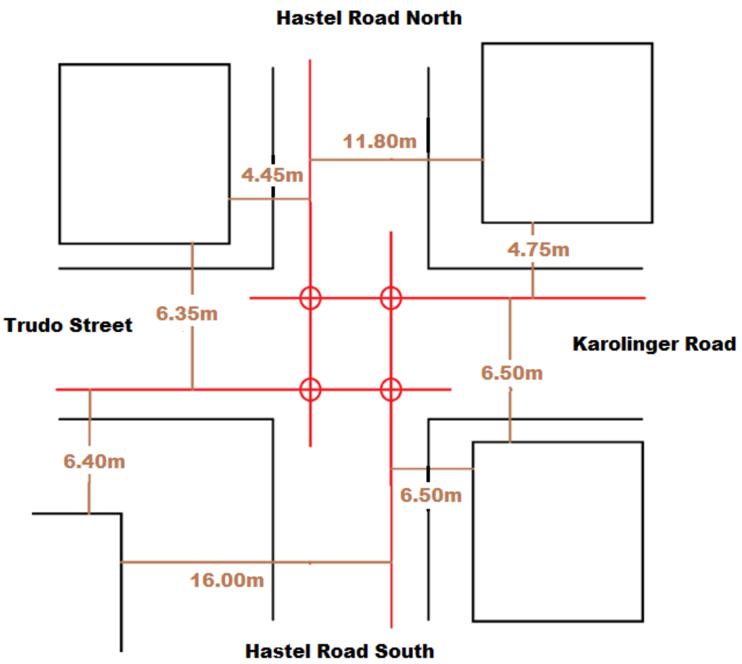

Figure 15 The simplified intersection with the locations of the view-blocking obstructions including the paths and possible impact points from all 4 directions.

Since the hazard level is now defined when a cyclist appears, the final step is to scale the hazard level with a probability of appearance (Cr) based on the cyclist flow data presented in Figure 12. It is not needed to provide the actual probability of a cyclist appearing at each time instant, but to represent the difference in cyclist flow in a ratio for the hazard level. Due to the statistical nature of this algorithm, it is not desired to have no hazard when the cyclist flow is low or even zero (e.g. at night), since it is still possible for a cyclist to appear. Furthermore, at a certain point it is busy enough to cap the hazard level at a maximum. Also, it is desired to increase the predicted hazard level faster at lower cyclist flow compared to higher cyclist flow to be conservative with a hazard warning. This can be achieved by using an exponential function:

(7) $$Cr = 1 - e^{\left(\frac{ln(1-a)-ln(1-b)}{Cc}C + ln(1-b)\right)} \quad [-]$$

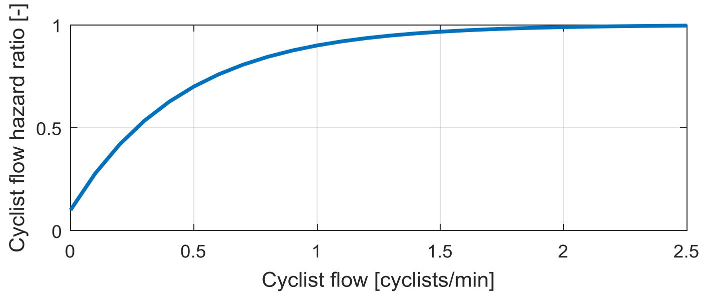

Figure 16 Cyclist flow hazard ratio (Cr) as a function of cyclist flow (C) based on (7) with a is.9, Cc is 1 and b is 0.1.

With (a) being the ratio reached at a cyclist flow of Cc and (b) the initial ratio at a cyclist flow of 0. C is defined as the cyclist flow relevant for the vehicle manoeuvre. For example, when the vehicle is moving straight, cyclists coming from the near side are assumed to be all relevant independent of their intended direction over the crossing, while cyclists from the far side are relevant only when they travel straight or turn left. All possible variations are depicted in Table 2.

Vehicle manoeuvre	Relevant from nearside	Relevant from farside
Straight	All	Straight/Left
Left turn	All	Straight/Left
Right turn	-	Straight/Left

Table 2 Relevant cyclist flow with respect to the vehicle manoeuvre for this specific crossing. Note that all manoeuvres are defined from their own point of view.

The final hazard level (H) then becomes:

(8) $$H = R * p * Cr \quad [-]$$

where H is a value between 0 and 1.

Summarizing:

- *R* is the ratio of the available time being used before entering a safe and smooth braking until standstill with 1 being fully used and 0 being fully available.

- *p* is the probability of a cyclist having a velocity which will lead to a critical situation if a cyclist would appear from behind the view-blocking obstruction (between 0 and 1).
- *Cr* is a qualitative measure to represent the probability of a cyclist appearing at each instant, with a value between 0 and 1.

3.4.2 Illustration of the static hazard model

To illustrate the above described methodology, an example is provided. Imagine a vehicle approaching the intersection from the Hastel Road North on a Monday morning at 9am and intending to move straight to the Hastel Road South. To show the difference in hazard level, the results are shown for two vehicle approaching speeds: 40km/h and 15km/h.

By using the velocity of the vehicle (V_v), a characteristic time (CT) of 3 seconds and a desired deceleration (a) of 4 m/s^2, the ratio of available time (**R**) can be computed using equation (4). Figure 16 shows the result as a function of the distance to the impact point.

Combining the view-blocking obstruction distances from Figure 15 with the vehicle velocity, a safety distance of 0.5 m and a cyclist length and vehicle width of 1.8 m, a cyclist velocity range can be computed using equation (5). By using the normal distribution of cyclist speed, the probability (**p**) that a bicycle meets a speed that results in a collision, can be determined. Figure 17 shows the result as a function of distance to the impact point.

The situation demonstrates that all cyclists from both the nearside Trudo street and the farside Karolinger road are relevant. From the cyclist flow, it can be determined that the cyclist flow at that moment is 1.5 and 0.1 cyclist per minute for the nearside and farside respectively. When using equation (7) with the parameters from Figure 17, the cyclist flow hazard ratio (**Cr**) becomes 0.97 and 0.28.

Combining all results by using equation (8) provides the hazard level, that is also visualized in Figure 18 as a function of distance to the impact point. Since the cyclist flow ratio (Cr) for nearside is almost 1, the hazard level is determined by the ratio of time available (R) and probability of a safety critical cyclist speed (p). It can be seen that the highest hazard level is when R is high (no time to brake anymore) at the same time when p is high as seen in the case with a vehicle velocity of 40 km/h and the near side view-blocking obstruction. When the velocity of the vehicle is lowered to 15 km/h the maximum of p occurs when R is still low, thus lowering the hazard level considerably. It can be seen, on the other hand, that the hazard level at the exact position of 10m before the impact point is higher for 15km/h than for 40 km/h. This is because at a distance of 10m with a car velocity of 40 km/h, an appearing cyclist will have to have a very high velocity to be safety critical, which is very unlikely. At this point the dynamic world model has to take over for the cyclists that come into the view of the car's sensor system.

The hazard level from the far side is much lower since Cr is low. Furthermore, the maximum of R and p does not occur at the same time, keeping the hazard level low. With 15 km/h these two parameters are even furthermore apart, lowering the hazard level even further to almost 0.

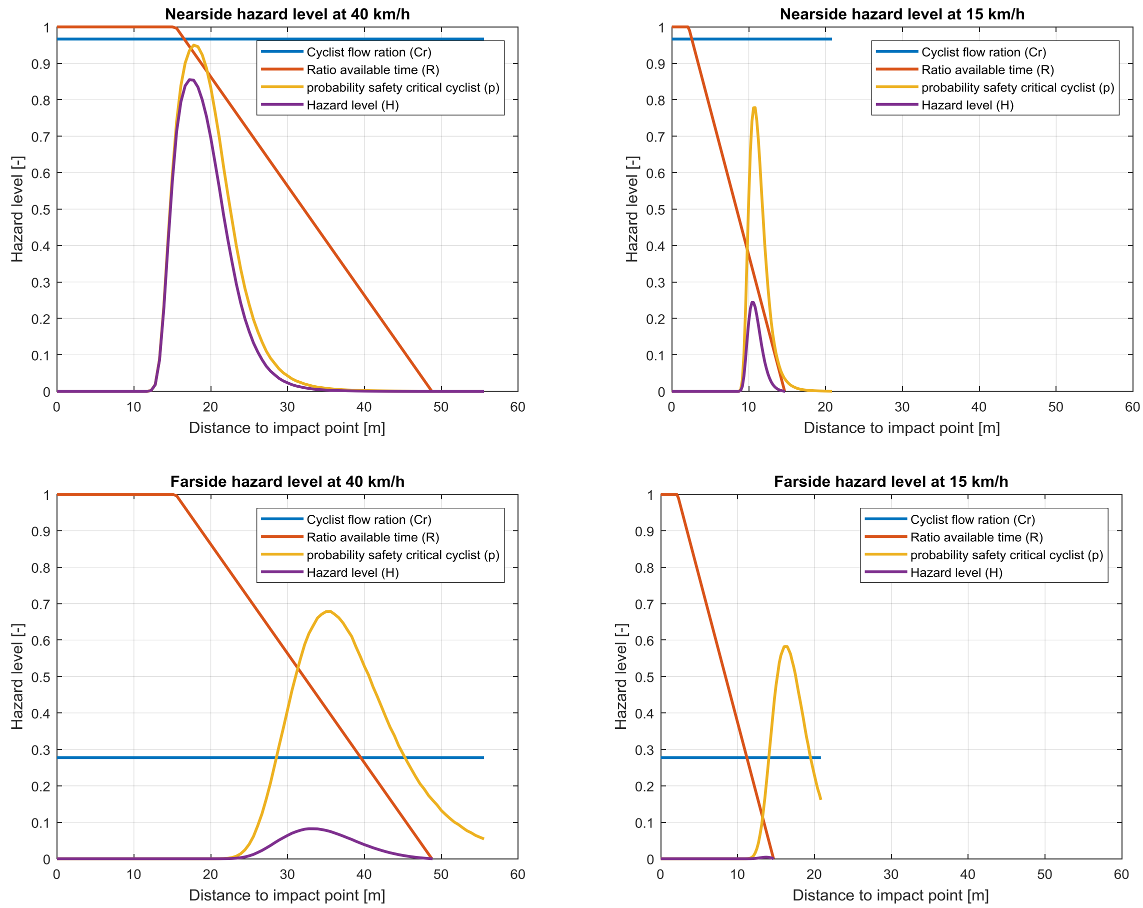

Figure 17 Hazard level (purple) with respect to the distance to impact point including all separate components. Top is from nearside and bottom from far side. Left at 40km/h and right 15 km/h.

The final step is to compute a suggested 'appropriate' velocity for the vehicle. This can be done by calculating the hazard levels for different velocities as shown in Figure 18 for the parameters used in the example above. By keeping the hazard level constant (at a low level), a suggested velocity can be obtained as shown in Figure 19.

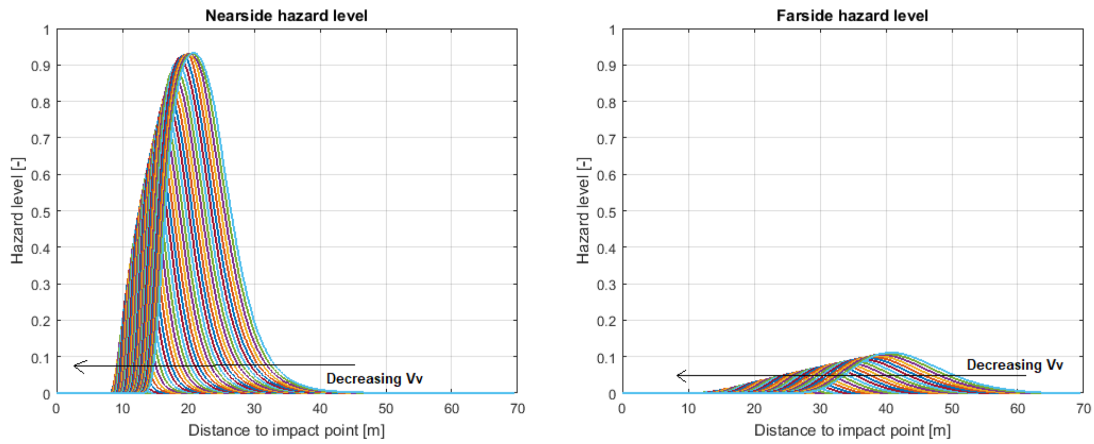

Figure 18 Hazard levels with respect to the distance to impact point for different vehicle speeds (10 to 50 km/h). Left for the nearside and right for the farside.

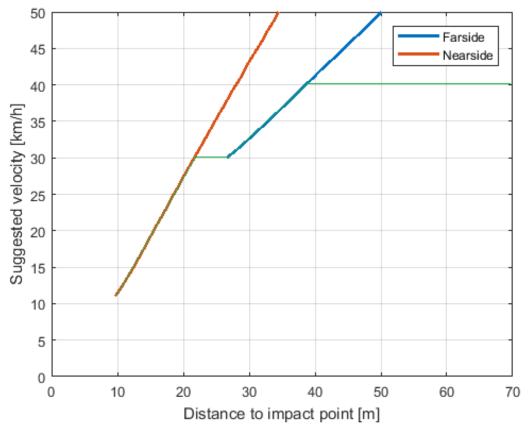

Figure 19 Suggested 'appropriate' velocity with respect to the distance to impact point(s) for the near (red) and farside (blue). The green line shows an example of what the suggested velocity for the vehicle would be if this vehicle approaches with 40 km/h.

3.4.3 Dynamic hazard algorithm

From the world model as described in Sections 3.1 and 3.2, trajectories (and their probabilities) of the vehicle and cyclist(s) are available. From these possible trajectories the minimum distance between vehicle and cyclist can be computed at any point in time. It is suggested that the dynamic hazard level is set to 1 when this distance is within the path of the vehicle and linearly decrease to zero when this distance is at the safety distance. By doing this, safety critical cyclists will immediately be a priority and once action is taken (by the vehicle or the cyclist) can be quickly decreased in priority preventing a high workload for the driver. Figure 20 shows an example of the combination of the dynamic and static hazard level for the nearside crossing with a 40 km/h vehicle speed. The static hazard level is identical to the example provided in Figure 11. In this scenario, a cyclist is present with a speed of 15 km/h and at collision course with the vehicle at 1 m offset to the farside of the vehicle centerline. It can be seen that the dynamic hazard becomes 1 immediately when the cyclist appears from behind the view-blocking obstruction at 20m from impact point. The dynamic hazard remains 1 in this case since in the example, both the vehicle and cyclist do not alter their velocity or direction.

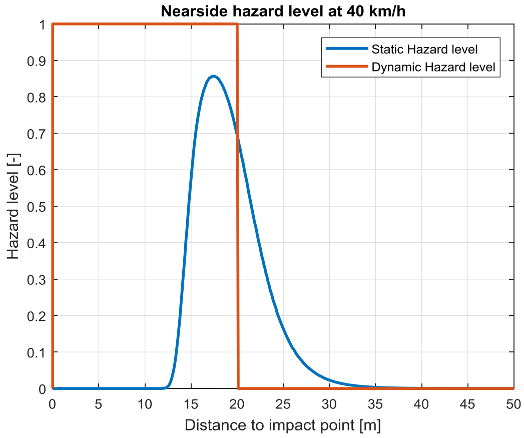

Figure 20 Dynamic and static hazard levels for nearside crossing at 40 km/h. The static hazard level is identical to the hazard level in Figure 16 (top left). The dynamic hazard level immediately becomes 1 when the cyclist is visible from behind the view-blocking obstruction.

3.5 Discussion

The static and dynamic world model elements have been described separately in the previous sections. In the final world model they are however, fully *integrated*. That is, the complete world model consists of static and dynamic elements within one representation. All elements, whether static or dynamic, can be associated with hazard levels in the same manner, and provide input accordingly for HMI attention (Chapter 4).

The presented hazard algorithms contain parameters that need to be tuned towards a good balance between nudging and comfort, where adaption of behaviour towards the hazard is desired without invasive warnings. The parameter values for the hazard model will be determined making use of a combination of simulation studies and drive simulator tests (task T2.5).

Any ADAS system, such as AEB, makes use of a world model. AEB systems are only activated in imminent critical (or close to critical) situations, to prevent that an undesired forceful braking action occurs posing a safety risk while the situation does not ask for it (false positive response). Static and dynamic objects in the world model of an AEB system objects are only considered in case of high confidence that these objects are actually present and need to be responded to. For nudging systems, a response on an actual or potential hazard at lower confidence is not required. The requirements of a world model for an AEB system are consequently different from those of a world model for nudging.

Although the requirements on the world model are different, they are both fed with information from the same sensor set, so that it might be expected that the same objects are used as input for both emergency and nudging system. In the project, no integration is foreseen between the AEB system that is present in the vehicle and the nudging system that is going to be implemented. Both systems can run simultaneously, and because the dynamic models receive input from the same sensor set, no contradiction in activation of the system is expected.

The AEB system will not be altered. In case of an imminent critical situation, the AEB will take control and act as a last resort accident avoidance/mitigation system.

4 Development of HMI options

In this chapter, options are described for the human-machine-interface (HMI) that is intended for Directing Driver Attention in a nudging style.

The aim of the HMI is to transfer knowledge about a current possible hazardous situation and nudge the driver into safer behaviour. The situation addressed here considers hazardous intersections (e.g. intersections with limited visibility where the driver of the car crosses the path of vulnerable road users on wheels). The in-vehicle HMI should nudge the driver into a safer behaviour several seconds before e.g. a cyclist is crossing the path of the host vehicle. In this case several behaviours can be considered as safe (as described in WP1 and in the proposal of the project):

1) Appropriate speed: If a driver is approaching a hazardous intersection, the driver should be nudged to drive slower (Connection to section 5.1.1 in D1.1).
2) Mental Preparation: When approaching a hazardous intersection, the driver should be mentally prepared to be able to react as fast as possible in case a cyclist crosses his path. (Connection to 5.1.3 in D1.1 -Inattention to possible risk)

To reach these aims, the HMI must consider several aspects:

- An appropriate nudging solution has to be chosen according to the decision logic and control model (T2.4.1)
- The nudge output has to be transferred to an appropriate information display (T2.4.2)

Both aspects are described in detail in the following sections starting with an initial version of the decision logic and control model in Section 4.1 and the development of different information displays in Section 4.2. The choice of potential candidates for future tests is explained in Section 4.3. The section closes with a short discussion of what is measured in future experiments and tests.

4.1 Decision and control logic

The decision logic and control model connects the output of the world and hazard perception model (T2.2 and T2.3) to the HMI.

Two basic outputs coming from the world model based hazard prediction have to be recognized: The prediction of what is likely to happen (e.g. the localization of road users with intended trajectories and the probability of such trajectories) and how hazardous this is (hazard level).

The hazard level can be used to make a decision for an appropriate speed (needed to be in line with aim 1). This information needs then to be transferred to the HMI to be shown as part of the nudge output. Additional values that need to be shown on the HMI component have been systematically derived.

For deriving the information elements that need to be perceived by the driver a hierarchical task analysis (HTA) has been conducted. This is quite important as "a thorough understanding of the behavioural pattern is considered an important prerequisite for developing appropriate nudge countermeasures." [22].

The task analysis is shown in Figure 21.

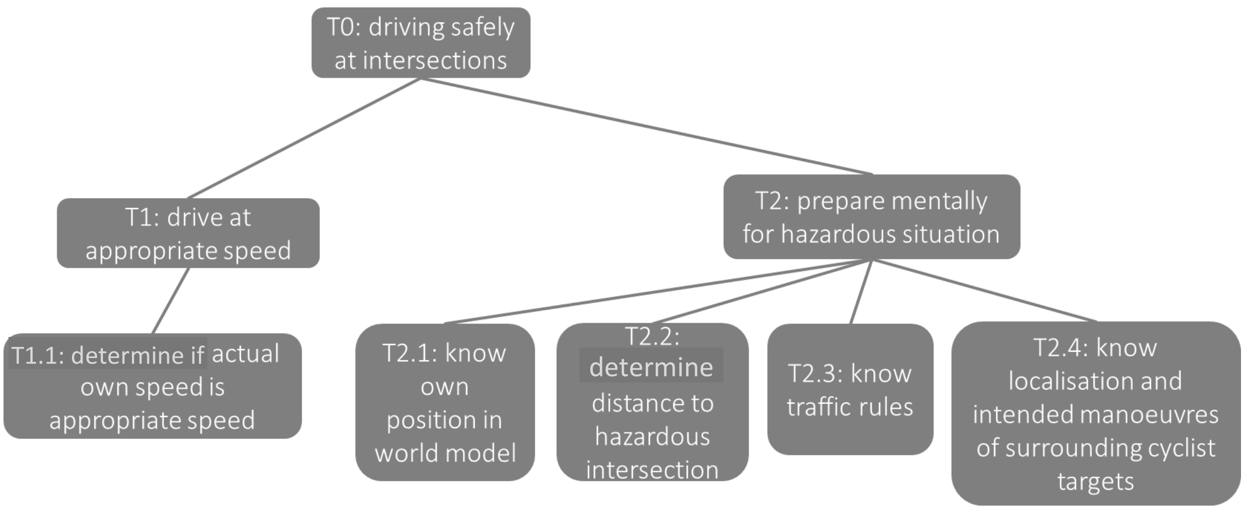

Figure 21 Task analysis for directing driver attention.

Based on the Task Analysis (TA) the information the driver needs to know for accomplishing the task of driving safely at intersections can be summarized in the *Table 3* Overview of needed information transfer to the driver.

With the information elements shown before, and especially the sources from which the information elements come from, an initial version of the decision logic and control model is shown in Figure 22. In this Figure it is highlighted which values needs

based on TA (Figure 21)	Information	Source
T1.1	own actual speed	Car sensor
T1.1	appropriate speed	Derived Value according to hazard level calculated in world and hazard perception model from T2.2 & T2.3
T2.1	own position in world model	World model from T2.3
T2.2	distance to hazardous intersection	World model from T2.3
T2.3	traffic rules	World model from T2.3
T2.4	localisation of surrounding cyclist targets	World model from T2.3
T2.4	*intended trajectory of cyclist targets*	*Prediction of cyclist intend from T2.2*

Table 3 *Overview of needed information transfer to the driver*

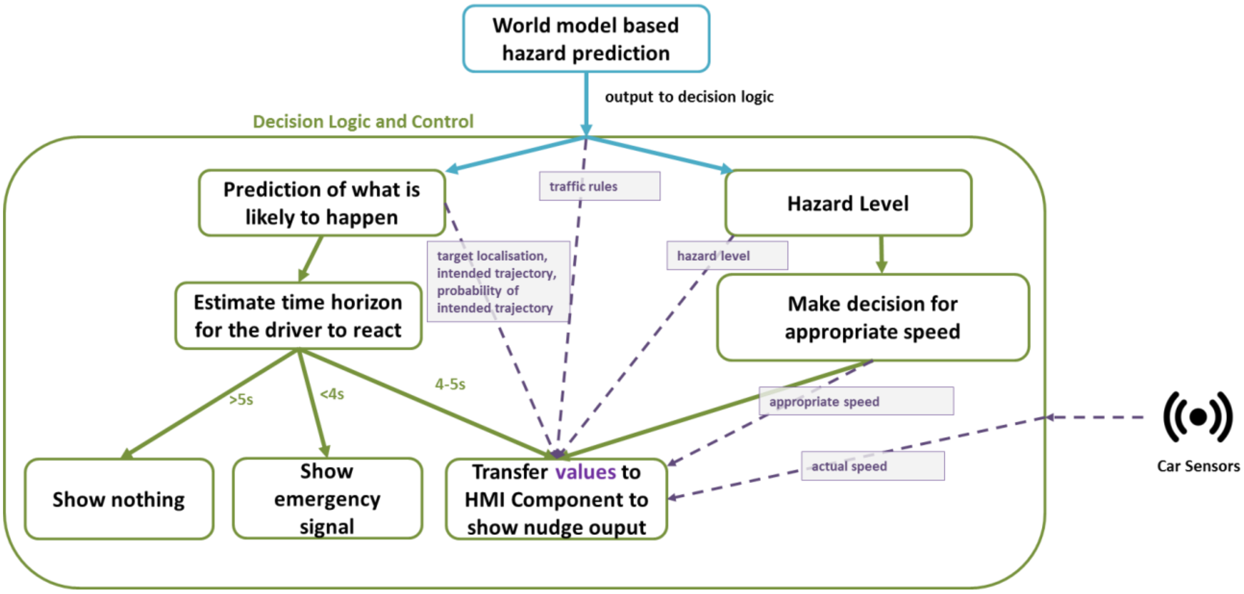

Figure 22 Initial draft for decision logic and control model.

to be transferred to the HMI to be usable as variables in the HMI visualization (shown in violet). The information transferred from the world model based hazard prediction via the decision logic and control model to the information display has to be visualized as nudge output to the driver. The process conducted for generating different design solutions and the concept of different possible design solutions are presented in detail in the upcoming section.

4.2 HMI design solutions

The design of the information display should follow the nudging approach as defined in WP1. In this regard the framework in WP1 has found that the safety by design approach corresponds with the basic ideas of nudging, that the task environment is designed in such a way that the desired behaviour is almost automatically followed ([5] section 6.2.1). As an example: the use of red colours for alert signals can be applied.

Furthermore, nudging focuses on subconscious processes (System 1 based on Shriffin and Schneider, [5] section 6.6) and thus also the information displayed should be perceivable in this way (automatic, low effort, rapid).

In WP1 cognitive biases are addressed. This is also important with regard to the information display and should be considered during the design phase. Especially the bias in perception of risk and severity (section 6.7.1 in D1.1) is important as the use case (Directing Driver Attention) deals with directing driver's attention to a possible hazardous situation and inducing a safer behaviour. This involves for instance the optimism bias (The tendency to be over-optimistic, overestimating favourable and pleasing outcomes as this is based on the experience that in most cases no accident happens).

The theoretical background about nudges has some design implications for the derivation of nudging measures and an appropriate information display (for details see [5] section 7.2.5): "Purely informing the drivers about the 'how' and 'why' of safe driving may not achieve the desired results because providing information to motivate behaviour change almost exclusively runs via cognitive effort (i.e. the reflective route)". In this aspect the cognitive effort should be kept as low as possible: "Cognitive effort is only attributed when people are motivated to do so and have the capacity" ([5] section 7.2.5) "Thus, information that enters the brain without demanding cognitive control or capacity is more likely to influence behaviour when the driving itself takes up a lot of effort". Thus the focus must lie on simple cues that are fast and easily perceived correctly.

To ensure that the driver can perceive the information needed as fast and correctly as possible, a model based approach called the Konect method [23] is applied. The Konect method has already been applied successfully for deriving HMIs for truck platooning to induce a fast reaction of the truck driver in case the ACC fails (see [23] and [24] for details).

Based on the Konect model based approach different design concepts have been developed in a workshop conducted with WP2 partners on 30.01.2018 at OFFIS in Oldenburg. The design concepts are described in the upcoming subsections: Concepts

1 – 3 have been developed in the WP2 Workshop on in-vehicle solutions. Concept 4 has then been derived post-hoc based on the ideas of nudging in concepts 1 – 3.

4.2.1 The „Nudging Blob“ Concept

The “Nudging Blob” Concept is shown in Figure 23 and represents the actual state via a symmetric circle – called “Nudging Blob”. In the middle of the circle a bar indicates the own actual speed of the car. The line above of this bar indicates the quantitative value of the appropriate speed. If the own speed complies with the appropriate speed the bar is coloured in green. Otherwise the bar changes its colour from green to red. In case targets occur e.g. a cyclist, convexities are appearing. The convexity position represents the position of the real world target seen from the top e.g. a cyclist coming from the right side at an intersection will first appear as a convexity at the upper right corner and will move from the upper right corner to the right side when the car is further approaching the intersection. The colour and size of the convexity indicates the hazard level. This can be further emphasized by a shape change of the convexity – starting from a more circular form for a lower level ranging to a sharp form for a high hazard level.

The “Nudging Blob” concept intuitively addresses humans preference for symmetric forms. In case a hazard occurs, the symmetry is disturbed and the driver will intuitively recognize that the situation is becoming hazardous. Furthermore the speed in the middle changes its colour to red. To further emphasize this, the size of the bar can be adopted. In this concept, the human desire for symmetry can be used to subconsciously nudge the driver into adapting their speed and looking at the hazardous targets.

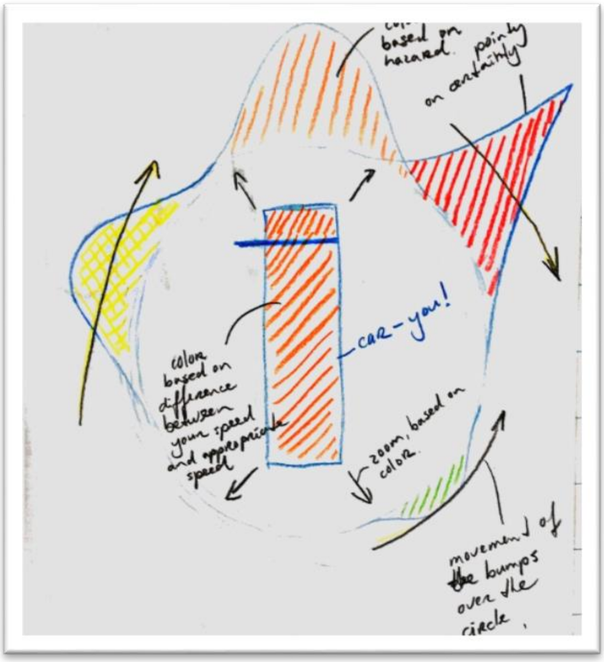

Figure 23 The Nudging Blob concept as sketched in the HMI design workshop.

The sketched design concepts have been transferred to digital versions of the different concepts in form of pictures showing different situations. This is shown in Figure 24.

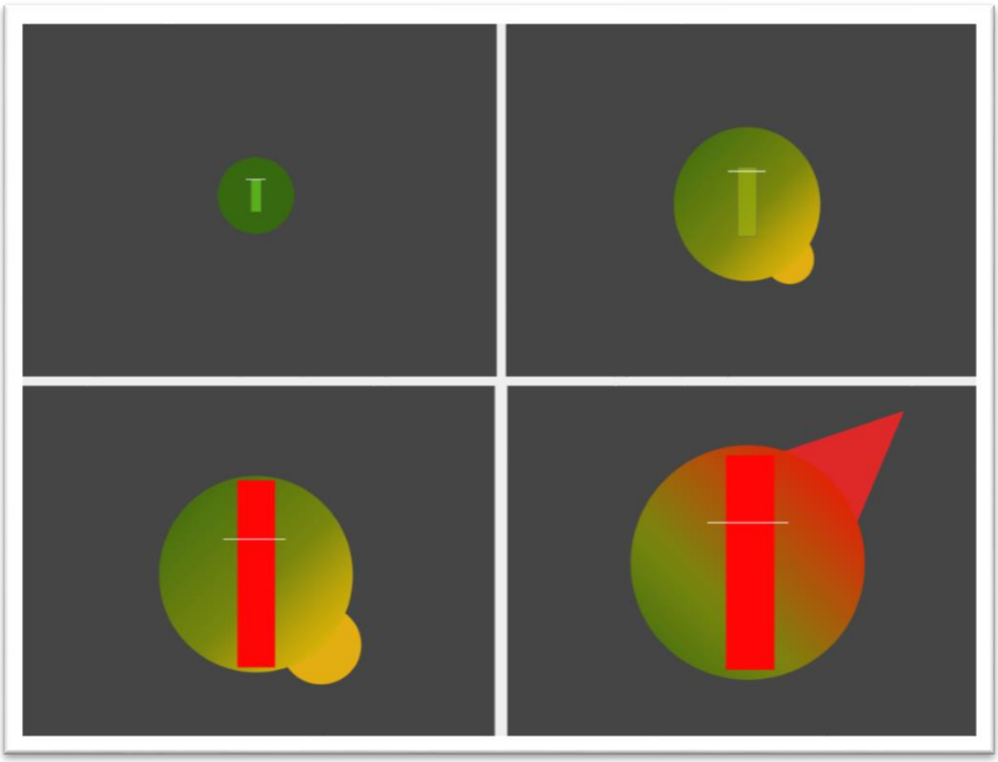

Figure 24 The nudging blob concept.

4.2.2 The Augmented Representational Scene Concept

The Augmented Representational Scene Concept is shown in Figure 25 as sketched in the WP2 design workshop. Compared to the “Nudging Blob” concept described in the prior section, the augmented representational concept aims for a more realistic representation of the scene. This has the advantage that the cognitive step for translating the visual scene to the realistic setup outside of the vehicle has not to be made by the driver. Thus the scene allows intuitive understanding. As disadvantage the desire for symmetry is addressed less compared to the “Nudging Blob”.

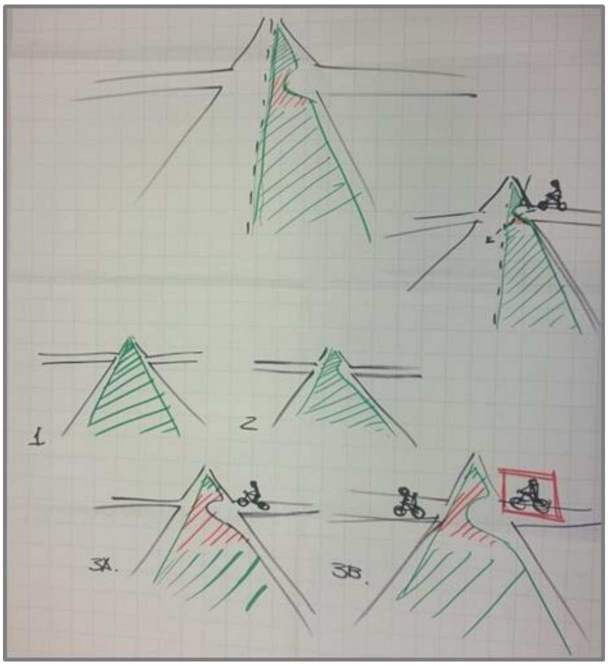

Figure 25 Augmented Representational Scene Concept sketched in WP2 design workshop.

In the Augmented Representational Scene Concept the own trajectory is visualized via a green corridor (this is commonly used to show free space e.g. see for example https://www.mobileye.com/our-technology/). Possible hazards will lead to indentation. Around these indentations the colour also indicates the hazard level. For further emphasizing hazardous targets, these targets can be surrounded with a red frame. The Augmented Representation Scene can be displayed as Cockpit display as shown in Figure 26 or as Head-up Display as shown in Figure 27. In case a slightly hazardous situation appears it seems to be difficult to perceive as the slight indentation is not salient in the distance. Thus for slightly hazardous situations in the

Head-up Display an additional colour (yellow/orange) is added for indicating slightly hazardous situations.

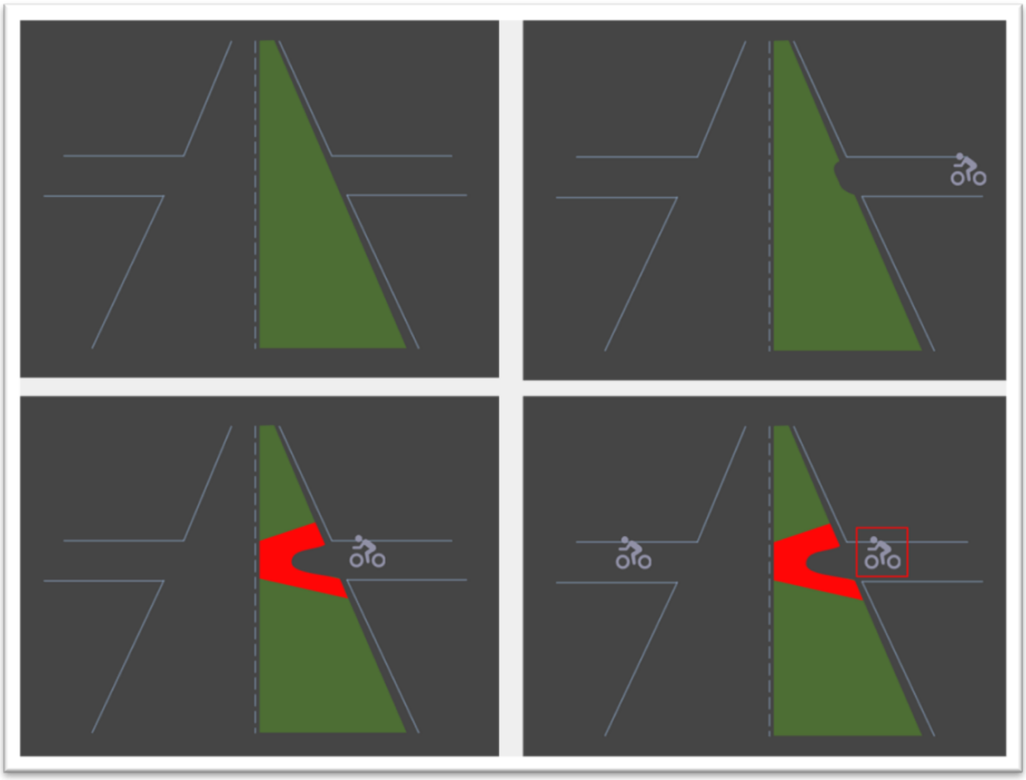

Figure 26 Augmented Representational Scene as cockpit display.

Figure 27 Augmented Representational Scene as HUD.

4.2.3 The Static Concept

The Static Concept is shown in Figure 28 and is quite similar to the augmented representational scene concept. Compared to the Augmented Representational Scene Concept, the Static Concept also represents the real spatial conditions. Instead of working with a dynamic corridor (as compared to the augmented representational scene concept), the Static Concept works with static indicators.

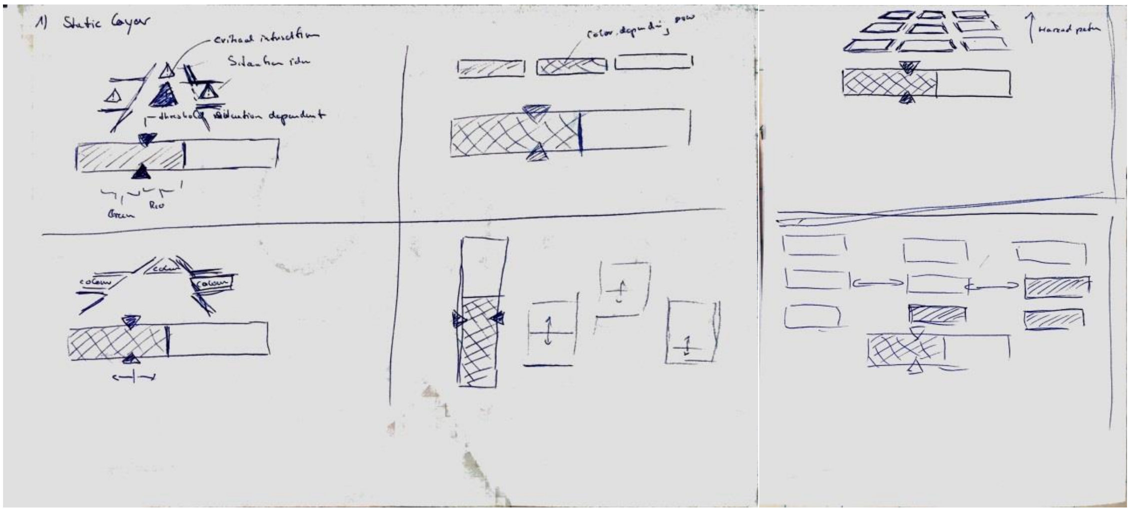

Figure 28 The Static Concept as sketched in the WP2 workshop.

Figure 29 shows the digital output of the sketches. The speed is visualized as a bar at the bottom. If the speed is appropriate the bar is coloured green, if it deviates from appropriate speed (which is indicated via the position indicators), it is coloured in red. In case a target is detected, a warning sign appears showing the position of a cyclist.

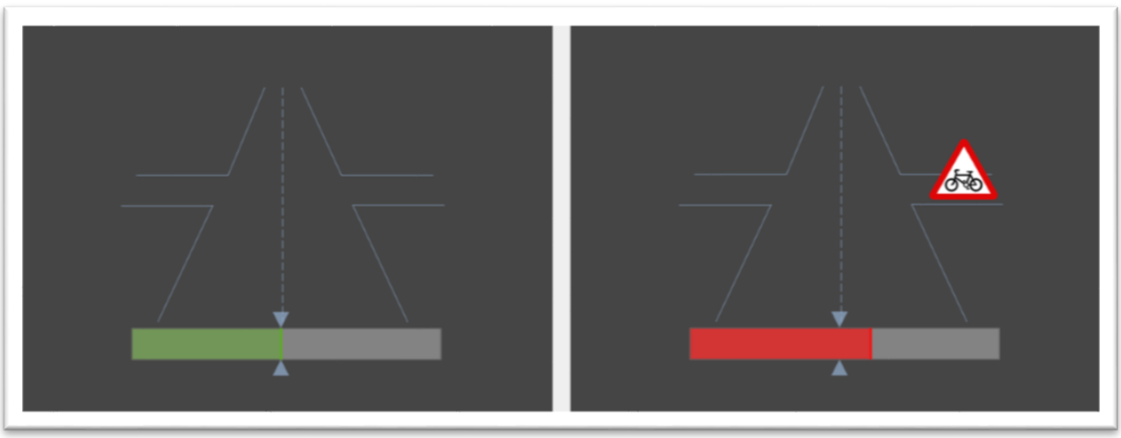

Figure 29 Example of the Static Concept

4.2.4 Mixed Concept (Nudging Cross)

Based on the ideas developed in the HMI Workshop, a mixed version recognizing the ideas for nudging of the workshop participants has been derived. This is shown in Figure 30.

This mixed concept is the "Nudging Cross". Similar to the Nudging Blob the Cross represents a symmetric form. Deviations from the symmetry appear at the side where a target occurs (e.g. a cyclist coming from the right side) as indentation. In this case it is more similar to the augmented representational scene concept as the cognitive translation step from the reality to the visualization in the cockpit display can be accomplished more easily. The size and form of the indentation represents the hazard level. The colour and size of the cross represents the speed. The higher the speed deviation is compared to the appropriate speed level the bigger the cross will appear and will be coloured in red. If the speed is ok and no hazard is foreseen, the upcoming intersection will be visualized as a small, symmetric green cross. This is shown in Figure 30.

Figure 30 Nudging Cross as cockpit display.

In all design options, the color was chosen based on the mental model of drivers in Europe (based on the color coding of traffic lights). This was done to allow intuitive understanding and thereby nudging without the need to learn a new meaning of color codes.

4.3 Decision for potential candidates

In the prior section, the potential candidates for the HMI solution are described. In this section a more detailed description is given why these candidates are chosen for future tests and why other candidates were excluded.

For a successful nudging solution it is important that the solution is quickly recognizable by the driver. At the same time, the solution should not be too obstrusive as this is one of the most important aspects that differentiate the nudging solution from warnings or alerts. Two criteria were considered during candidate selection: the solution should provide **fast intuitive recognition** and at the same time it should be **not too obstrusive**. This is shown in Figure 31. The previously described candidates are fast to recognize (due to the application of the Konect method and usage of visual attributes that are perceived preattentively by humans, e.g. color, shape). Other visual designs, that are possible but have been excluded, are those that use text for encoding information (e.g. displaying the text "hazardous intersection ahead" etc.). These have been excluded as they are not fast perceivable (not preattentive). This will not only result in slow perception of information but might also prevent the driver from looking at the information as this is associated with effort.

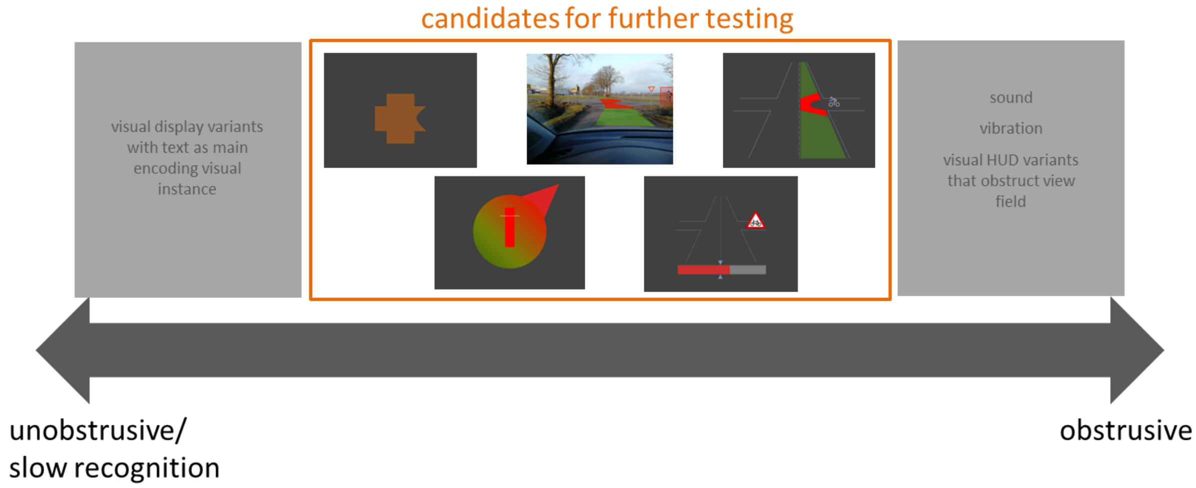

Figure 31 Design space and choice for candidates for further tests.

On the other hand, there exist design variants that have been discussed in the workshop as being too obstrusive, e.g. using sound, vibration or head-up display variants that block the field of view of the driver (see Figure 32 for example). In these cases, the design might provoke a harsh breaking event (with possible safety risk) as instinctive reaction. This is not intended by the design solution as the design is not used in emergency situations but to nudge the driver to slow down and thereby prevent the emergency case. This was the reason why these design possibilities were excluded as candidates for further testing.

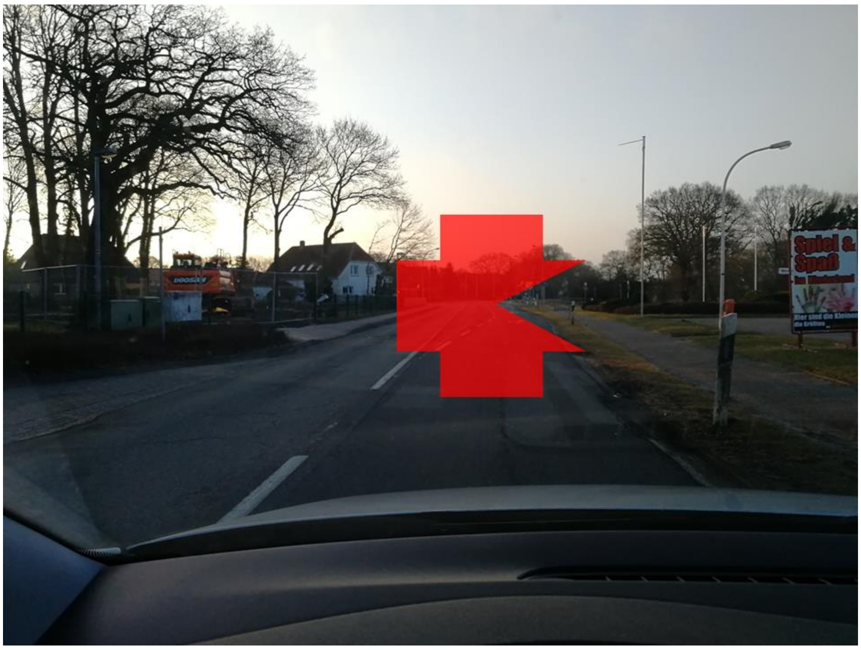

Figure 32 Obstrusive visual HUD.

4.4 Discussion

Chapter 4 presented the initial draft for the decision logic and control model and different HMI design solutions for directing driver attention at hazardous intersections. In future tests or simulations, the following parameters need to be determined or tested:

- For the decision logic and control model, it has to be determined at what time the nudging output should be shown to the driver and what is an appropriate speed (c.f. section 4.1) according to different predicted hazard levels. The following input is needed for the HMI: hazard level, appropriate speed based situation and hazard level, actual speed of the vehicle, position and intended trajectory of cyclists. These inputs are needed from the cyclist trajectory and hazard prediction model.

- For the different design solutions it has to be determined how far they are able to achieve the main aim of the directing driver attention HMI (appropriate speed and mental preparation as stated in the introduction) and the distraction that such HMI might pose. This has to be determined in future tests (task T2.5).
- How is the reaction of the drivers in different situations with the different design solutions? Does the driver adapt speed in the right manner? Does the driver adapt gaze behaviour and feet position according to the situation (to assess the influence of the design on the mental preparation)?

These tests are scheduled to be performed in Task 2.5, e.g. to make a selection of the best possible option to be implemented into the FIAT test vehicle in MeBeSafe. Also in WP5, during the field tests, the performance of the selected nudging solution will be evaluated.

5 Concluding remarks

This document described the development of in-vehicle nudging solutions that target the objectives:

- Stimulate the use of safety functions onboard vehicles;
- Direct driver attention to potential hazards, specifically regarding cyclists.

The models and information flows have been described, starting at the level of in-vehicle sensors up to and including possible options for the HMI.

The options regarding the use of safety functions onboard vehicles, incl. the function that indicates a driver to take a break when the system senses a certain level of drowsiness of the driver, are further elaborated in WP4 (Driver Coaching).

The options of in-vehicle nudging solutions to direct driver attention to potential hazard will be evaluated in Task T2.5, in driving simulator tests and virtual simulations. The objective of this task is to propose one most feasible and promising solution to be implemented in a vehicle in Task T2.6 (a FIAT 500X). Moreover, simulations will be used to determine an initial set of values for some of the parameters that need to be selected in the hazard prediction model, such as the used characteristic times in equations (4) and (7). Similarly, simulation studies will be used to support design choices in the HMI, before implementing a limited number for testing in driver simulator studies.

Simultaneously, the basic models developed in tasks T2.2 and T2.3 will be further extended and improved based on the feedback received from the observation study performed as part of T2.2.

6 References

[1] http://www.udrive.eu

[2] Malta, L., Ljung Aust, M., Faber, F., Metz, B., Saint Pierre, G., Benmimoun, M., & Schäfer, R. (2012). EuroFOT final results: impacts on traffic safety. Retrieved from http://www.eurofot-ip.eu/ Source Date: 11/21/2012.

[3] Department for Transport, 2016. Reported Road Casualties Great Britain: 2015 Annual Report. , (September), pp.1–409. Available at: https://www.gov.uk/government/statistics/

[4] Victor, T.; Dozza, M.; Bärgman, J.; Boda, C.-N.; Engström, J.; Flannagan, C.; Lee, J.D.; Markkula, G. (2015). Analysis of Naturalistic Driving Study Data: Safer Glances, Driver Inattention and Crash Risk; Transportation Research Board: Washington, DC, USA, 2015; Available online: http://onlinepubs.trb.org/onlinepubs/shrp2/SHRP2_S2-S08A-RW-1.pdf

[5] "Sleepiness at the wheel", Institut National du Sommeil et de la Vigilance

[6] Karlsson M., et al. "Integrated Framework", D1.1 delivery report for MeBeSafe, October 2017.

[7] Schubert, A., Erbsmehl, C. & Hannawald, L., 2013. Standardized pre-crash scenarios in digital format on the basis of the VUFO simulation input data from GIDAS. In 5th International Conference on ESAR „Expert Symposium on Accident Research". pp. 366–372.

[8] Uittenbogaard, J., Op den Camp, O., Montfort, S. van, CATS car-to-cyclist accident parameters and test scenarios, CATS deliverable D2.2, TNO 2014, R11705, Helmond, September 2016.

[9] ISO 45001:2018 Occupational health and safety management systems – requirements with guidance for use, International Organization for Standardization, https://www.iso.org/obp/ui/#iso:std:iso:45001:ed-1:v1:en.

[10] Sanz, L., et al. "Accident Analysis, Naturalistic Observations and Project Implications", Deliverable D2.1 of the EU project PROSPECT, www.prospect-project.eu, 2016.

[11] Schleinitz, K., The German Naturalistic Cycling Study- Comparing cycling speed of riders of different e-bikes and conventional bicycles, International Cycling Safety Conference, 2014, Göteborg.

[12] Kühn, M., Hummel, T., Lag, A., Cyclist-car accidents – their consequences for cyclists and typical accident scenarios, The 24th International Technical Conference on the Enhanced Safety of Vehicles (ESV), Gothenburg, Sweden, June 8-11, 2015.

[13] Twisk, D., Nes, N. van, Haupt, J. "Understanding safety critical interactions between bicycles and motor vehicles in Europe by means of Naturalistic Driving techniques", Proceedings, International Cycling Safety Conference, 2012, Helmond.

[14] Helbing D. and Molnar, P. Social force model for pedestrian dynamics. Physical review E, 51(5):4282, 1995.

[15] Redmon, J., & Farhadi, A. YOLO9000: Better, Faster, Stronger. Paper presented at IEEE International Conference on Computer Vision and Pattern Recognition (CVPR 2017).

[16] Ren, S., He, K., Girshick, R., & Sun, J. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. IEEE Transactions on Pattern Analysis and Machine Intelligence, 39(6), 1137-1149, 2017.

[17] Reid, D. An algorithm for tracking multiple targets. IEEE Transactions Automatic Control, 24(6), 843–854, 1979.

[18] Leal-Taixe, L, Milan, A., Reid, I., Roth, S., & Schindler, K. MOTChallenge 2015: Towards a benchmark for multi-target tracking. arXiv:1504.01942, 2015.

[19] Wojke, N., Bewley, A. & Paulus, D. Simple Online and Realtime Tracking with a Deep Association Metric. University of Koblenz-Landau, Queensland University of Technology. arXiv:1703.07402, 2017.

[20] Cao, Z., Simon, T., Wei, S.E., and Sheikh, Y. Realtime multi-person 2d pose estimation using part affinity fields. CVPR 1 (2), 7, 2017.

[21] Pool, E.A.I., Kooij, J.F.P., and Gavrila, D. M. Using Road Topology to Improve Cyclist Path Prediction. Proc. of the IEEE Intelligent Vehicles Symposium, Redondo Beach (CA, USA), 2017

[22] Alahi , A., Goel, K., Ramanathan, V. Robicquet, A., Li, F.F., and Savarese, S. Social LSTM: Human Trajectory Prediction in Crowded Spaces, CVPR, 2016.

[23] Op den Camp, O., Hair, S. de, Gelder, E. de, Cara, I., Observation study into the influence of a view-blocking obstruction at an intersection on bicycle and passenger car velocity profiles, International Cycling Safety Conference, 2015, Hannover.

[24] http://inudgeyou.com/en/could-a-nudge-improve-safety-at-railway-platforms/

[25] M. C. Ostendorp, S. Feuerstack, T. Friedrichs, A. Lüdtke. Engineering Automotive HMIs that are optimized for Correct and Fast Perception. ACM SIGCHI Symposium on Engineering Interactive Computing Systems, 2016.

[26] M. C. Ostendorp, T. Friedrichs, A. Lüdtke. Supporting Supervisory Control of Safety-Critical Systems with Psychologically Well-Founded Information Visualizations. Proceedings of the ACM SIGCHI 9th Nordic Conference Human-Computer Interaction, 2016.

[27] Op den Camp, O., van Montfort, S., Uittenbogaard, J., Welten. CATS D6.1 Final project summary report, TNO 2016 R10921, www.tno.nl/CATS, July 2016.