

Background Traffic Agents for Driving Simulators

Simulating Traffic in Multiple Environments

M.F. de Goeij

Master thesis

ICA-3507637

t_de_goey@hotmail.com

Utrecht University

Supervisors:

Dr. Roland J. Geraerts

Dr. Ir. A. Frank van der Stappen

Abstract

In the field of agent models for driving simulators, there are few models aimed at simulators that teach students how to drive. By approaching the problem from a students' perspective, we hope to increase the learning capacity of the driving simulator. Furthermore, none of the existing agent models use the Belief, Desire and Intention software model, which forms the basis of our work. By using BDI, we can exert more control over the agents while remaining to display realistic behaviour. We validate our method by presenting the behaviour of our agents to both driving school students and driving instructors. Results show that our model can produce significantly deviating and realistic behaviour. Although surprisingly, it is deemed more cautious than intended.

This thesis project was conducted as part of a collaboration between Utrecht University and GreenDino B.V.

Contents

1	Introduction	5
1.1	Research Questions	6
1.2	Structure	6
2	Related Work	7
2.1	Agent Models for Driving Simulators	7
2.2	Controlling Agents	8
2.3	Conclusion Chosen Model	9
3	Preliminaries	11
3.1	BDI	11
4	Method	13
4.1	RoadNet	13
4.2	Belief base	14
4.3	Plans	15
5	Validation	16
5.1	Motivation	16
5.2	Questionnaire design	16
5.3	Questionnaire design motivation	20
5.3.1	Personal information motivation	20
5.3.2	Usage of movie clips motivation	20
5.3.3	Question design motivation	21
5.3.4	Movie clip content	21
6	Analysis	24
6.1	Biased intersection	24
6.2	Leaving traffic light	26
6.3	Approaching traffic light	28
6.4	Highway left	29
6.5	Highway right	31
7	Implementation	33
7.1	Alterations to BDI	33
7.2	Roadnet	34
7.3	Path planning and route following	36
7.4	Belief Base	37
7.4.1	Vehicle	37

7.4.2	Driver	37
7.4.3	External	38
7.5	Plans	38
7.5.1	Follow other	39
7.5.2	Lane changing	39
7.5.3	Handle intersections	40
8	Conclusion	43
8.1	Discussion	43
8.2	Contribution	43
8.3	Future work	44
A	Questionnaire Example	45
B	Analysis Graphs	46

1. Introduction

Advantages of learning to drive in a driving simulator instead of in a real car are numerous: A simulator is safer, costs less, is more environmentally friendly and is proven to be more efficient than regular training [11]. Another advantage of driving simulators that has had little research so far, is the potential for control of the background traffic. Driving instructors in driving lessons on the real road cannot influence their surroundings; students need to drive around if the instructor wishes to encounter a specific situation. A student that has difficulty with vehicles coming from the right could encounter vehicles from the right at every intersection. These specific situations could be created when needed in a simulator.

Further benefits might be achieved by taking the students' overall skill level into account. Students' skill level can be categorized with driving skill and their willingness to take risks. Novice students usually start with a low level of driving skill. The goal of the lessons is to increase that level to high or very high. To better accommodate novice drivers the surrounding traffic could be made more docile, keeping more distance, driving slower and behaving exactly according to the traffic rules. More advanced students could encounter more realistic traffic that mimics real-world traffic. The willingness to take risks varies per student and is an important when it comes to accidents amongst novice drivers [12]. Slower driving traffic could be used to encourage cautious students to take more risks. On the other hand, fast and aggressive traffic could be used for students who take too much risk.

According to De Winter et al. [11], the greatest challenge of driving simulators is creating a simulator that mimics the real world. Realism should therefore not be forgotten when giving users more control over the background traffic. The traffic behaviour should lie within the behavioural spectrum of real world vehicles. Individual vehicles, or agents, are allowed to drive fast or break traffic rules but they are not allowed to drive through objects or disappear suddenly. This has a consequence for the handling of the human driver, who could behave in an unpredictable manner. Similar to the regular behaviour of the agents, their behaviour to deal with the human driver must also lie within the behavioural spectrum of real-world traffic. The agents should therefore be highly autonomous to deal with complex traffic situations and unpredictable human drivers. However, the agents should still be able to follow orders to allow the high levels of control needed from an educational standpoint. We facilitate all these wishes by introducing a realistic, flexible and adaptable agent model for the background traffic of a driving simulator. Our model is realistic because the behaviour falls inside the scope of real-world traffic. It is flexible since it supports a wide range of environments, and it is adaptable because it displays a wide range of behaviour. A control system that guides the agents, and

configuring the agents to support students during driving lessons falls outside the scope of our thesis.

1.1 Research Questions

Considering our wish of creating a realistic, flexible and adaptable agent model for the background traffic of a driving simulator, it is possible that there are no existing models that can be directly built upon. Although the field of traffic simulation is large, the field of agent models for driving simulators is relatively small. Considering these factors, we present the following research questions:

1. How can we improve the realism, flexibility and adaptability of the background traffic for a driving simulator?
 - (a) Is there a model that can be built upon?
 - (b) If not, which model or technique comes closest?
2. What model supports the factors of question 1?
3. How can we validate this method for use among driving school students?

1.2 Structure

Chapter two provides an overview of the related work to this thesis. In chapter three we will discuss the work by Rao and Georgeff [29, 28] in more detail as it is directly related to our own method. Chapter four is a description of the theory behind our method. Chapter five explains how we will validate our method followed by the results of that validation in chapter six. Chapter seven describes our implementation and any alterations to the theory of chapter four. Lastly, we give our conclusion and possibilities for future work in chapter eight.

2. Related Work

As Bazzan et al. explain in their review on agent-based traffic simulation [2], the field is large and encompasses many applications related to traffic. It is used for transport optimization, traffic jam analysis, traffic flow, and driving models. Driving models, a survey of which can be found in Kesting et al. [23], aim to create realistic driving behaviour that mimics human behaviour. Kesting et al. [23] introduce the notion of the driver-vehicle model, which is a reference to the idea of having agents with human and vehicular properties. However, these models do not take a human driver into account because they only focus on the simulation of traffic consisting entirely of agents.

2.1 Agent Models for Driving Simulators

Driving simulators are designed to imitate a real vehicle and are operated by a human driver. The driver can perceive the environment through the use of one or more screens, usually to simulate the windows of a real car. In that environment, the traffic surrounding the human driver takes a leading role. Most models that simulate these traffic agents focus on creating realistic and human-like behaviour. Al-Shihabi and Mourant [32] use fuzzy variables and fuzzy if-then rules to simulate human traffic behaviour. However, the method is limited to two-lane highways, with future work for multi-lane highways. The authors state that the addition of other driving environments requires too much work on several levels in their framework, a common argument against the use of fuzzy logic. Demir and Çavuşoğlu [13] use Hierarchical Concurrent State Machines to simulate the background traffic. The agents can display different behavioural styles such as slow, normal and fast driving styles. Furthermore, the agents are able to display aggressive and rule-breaking behaviour, for example tailgating and running through red lights. However, the authors do not state how the agents deal with the human driver, or whether they do at all. It is also not clear how the different driving behaviours are incorporated into the model, making it hard to judge the actual merits of the method.

Seele et al. [31] base their model on psychological factors, giving a personality profile to each agent. The agents reason about their environment and make decisions to advance their goals. Unfortunately, it is not mentioned how many environments the agents can operate in, since their testing environment is only limited to a single intersection. Furthermore, it is unclear how agents deal with the decision making process, incorporate the personality, or deal with conflicting decisions. Lacroix et al. [24] use behavioural patterns based on configurable parameters for their agent model. Different behaviours can be introduced by selecting the right percentages for each of the patterns, or initiating an agent of a certain pattern with different

parametric values. These behavioural pattern configurations can also be drawn from real or simulated data for increased ease and control. One of the drawbacks of the model is the lack of deviant behaviour for urban traffic, an issue that is mentioned as future work.

The agent model of Olstam et al. [26] lacks variety in behaviour or environment but tackles the problem of when and where to generate the vehicles. Their method employs a different model for agents close to the human driver than for agents that are further away. The agents themselves are generated outside the view of the human driver. However, their behaviour is limited and based on equations, making it difficult to add new behaviour with new parameters. Nevertheless, the basic theory of only generating agents outside the view of the human driver is useful to consider. Yin et al. [38] handle human unpredictability with regard to turning signals. Their model sends signals from the human controlled vehicle to the agents to indicate the intention of the human driver. In the event that the human driver does not use his turning signals properly, the model tries to predict the signal based on the location of the human driver. For example, when the human driver is at the rightmost lane at an intersection, a right turning signal is given to surrounding agents. However, not all behaviour can be captured in this manner and the more complex behaviour of human drivers remains difficult to predict.

2.2 Controlling Agents

Driving simulators for driving schools would benefit greatly from more control over the agents, both to create general traffic as well as specific situations. Unfortunately, such levels of control are usually not the focus of traffic agent models. Demir and Çavuşoğlu [13] for example mention briefly that their existing driving simulator, TRAFIKENT, can generate vehicles through a ‘Vehicle Generation Unit’ or by the directives of the ‘Traffic Director’. The ‘Vehicle Generation Unit’ can randomly generate vehicles according to a predefined distribution function. How the ‘Traffic Director’ influences the generation of vehicles is not mentioned. However, it should be possible to create specific situations with the ‘Traffic Director’, as the name and description is similar to papers about controlled events. Seele et al. [31] make use of the IDM, Intelligent Driver Model, which allows for the setting of several parameters per individual agent. However, with IDM it is probably not possible to create specific situations without extensions.

In both models, the control of the agents is handled by a dedicated controller system. Although the working of such a controller system is outside the scope of this thesis, the agent model does need to facilitate extensive levels of control. To this end, several papers that deal with controlling agents have been examined.

The first one is from Olstam et al. [27], and it aims to combine autonomous vehicles and controlled events. The analogy used is a theater play: the stage is first prepared and when the agents are ready the play starts. The largest problem is to not let the human driver notice the preparations, which is called the play preparation problem. Agents can transition between being fully autonomous outside the stage and fully controlled when on the stage. Using a theater play as metaphor is not new. Wassink et al. [36] use the movies to explain a dynamic scenario generation algorithm for driving simulators.

While outside the scope of driving simulators, both Si et al. [34] and Shoulson et al. [33] give interesting views from the field of interactive storytelling. Both papers explain methods to guide a user and the agents surrounding the user to tell a story. Shoulson et al. [33] argue that agents should not be allowed to grow too complex. High-dimensional agents are difficult to handle and process. Letting the controller do all the work is also undesirable for roughly the same reasons. This can be directly applied to the agents in driving simulators, who should not be relied upon to create the desired scenario on their own. Instead, they should be given a path through which the agent would automatically fulfill its part for the desired scenario. When driving along this path, the agent can operate autonomously, which allows the agent to deal with sudden changes or irregularities.

Si et al. [34] try to predict what the user would do in an interactive story, not unlike Yin et al.[38] who have the same goal, only for driving simulators. Si et al. [34] tackle this problem by having the controller and the agents keep track of the user through a set of beliefs, reasoning about what they think the user thinks. This model is called a decision-theoretic goal-driven agent, with a personality and motivation as the agents goals. This way of thinking has its origin in the Belief Desire Intention (BDI) model, which is an important area of research in AI [28]. The belief set of an agent contains information about itself, the environment and other agents. Desire contains the goals of the agents, what the agent wants to achieve. Intentions describes how the agent plans to achieve those desires. The model is powerful and has been used countless times since Rao et al. [28]. The use of BDI in traffic simulation is both mentioned in the overview by Chen & Cheng [9] as well as in the overview by Bazzan and Klügl [2]. However, the method seems to be absent when it comes to using it for the background traffic in a driving simulator. One of the reasons for this is mentioned by Rossetti et al. [30] who implement a BDI model for traffic simulation but argue that the computational load might be too high when using a large amount of agents. This problem might be less apparent in a driving simulator where there is no need for a large and intricate traffic simulation with dozens or maybe hundreds of vehicles. Especially when driving in an urban environment the number of visible vehicles is limited. Another issue might be the fact that most BDI systems are implemented using special languages which are not always fit for virtual environments. More precisely, the creation of virtual environments is usually done in object oriented programming languages, which is not the standard technology for BDI systems. This problem and a possible solution is described by Dastani and Testerink [10], who translate the traditional BDI concepts to an object oriented-language.

2.3 Conclusion Chosen Model

Most agent models discussed earlier seem to lack one aspect or another. Models that display varied or realistic behaviour only work in a limited set of environments [32, 31]. Models that do work for a varied set of environments lack varied behaviour [24]. Furthermore, the human element is frequently overlooked [13]. Models that do mention the human element lack aspects in other fields, having no behavioural diversity or a wide range of environments [38]. An agent model that can display varied behaviour, operates in a wide range of environments and takes the human driver into account would be a significant contribution to the field.

Almost every discussed model uses a different technique or approach when modelling the behaviour, but BDI is not used in any of them. The reasons for not using BDI are clear: BDI has high computational load and lacks OO-language support. However, the complication of autonomous agents that remain highly controllable while being able to handle an unpredictable human student gives reason to explore BDI as model for the background traffic in a driving simulator.

3. Preliminaries

As the BDI model is directly related to our own method, we summarize previous research on this topic.

3.1 BDI

Rao and Georgeff present their formalization of the BDI architecture in ‘Modeling rational agents within a BDI architecture [29]. However, in their follow up paper [28], they mention that their previous work is too theoretical. Therefore, they give some considerations and improvements to make BDI more fit for practical use. As done in the paper, we first explain the basis of BDI and the theoretical model, after which the improvements of Roa and Georgeff are given. Lastly, our own alterations are discussed.

At the basis of the model lie the dynamic data structures representing the beliefs, desires, and intentions of the agent. Beliefs hold the state of the environment, and can be viewed as the informative component. Desires are the objectives to be completed, and can be viewed as the motivational component. Intentions are the currently chosen course of actions, and can be viewed as the deliberative component. Complementing these data structures is an event queue holding both internal and external events. These events act as input of the system. Furthermore, update and query operations are allowed on the three data structures. Now that the data structures are established the main interpreter loop of an agent can be given:

Algorithm 1 BDI interpreter loop

```
Initialize-state();  
repeat  
  options := option-generator(event-queue);  
  selected-options := deliberate(options);  
  update-intentions(selected-options);  
  execute();  
  get-new-external-events();  
  drop-successful-attitudes();  
  drop-impossible-attitudes();  
until end repeat
```

The first step in the interpreter loop is the option generation based on the event queue. Next, the deliberation step selects a subset of these options to be incorporated in the intentions data structure. After the intentions are updated, any actions produced by the intentions are executed. New external events are added to

the queue, while internal events are added when they occur. The last step is the modification of the intention and desire structures by dropping all successful and impossible desires and all satisfied and unrealizable intentions.

Roa and Georgeff give three important alterations that are meant to make the system more practical. The first alteration is to only represent beliefs about the current state of the world. Furthermore, these beliefs are expected to change over time. Second are the introduction of plans, which hold both the information to achieve a new state as well as the options available to the agent. Each plan has certain subgoals that need to be achieved for the plan to be executed successfully. Plans also contain certain conditions that need to be fulfilled before the plan can be chosen as option. These conditions can either be invocation conditions or pre-conditions. The invocation condition specifies an event that acts as trigger while a pre-condition specifies the situation that must be true before the plan can be executed. The third alteration is the use of a conventional run-time stack of hierarchically related plans. This stack implicitly represents each intention that the system forms by adopting certain plans. Multiple stacks can co-exist and they can be suspended until some condition occurs, or ordered for execution. Unfortunately, these stacks are not thoroughly explained and are not further mentioned by Roa and Georgeff. Apart from these three alterations, the authors make a suggestion for the interpreter loop, which they otherwise leave completely intact. That suggestion is an extra procedure that delays posting any events on the event queue regarding intentions until the end of the interpreter loop. By implementing this delay, the system can determine which changes need to be noticed by the option generator. This not only results in a faster system, but also allows for various levels of commitment, which results in different behaviours.

4. Method

In this section we propose our agent model for the background traffic of a driving simulator. The following sections explain the theory and ideas behind the model.

4.1 RoadNet

Multiple descriptions exist for road networks [8, 14]. Both are trying to standardize the logical description of road networks. RoadXML has its origin in the French car industry and is backed by the former French national institute for transport and safety research.

OpenDrive on the other hand has its origin in the German car industry. After the launch in 2006 other German car companies joined, as well as German, Dutch and Swedish research institutes. Like RoadXML, the description of OpenDrive is in XML. Both contain an extensive description of all features of a road network.

Such an extensive description is not needed for our method, which requires something much more lightweight and flexible. The purpose of a standard is to be used by multiple products, therefore the product must either adapt to the standard or the standard incorporates all aspects of the products it supports. Since the exploration of BDI for background traffic agents in a driving simulator is new, it might not be wise to adapt it to an existing standard. Furthermore, it cannot be expected of these standards to incorporate the requirements of a new type of driving simulator agent. Therefore, we have created our own logical description of a road network that is more suited for the agents using it. However, the workings and descriptions of both standards can work as an inspiration for our own description.

Our *RoadNet*, which will be explained more thoroughly in Section 7.2, is a collection of straight lines in 3D space that represent real-world traffic lanes. We assume that these lines lie in the centre of such a traffic lane. Curved lanes can be created by connecting several shorter straight lines at increasing angle. Since the length of a line does not have any limitations, it is possible to accurately follow the centre of any curve. However, this would increase the number of lines severely, making it more difficult to handle the entire collection.

Furthermore, each line stores all relevant traffic rules related to that line. This means that agents can request the correct speed limit or any other traffic rule from the road they are driving on. By storing the traffic rules for each separate line, we have absolute control on where a certain traffic rule begins by rearranging the lines on that road. However, this is also a weakness, as each time a traffic sign changes location, the lines have to be rearranged. Another limitation is a lack of any direct link to parallel traffic lanes. A single line does not know that it is part of a multi-lane

highway or if there are any parallel lanes in opposite direction.

As already explained, lines can be linked together to form chains. However, they can have more than one link on either end, creating a split. This functionality is used to create intersections which are assumed to be the only location where lines separate. The lines contained in an intersection hold the information regarding its type. For example, if it is a equal or biased intersection. By requesting the type of intersection, agents know which right of way rules apply.

Although the description of our *RoadNet* is simple and has far fewer options than RoadXML or OpenDrive [8, 14], it is complete, flexible and can be easily expanded by adding new traffic rules. Furthermore, due to its simplicity it can be used to recreate almost any real-world network.

4.2 Belief base

The Belief base of the agent contains everything it knows about the world. Information about itself, others and the road it drives on. The Belief base is split into three sections: *External*, *Driver* and *Vehicle*. *Vehicle* holds all physical vehicular information: the type of the vehicle, its dimensions, current velocity and positional information. These properties can be requested by other agents, and can be visually observed by the human player. *Driver* holds the personality and memory of the agent: the level of haste, aggression, the preferred following and tail distances, preferred acceleration and deceleration, look ahead distance, which traffic rules the agent obeys and the route the agent takes. The *Driver-Vehicle* pair can be viewed as the human driver and the car itself. What the human thinks, sees, and feels is unobservable by others. What the car does and how it reacts is observable. The same distinction is made for the belief bases of the agent. *External* holds the information on other agents, and the road network. In other agents the vehicle belief base of surrounding agents is stored. This is done to assure equality between agents and the player. By only storing the observable information, the agents do not need to make a distinction between the human player and other agents. When an agent checks for agents in their surroundings, the player is seen as another agent. From every agent the vehicle belief base is stored for future reference. This is the same for the human player; the player vehicle has a vehicle belief base that stores the same information as the vehicle belief base of an agent.

The belief base is regularly updated to keep the information it holds up to date. However, it is possible that a large belief base with many updated variables decreases performance. Therefore care must be taken when to update which variables. For example, one of the most important properties that require frequent updates is the lane on which the agent is driving. Delays here could result in an agent thinking that it is safe to cross an intersection because the agents thinks others have not arrived yet. Such situations could easily result in crashes, since they have difficulty detecting each other. However, if the *RoadNet* is large, it might be too expensive to determine the right lane every frame. Furthermore, it is also important to limit the number of variables in the belief base. As new traffic situations are introduced it is tempting to incorporate new variables into the belief base. However, that can have a severe impact on performance. Our own implementation of the Belief base,

including which variables we have used, is explained in Section 7.4.

4.3 Plans

Plans are a crucial part of the agent model. As explained in the preliminaries, a plan is a sequence of actions to achieve a goal. For vehicles, the actions that take up a sequence are limited. At the basis all vehicles have only three forms of output: accelerate, decelerate and steer. Therefore, the action that each plan returns consists of three values, a value for gas, brakes and steering. Vehicles are mostly used for transportation, therefore the most important goal of any vehicle is to reach its destination. Directly linked to that is the goal of reaching that destination safely, without crashing. These two goals are fixed for every agent, they always have a destination they try to reach and they always try not to crash. Furthermore, a subgoal of most drivers and thus agents, is to also follow the traffic rules. However, this goal is not fixed, and agents can be adjusted to ignore any or all traffic rules. Each plan that the agent has adheres one or more of these (sub)goals, depending on what the plan is for.

All plans are divided in three layers: the strategic, tactical and operational layer. This approach is common for traffic simulation [32] and is used to divide actions in traffic depending on the time window in which they occur. Plans in the operational layer take place in milliseconds, plans in the tactical layer take several seconds or a few minutes, and strategic plans a few dozen minutes or several hours. The layers are also hierarchical, plans in the operational layer have the highest priority, followed by the tactical layer, and lastly the strategic layer. Such a division guarantees that events that occur in a very small time frame are handled immediately.

Each plan is entirely free to draw information from the belief base and/or calculate new information. This is a strength as well as a weakness. By allowing such freedom it is possible to create a plan for any possible traffic situation. This makes the agent model very flexible and adaptable by either expanding or creating new plans for new situations. However, if such new plans calculate too much on their own, it is possible that the agent model becomes slow. Moreover, it is equally as easy to create a plan that 'breaks' the agent by introducing faulty behaviour or behaviour that interferes with other plans. Therefore care must be taken to the design and structure of each new plan to make sure it works. Fortunately, a faulty plan can be easily removed with no loss except the intended behaviour of the plan itself. Our own implemented plans are explained in Section 7.5. Although the number of plans might seem limited, only a small number is sufficient for basic traffic behaviour. Furthermore, plans can be build upon by other plans. This creates even more

5. Validation

Validation of the proposed agent model is done via an online questionnaire aimed at driving school students and driving instructors. The questionnaire participants are shown a series of clips displaying different traffic situations. After each clip the participant is asked to judge the behaviour of the traffic in the clip.

5.1 Motivation

According to Green Dino BV [7], student drivers can be categorized using two variables: their driving skill and their willingness to take risks. Both variables can range from *Very Low* to *Very High*. A student is ready to take the driving examination when their driving skill is High/Very High and their willingness to take risks is Medium. Any deviation from the examination state needs to be corrected. Green Dino BV [7] states that the surrounding traffic can help in achieving this state. For example, a student that is too cautious has a very low or low willingness to take risks. To achieve the medium status they need to be encouraged. In the driving simulator this might be achieved by introducing slow and/or calm traffic in the student's lessons. The student then hopefully notices that it is safe to take more risks. Students that are too aggressive and take too much risk need to be shown the error of their ways. This might be done by configuring the traffic to match their own behaviour. If the aggressive student's skill is low they might crash, indicating that they are doing something wrong. However, it is not yet known to what extent, if at all, students can recognize different behaviour in the surrounding traffic. If they do not notice that the other traffic is driving slow or aggressive, the surrounding traffic might not contribute to their lessons as intended.

5.2 Questionnaire design

Participants of the questionnaire will first be asked their age, sex, whether they are a driving instructor or student, and in the latter case how many hours they have spent on the simulator, and how many hours they have on the real road. The options for hours on the simulator range from 1 to 9 and 10+. The options for hours on the real road range from 0 to 49 and 50+. Each consecutive question, given in a random order, contains a clip and three questions about the clip. In the first question the participant is asked to rate the behaviour in the clip using a seven point scale. The option to the far left is cautious, the middle is normal and the far right is aggressive. In the second question the participant is asked to motivate their answer for question one. The third question is a Likert scale with the statement: 'The behaviour in the

clip was realistic’ and five points ranging from strongly disagree to strongly agree.

Bosnjak and Batinic [4] explain the factors that improve the willingness to participate in online questionnaires. In giving advance information, the most important to least-important factors are: information on access to email address, guarantee of feedback about the results, information on the exact aims of the investigation, complete anonymity of the answers, and lastly, the personal request by the researcher. These factors can easily be incorporated in the advance information mail since any information provided will not influence their answers. Knowledge about the aims of the questionnaire do not give clues to which movie clip is the ‘correct’ one. When it comes to the amount of time voluntarily provided for scientific online questionnaires, most people lie within the 11-15 and 21-30 minutes groups.

The introduction to the questionnaire begins by explaining that the questionnaire is about the behaviour of virtual traffic, and that they, the participants, have to assess that behaviour. It is also explained that the evaluation process is done by showing a clip of a traffic situation involving one or more cars. Before each clip is shown the exact situation is explained, including which cars are to be followed and judged. Lastly, the participant is told that he has to answer a few questions after each clip to the purpose of rating the behaviour of the cars in the clip.



Figure 5.1: The approaching traffic light scenario.

The questionnaire contains five scenarios and 22 clips. The first scenario contains a biased intersection and two cars approaching that intersection. One car is driving straight on the major road and the other car wants to enter the major road in the same direction as the first car. The clip ends when both cars are out of view. Variables to change are aggression and approach distance. This scenario contains



Figure 5.2: The leaving traffic light scenario.



Figure 5.3: The biased intersection scenario.

four clips plus two control clips for a total of six clips. The second scenario contains a red traffic light with three cars waiting for the light to turn green. After a few seconds the light turns green and the vehicles drive away. The first and third cars go right and the second car straight. The variable changed in this scenario is preferred

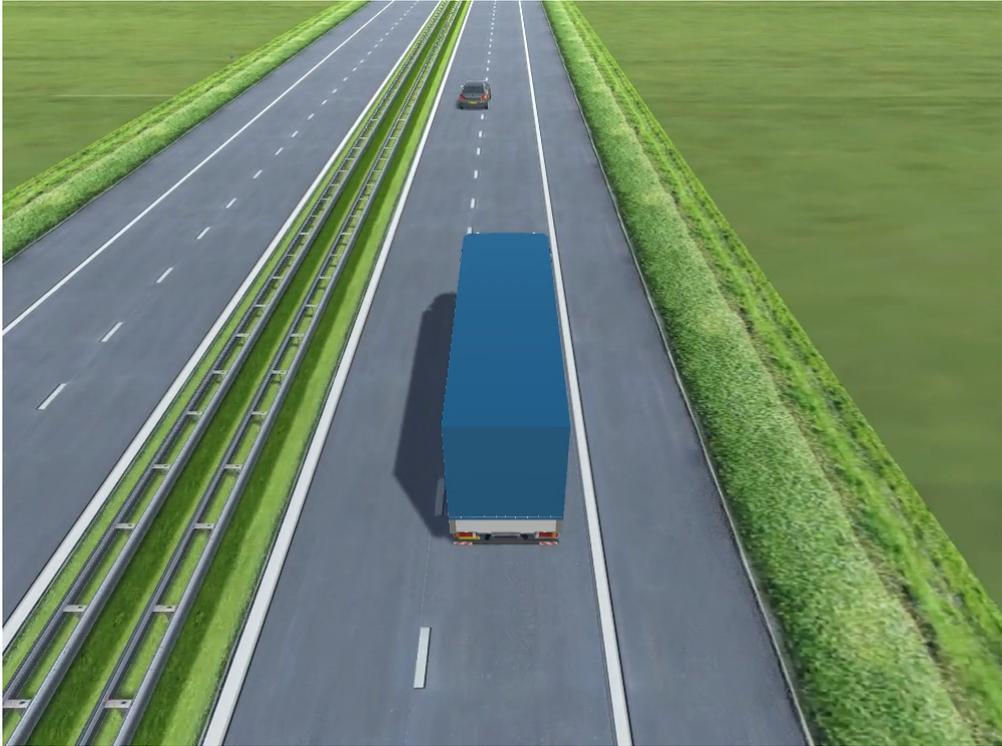


Figure 5.4: The highway right scenario.



Figure 5.5: The highway left scenario.

acceleration. This scenario contains three clips plus a control for a total of four clips. The third scenario contains a red traffic light and three cars approaching that red light. The clip ends when all cars stand still for the red traffic light. Variables changed in this scenario are following distance and preferred deceleration. This sce-

nario contains five clips plus a control for a total of six clips.

The fourth and fifth scenario contains a highway with a truck driving on the right lane and a car approaching from behind that will overtake the truck. Both scenario's contain three clips each. The fourth scenario encompasses the lane changing manoeuvre of the car towards the left lane. The changed variable is the preferred following distance and the clip ends when the car is driving alongside the truck. The fifth scenario encompasses the lane changing manoeuvre back towards the right lane. The changed variable to change is the preferred tail distance and the clip ends when the car has completed the manoeuvre.

5.3 Questionnaire design motivation

It is possible that student drivers are not capable of recognizing the different sets of behaviour, since they are not experienced enough. Therefore, both student drivers and driving instructors are asked to participate in the questionnaire. If the students fail to recognize the different behaviours and the driving instructors do recognize it, the students probably lack experience. However, the difference might be less for more experienced students with considerable road experience. Students are therefore asked for the number of hours they have spent on the simulator and on the road. Should the driving instructors also fail to recognize the different behaviours, then the agent model itself has failed.

5.3.1 Personal information motivation

The standard simulator curriculum employed by most driving schools is eight hours and students always start with at least one lesson on the simulator. However not all driving schools and/or students follow this curriculum from start to finish; some stop halfway, retake lessons, or alternate between simulator and a real car. The same applies for lessons in the real car. The number of lessons and hours needed vary greatly per student. To limit complexity, the questions on simulator and real world experience are therefore expressed in hours. For the simulator this means each number from 1 to 9 and 10+. All students have spent at least one lesson on the simulator, resulting in one hour. Eight is the length of all lessons together, plus one hour to account for retakes. It is unlikely that students spent more than 10 hours on the simulator, so everything above that is grouped together. The number of hours on the real road range from 0 to 49 and 50+. It is possible that students do not have real road experience yet, which accounts for the 0. The average number of hours it takes for a student to achieve their driver's license is 43 with a standard deviation of 6.[Reference Needed] It is unlikely that a student requires more than 50 hours of real road lessons in addition to their lessons on the simulator.

5.3.2 Usage of movie clips motivation

By using short movie clips to show the behaviour of the traffic agents, it is guaranteed that the participant is in the best position to observe that behaviour. Should the participants judge the same behaviour while driving in an actual simulator they might not notice the traffic. Furthermore, it is not guaranteed that the same situation occurs for every participant. Moreover, the organization around using sim-

ulators will result in less participants when compared to an online questionnaire. Therefore, using movie clips works best to maximize the number of participants and guarantees that the participant is able to see the intended situation and vehicles.

5.3.3 Question design motivation

After each clip, the participant is asked to answer three questions with relation to the behaviour in the clip. The number of questions is limited to three to limit the length of the questionnaire. Long questionnaires have lower response rates. There are many options for question formulation and design in questionnaires. The most unbiased answer is given with an open question. However, this might result in completely undesired and unrelated answers. In the case of traffic behaviour they might comment that the indicator lights were not on, or that the traffic is ‘beautiful’. Such answers are difficult to interpret and categorize, which can result in errors. Another option is a limited list of possible answers. However, the options might not cover the intended answer or, if the list is too large, they have a difficulty choosing one. Therefore, it is best to have a limited list of options but with a small gradation between them, to give more options than just ‘aggressive’ or ‘cautious’. The choice of a seven point scale over a five point scale is to increase diversity. From peer experience it is known that participants avoid the far left and far right choices, resulting in only three choices within a five point scale. The second question, why they have chosen their answer to question one, forces the participant to think about their answer and might indicate areas of improvement. Question two also supports question one by allowing the participant to give an open answer. The third question, indicate level of realism, uses the common Likert scale to make answering and evaluation easier. The use of a five point scale is standard for Likert scales and is therefore not increased to a seven point scale.

5.3.4 Movie clip content

Most scenarios have multiple clips where at least one variable is changed. A scenario contains a specific situation, for example a red traffic light with three approaching cars. The variable to change is, for example, the deceleration rate and is set to low, normal and high, resulting in three clips. If possible, a control clip is added that shows the same situation only in the engine of the Green Dino BV [7] driving simulator. All clips have a length between 10 and 30 seconds, and usually start just before the first car enters the screen or starts moving, and stops after the last car leaves the screen or has stopped moving. The only exception is overtaking on the highway scenario, which stops after the manoeuvre is complete.

The variables that change in the clips are: following distance, tail distance, preferred deceleration, preferred acceleration, haste, and aggression. The configuration of the variables is drawn from literature on driving behaviour, government safety advices and expertise from Green Dino BV [7]. When configuring the variables, the goal is to find a minimum realistic value, average realistic value, and a maximum realistic value. A minimum, average, and maximum is taken because these values usually match with what drivers see as cautious/slow, normal or aggressive/fast.

Following distance is the time in seconds between the vehicle in front of the agent

and the agent itself. It is also the minimum distance in meters when the speed approaches 0. The following distance in seconds as advised by the Dutch government is two [19]. However, in most real world situations this value is between one and two seconds [5], or at higher speeds as low as 0.55 seconds [16]. Since we mostly do not deal with high speeds the preferred following distances will be set at one second for the minimum value, on and a half as average and two seconds for the maximum value. For highway scenarios an extra extreme minimum value is added, namely 0.55.



Figure 5.6: Example of short following distance at the approaching traffic light scenario.



Figure 5.7: Example of long following distance at the approaching traffic light scenario.



Figure 5.8: The control of the approaching traffic light scenario.

Tail distance is the time in seconds the agent allows between itself and another agent behind it. In the model, it is mostly used for overtaking. Since no clear data

is available for this value, and it is closely related to following distance, the same values for following distance are used for tail distance. For highway scenarios this also means that the extra extreme minimum of 0.55 is added.

The preferred acceleration and preferred deceleration values are thresholds that the agent does not break when accelerating or braking. Acceleration and deceleration values differs greatly between different vehicle types and countries [1, 3, 25]. Akçelik et al. [1] study acceleration and deceleration in India. Their model is based on real world data, and shows a maximum acceleration of $2.7 m/s^2$ with an average of $1.5 m/s^2$ for a vehicle starting from $0 km/h$ and accelerating to $60 km/h$. Bogdanović et al. [3] study the acceleration at traffic signals in Serbia. Their conclusion was that under normal conditions the acceleration values lie between $1.7-2.0 m/s^2$. They purposefully did not take the extremes into account since their aim was to provide data for traffic flow planning. The authors state that 95% of the values lie between $0.87-3.26 m/s^2$ for their close measuring point, between $1.17 - 2.58 m/s^2$ for their far measuring point and between $0.71-3.4 m/s^2$ for their second study a few years later. Considering this data, the minimum acceleration is set to $1 m/s^2$, the average to $2 m/s^2$ and the maximum to $3 m/s^2$.

Akçelik et al. [1] give an average deceleration of $1.8 m/s^2$ and a maximum of $3.1 m/s^2$ for a car travelling at $60 km/h$ and coming to a full stop. Maurya and Bokare [25] not only provide their own findings but also the findings of several other studies on deceleration. Those studies, dating from 1960 to 2005, give a minimum deceleration rate of $0.28 m/s^2$ and a maximum of $4.9 m/s^2$. The most recent studies give values between 2 and $3 m/s^2$. However, those values are for sudden deceleration at a signalized intersection. The values given by the authors themselves have a maximum of $1.6 m/s^2$ and a mean of $1.2 m/s^2$ for cars travelling between $90-100 km/h$. Considering this data, the minimum deceleration rate will be $1 m/s^2$, the average will be set to $2 m/s^2$, and the maximum to $4 m/s^2$.

Aggression influences the distance perceived to others. For example, when determining the distance between itself and an intersection and another agent with the same intersection then aggression increases or lowers those distances. There are no exact values known for aggression since it is an unique value. It is therefore set according to the needs of the scenario. Only the average setting can already be given, which is 0.

6. Analysis

The completed questionnaires are analysed using a repeated measures MANOVA test with Wilks' Lambda method [15]. Each scenario is treated as a separate test, scenarios cannot be properly tested with each other due to the large differences between them. Furthermore, the behaviour score and realism score are separately tested per scenario to limit complexity. Moreover, one test is done with both students and instructors and another test with only students. The tests with both students and instructors aims to measure the effect of clip on behaviour and realism scores as well as the interaction between profession and clip. In the tests with only students, the effect of clip on behaviour and realism scores is measured as well as the interaction between hours in the simulator on clip and hours on the road on clip. The hours on simulator data have been grouped into two data sets: group one with hours from 1 to 5 and group two with hours from 6 to 10+. The hours on the road data has been grouped into five groups: group one with hours from 0 to 9, group two with hours from 10 to 19, group three with hours from 20 to 29, group four with hours from 30 to 39 and group five with hours from 40 to 50+. In all tests both age and sex are taken as covariants.

The questionnaire was sent to 1500 driving school students who have been active on the simulator in the last year. The questionnaire was also sent to 80 driving schools that employ simulators of Green Dino BV. Furthermore, one driving school actively asked simulator students to fill in the questionnaire after their lessons. In total, 182 students and instructors started the questionnaire. Roughly half of them filled it in from beginning to end. Results of participants that have stopped halfway are only taken into account when a full set of clips from one scenario was completed. This creates fluctuations in the number of participants per scenario. Each of the following sections gives the results of one of the five scenarios. As explained earlier, each scenario has four tests: one for behaviour amongst students and instructors, one for behaviour amongst only students, one for realism among students and instructors, and lastly one for realism amongst only students.

6.1 Biased intersection

The test with the behaviour score amongst both students and instructors was done with 78 subjects; of those 78, 14 were instructors and 64 were students. The effect of clip on behaviour score was significant ($F(5, 70) = 4.963$, $p < .05$) with a strong effect ($\eta^2 = .262$). The interaction between clip and profession was not significant ($F(5, 70) = .772$). The other interactions, sex and age, were likewise not significant. The pairwise comparison of each clip to all other clips reveals that only clip 1 & clip

4 and clip 3 & clip 6 are not significant.

The test with the behaviour score amongst only students had 29 participants in the first group of hours on simulator and 35 participants in the second group. The first group of hours on road had 15 participants, the second group 7 participants, the third group 9 participants, the fourth group 19 participants, and the fifth group 14 participants. The effect of clip on behaviour score was significant ($F(5, 52) = 8.237, p < .05$) with a strong effect ($\eta^2 = .442$). The interaction between clip and hours sim was not significant ($F(5, 52) = .159$). However, the interaction between clip and hours road was significant ($F(20, 173.414) = 1.684, p < .05$) with a strong effect ($\eta^2 = .137$). Likewise, the interaction between clip and age was significant ($F(5, 52) = .748, p < .05$). The other interactions were not significant. The pairwise comparison of each clip to all other clips reveals that only clip 1 & clip 4 and clip 1 & clip 2 are not significant.

In combination with the means of Table 6.1, it is revealed that there is a difference between the clips, and that the intended behaviour of the non-control clips was recognized as intended. Amongst the non-control clips, the biggest surprise is that normal far differs significantly from normal close, with normal close being seen as more cautious than normal far. Amongst the control clips, control close is seen as equally cautious to normal close. Furthermore, amongst only students there is no significant difference between control close and cautious far and amongst both students and instructors there is no significant difference between normal far and control far. Although there is no significant relation between profession and clip score, these differences seem to give an indication that there is a difference between instructors and students. However, the small number of instructors could influence the results. Another surprising result was the significant relation between hours road and behaviour score amongst students. The means of Table 6.2 reveal that the more road hours a student had the less cautious he thought the clips were. The exception is the first group which had the least road hours. However, like the instructors some groups had a very low number of subjects, which could influence the results.

clip	description	behaviour S + I	behaviour S	realism S + I	realism S
1	control close	-1.052	-1.423	-.342	-.407
2	cautious far	-2.031	-2.050	.282	.142
3	normal far	.007	.109	-1.210	-1.171
4	normal close	-.947	-1.214	-.319	-.515
5	aggressive close	1.537	1.450	-.195	-.416
6	control far	.586	.758	-.400	-.369

Table 6.1: Estimated marginal means of behaviour and realism scores for the biased intersection scenario amongst students (S) and instructors (I).

The test with the realism score amongst both students and instructors was done with 77 subjects; of those 77, 11 were instructors and 66 were students. The effect of clip on realism was significant ($F(5,69) = 2,574, p < .05$) with a strong effect ($\eta^2 = .157$). The interaction between clip and profession was not significant ($F(5, 69) = .296$). The other interactions were likewise not significant. The pairwise comparison

Group	Hours	Mean
1	0-9	-.404
2	10-19	-.630
3	20-29	-.422
4	30-39	-.293
5	40-50+	-.228

Table 6.2: Means of hours on the road for behaviour score amongst students

of each clip to all other clips reveals that only clip 3 has a significant difference with all other clips.

The test with the realism score amongst only students had 30 participants in the first group of hours on simulator and 36 participants in the second group. The first group of hours on road had 16 participants, the second group 7 participants, the third group 9 participants, the fourth group 22 participants and the fifth group 12 participants. The effect of clip on realism was not significant ($F(5,54) = 1,896$). The same applies for the interaction between clip and hours sim ($F(5, 54) = 1.084$), as well as the interaction between clip and hours road ($F(20, 180.048) = 1.593$). Furthermore, all other interactions are also not significant.

In combination with the means of Table 6.1, it is revealed that the normal far clip is seen as the most realistic of all clips. All other clips are judged equally realistic. Furthermore, most clips have a negative realism score, indicating that they are more realistic than unrealistic. The only exception is cautious far, which is seen as more unrealistic than realistic. Cautious far was also judged as the most cautious, giving rise to the hypothesis that the more extreme the behaviour, the more unrealistic it is perceived. Unfortunately, cautious far is not significantly different to any other clip. Furthermore, there is no significant effect of profession on clip. However, similar to the behaviour score tests there is a difference between the two groups. The test amongst both students and instructors shows a significant effect of clip on realism, while the test amongst only students does not show this significance.

6.2 Leaving traffic light

The test with the behaviour score amongst both students and instructors was done with 85 subjects; of those 85, 15 were instructors and 70 were students. The effect of clip on behaviour score was significant ($F(3, 79) = 6.045$, $p < .05$) with a strong effect ($\eta^2 = .187$). The interaction between clip and profession was not significant ($F(3, 79) = 2.400$). The other interactions were likewise not significant. The pairwise comparison of each clip to all other clips reveals that only clip 1 & clip 3 and clip 2 & clip 4 are not significant.

The test with the behaviour score amongst only students had 31 participants in the first group of hours on simulator and 39 participants in the second group. The first group of hours on road had 17 participants, the second group 7 participants, the third group 10 participants, the fourth group 23 participants and the fifth group

13 participants. The effect of clip on behaviour score was significant ($F(3, 60) = 3.879$, $p < .05$) with a strong effect ($\eta^2 = .162$). The interaction between clip and hours sim was not significant ($F(3, 60) = .224$). The same applies for the interaction between clip and hours road ($F(12, 159.037) = .994$), as well as the other interactions. The pairwise comparison of each clip to all other clips reveals that only clip 1 & clip 3 and clip 2 & clip 4 are not significant.

The marginal means of Table 6.3 reveal that similar to the biased intersection scenario, there is a significant difference between the clips and that the intended behaviour of the non-control clips is mostly recognized as such. This leads to the conclusion that the higher the preferred acceleration is, the more aggressive the behaviour is perceived. However, the aggressive and normal clips do not differ significantly on behavioural score. Since both means are around zero, the behaviour in the aggressive clip is seen as normal. With regards to the preferred acceleration this indicates that the highest setting was not high enough. The control and cautious clips are seen as significantly more cautious than the normal and aggressive clips, although there was no significant difference between them. Furthermore, there was again no significant effect of profession on behavioural score and no significant differences between the two groups.

clip	description	behaviour S + I	behaviour S	realism S + I	realism S
1	cautious	-1.969	-1.381	.190	-.158
2	normal	-.091	-.256	-.951	-.841
3	control	-1.738	-1.405	.492	.166
4	aggressive	-.084	.005	-.862	-.774

Table 6.3: Estimated marginal means of behaviour and realism scores for the leaving traffic light scenario amongst students (S) and instructors (I).

The test with the realism score amongst both students and instructors was done with 84 subjects; of those 84, 15 were instructors and 69 were students. The effect of clip on realism score was significant ($F(3, 78) = 3.928$, $p < .05$) with a strong effect ($\eta^2 = .131$). The interaction between clip and profession was not significant ($F(3, 78) = 1.456$). The other interactions were likewise not significant. The pairwise comparison of each clip to all other clips reveals that only clip 1 & clip 3 and clip 2 & clip 4 are not significant.

The test with the realism score amongst only students had 30 participants in the first group of hours on simulator and 39 participants in the second group. The first group of hours on road had 17 participants, the second group 7 participants, the third group 10 participants, the fourth group 22 participants and the fifth group 13 participants. The effect of clip on realism score was not significant ($F(3, 59) = 2.303$). The interaction between clip and hours sim was not significant ($F(3, 59) = 1.263$). The same applies for the interaction between clip and hours road ($F(12, 156.391) = .546$), as well as the other interactions.

Table 6.3, reveals that the cautious and control clips are seen as less realistic than the aggressive and normal clips. Furthermore, like with the behaviour scores, the cautious and control clips and the normal and aggressive clips have no significant

difference between them. Also, the normal and aggressive clips are seen as more realistic than unrealistic due to the fact that their means are negative. The cautious and control clips are seen as more unrealistic than realistic since their means are positive for students and instructors. However, there is no significant result in the test with only students, even though there is no significant effect of profession on the realism scores. This leads to the same conclusions as the previous scenarios; instructors do influence the results but this might be the result of too few subjects. When comparing the realism with behavioural scores, the extreme behaviours are again seen as more unrealistic than the normal behaviours.

6.3 Approaching traffic light

The test with the behaviour score amongst both students and instructors was done with 82 subjects; of those 82, 15 were instructors and 67 were students. The effect of clip on behaviour score was not significant ($F(5, 74) = 1.792$). The interaction between clip and profession was also not significant ($F(5, 74) = .481$), similar to the other interactions which were also not significant.

The test with the behaviour score amongst only students had 30 participants in the first group of hours on simulator and 37 participants in the second group. The first group of hours on road had 17 participants, the second group 7 participants, the third group 11 participants, the fourth group 21 participants and the fifth group 11 participants. The effect of clip on behaviour score was significant ($F(5, 55) = 4.004$, $p < .05$) with a strong effect ($\eta^2 = .267$). The interaction between clip and hours sim was not significant ($F(5, 55) = 1.907$). The same applies for the intersection between clip and hours road ($F(20, 183.364) = 1.326$), as well as the other interactions. The pairwise comparison of each clip to all other clips reveals that only clip 2 & clip 3, clip 2 & clip 5 and clip 3 & clip 5 are not significant.

In combination with the means of Table 6.4, it is revealed that unlike the previous scenarios there is a significant effect of clip on behavioural score amongst only students while this effect is absent in the test with both students and instructors. Furthermore, there is no significant effect of profession on score. This leads to the conclusion that in this scenario the instructors seem to have a dampening effect while in the previous two scenarios the instructors have an amplifying effect. As for the differences between the behavioural scores themselves: cautious short, normal average and aggressive long do not differ significantly from each other. Furthermore, these three clips are seen as more aggressive than the cautious long clip and less aggressive than the aggressive short clip, indicating that the intended behaviour of the clips was mostly recognized as such. This leads to the conclusion that the following distance does have an influence on the perceived behaviour since the cautious short and aggressive long differ significantly from their behavioural counterparts. Lastly the control clip differs significantly from all other clips and is seen as the most cautious.

The test with the realism score amongst both students and instructors was done with 78 subjects; of those 78, 15 were instructors and 63 were students. The effect of clip on realism was not significant ($F(5,70) = .326$). The interaction between clip and profession was also not significant ($F(5, 70) = .712$), similar to the other

clip	description	behaviour S + I	behaviour S	realism S + I	realism S
1	cautious long	-.229	-.482	-.790	-.693
2	cautious short	.369	.260	-.552	-.296
3	normal average	.404	.515	-.858	-.809
4	control	-1.283	-1.614	-.069	.197
5	aggressive long	.757	.880	-1.028	-.952
6	aggressive short	1.268	1.408	-.708	-.298

Table 6.4: Estimated marginal means of behaviour and realism scores for the approaching traffic light scenario amongst students(S) and instructors (I).

interactions which were likewise not significant.

The test with the realism score amongst only students had 28 participants in the first group of hours on simulator and 35 participants in the second group. The first group of hours on road had 14 participants, the second group 7 participants, the third group 11 participants, the fourth group 21 participants and the fifth group 10 participants. The effect of clip on realism was not significant ($F(5,51) = 1,668$). The same applies for the interaction between clip and hours sim ($F(5, 51) = .768$) and the interaction between clip and hours road ($F(20, 170.098) = 1.059$). Furthermore, there is no significance on the other interactions.

Since none of the effects in the realism tests are significant, it is difficult to draw any proper conclusions. However, the means of Table 6.4 do give an indication that the control clip is seen as less realistic than the non-control clips. Since this effect was one of the main research questions of our work, we will mention the pairwise comparison of the control clip with the non-control clips. Amongst both students and instructors there is a significant difference between the control clip and the normal and aggressive long clips. Amongst only students there is also a significant difference between the control clip and the cautious long clip. These results add to the evidence that the control clips are seen as less realistic than the non-control clips but this addition is slim due to their origin.

When comparing the realism scores with the behavioural scores we again see that the most cautious clip is the most unrealistic. However, the most aggressive clip does not share this tendency and is perceived as equally realistic as the other non-control clips.

6.4 Highway left

The test with the behaviour score amongst both students and instructors was done with 85 subjects; of those 87, 15 were instructors and 72 were students. The effect of clip on behaviour score was significant ($F(2, 82) = 5.489$, $p < .05$) with a strong effect ($\eta^2 = .118$). The interaction between clip and profession was not significant ($F(2, 82) = .362$). The other interactions were likewise not significant. The pairwise comparison of each clip to all other clips reveals that each clip has a significant difference with all other clips.

The test with the behaviour score amongst only students had 32 participants in the first group of hours on simulator and 40 participants in the second group. The first group of hours on road had 17 participants, the second group 7 participants, the third group 13 participants, the fourth group 22 participants and the fifth group 13 participants. The effect of clip on behaviour score was significant ($F(2, 63) = 9.295$, $p < .05$) with a strong effect ($\eta^2 = .228$). The interaction between clip and hours sim was not significant ($F(2, 63) = .985$). The same applies for the interaction between clip and hours road ($F(8, 126) = .694$), as well as the other interactions. The pairwise comparison of each clip to all other clips reveals that only clip 1 & clip 2 are not significant.

The marginal means of Table 6.5 reveal that amongst students and instructors the shorter the following distance is, the more aggressive it is perceived. The least aggressive clip is the clip with average following distance, although it does not differ significantly from the minimum clip amongst only students. Furthermore, the average clip is seen as normal, creating a lack of any obvious cautious clip. The extreme minimum and minimum clips do differ significantly and the former is judged as most aggressive. Although there is no significant effect of profession on behavioural score, there is again a slight difference between the groups when it comes to number of significant results.

clip	description	behaviour S + I	behaviour S	realism S + I	realism S
1	minimum	.414	.321	-.827	-.838
2	average	-.034	.044	-.862	-.871
3	ex. minimum	1.764	1.666	-.508	-.215

Table 6.5: Estimated marginal means of behaviour and realism scores for the high-way left scenario amongst students (S) and instructors (I).

The test with the realism score amongst both students and instructors was done with 87 subjects; of those 87, 15 were instructors and 72 were students. The effect of clip on realism score was not significant ($F(2, 82) = .105$). The interaction between clip and profession was not significant ($F(2, 82) = 2.429$). The other interactions were likewise not significant.

The test with the realism score amongst only students had 32 participants in the first group of hours on simulator and 40 participants in the second group. The first group of hours on road had 17 participants, the second group 7 participants, the third group 13 participants, the fourth group 22 participants and the fifth group 13 participants. The effect of clip on realism score was not significant ($F(2, 63) = .890$). The interaction between clip and hours sim was also not significant ($F(2, 63) = .025$). The same applies for the interaction between clip and hours road ($F(8, 126) = 1.349$), as well as the other interactions.

Since none of the results are significant and there are no control clips in this scenario all clips are deemed equally realistic or unrealistic. With the mean scores of Table 6.5 we can observe that the scores of all clips is negative, meaning that the clips are seen as more realistic than unrealistic. Furthermore, there was no significant effect of profession on realism score and no difference in significant results between the

groups.

Due to the fact that there are no significant results for realism and no control clips, a comparison between realism and behavioural scores will prove difficult. The most aggressive clip does have a lower realism score but we cannot say anything about whether this difference is significant or not.

6.5 Highway right

The test with the behaviour score amongst both students and instructors was done with 85 subjects; of those 85, 15 were instructors and 70 were students. The effect of clip on behaviour score was significant ($F(2, 80) = 3.435$, $p < .05$) with a medium effect ($\eta^2 = .079$). The interaction between clip and profession was not significant ($F(2, 80) = .231$). The other interactions were likewise not significant. The pairwise comparison of each clip to all other clips reveals that each clip has a significant difference with all other clips.

The test with the behaviour score amongst only students had 32 participants in the first group of hours on simulator and 38 participants in the second group. The first group of hours on road had 17 participants, the second group 8 participants, the third group 10 participants, the fourth group 23 participants and the fifth group 12 participants. The effect of clip on behaviour score was significant ($F(2, 61) = 5.716$, $p < .05$) with a strong effect ($\eta^2 = .158$). The interaction between clip and hours sim was not significant ($F(2, 61) = 1.171$). The same applies for the interaction between clip and hours road ($F(8, 122) = .621$), as well as the other interactions. The pairwise comparison of each clip to all other clips reveals that only clip 1 & clip 2 are not significant.

In combination with the means of Table 6.6, it is revealed that the shorter the tail distance the more aggressive the behaviour is perceived. However, unlike the highway left scenario, the scores are all in the negative, lacking any obvious or even slightly aggressive behaviour. Even the extreme minimum clip is seen as normal where it was intended as aggressive. Furthermore, there was again no significant effect of profession on behavioural score and there are no differences on the number of significant results between the groups.

clip	description	behaviour S + I	behaviour S	realism S + I	realism S
1	ex minimum	-.035	-.075	-.969	-1.071
2	average	-1.292	-1.322	-.123	-.207
3	minimum	-.822	-.776	-.394	-.609

Table 6.6: Estimated marginal means of behaviour and realism scores for the highway right scenario amongst students (S) and instructors (I).

The test with the realism score amongst both students and instructors was done with 83 subjects; of those 83, 15 were instructors and 68 were students. The effect of clip on realism score was significant ($F(2, 78) = 3.234$, $p < .05$) with a medium effect ($\eta^2 = .077$). The interaction between clip and profession was not significant ($F(2, 78) = .649$). The other interactions were likewise not significant. The pairwise

comparison of each clip to all other clips reveals that only clip 2 & clip 3 are not significant.

The test with the realism score amongst only students had 31 participants in the first group of hours on simulator and 37 participants in the second group. The first group of hours on road had 17 participants, the second group 8 participants, the third group 9 participants, the fourth group 22 participants and the fifth group 12 participants. The effect of clip on realism score was not significant ($F(2, 59) = 2.931$). The interaction between clip and hours sim was also not significant ($F(2, 59) = 1.713$). The same applies for the interaction between clip and hours road ($F(8, 118) = 1.745$), as well as the other interactions.

From the results we can observe that there is a significant effect of clip on realism amongst students and instructors which lacks amongst only students. Furthermore, like the previous scenarios, there is no significant effect of profession on realism score. When taking the means scores of Table 6.6 into account, it is revealed that the extreme minimum clip is seen as more realistic than the other clips. However, since all clips have a negative realism score they are seen as more realistic than non realistic.

Combining the realism with behavioural scores reveals that the clip with the most normal behaviour, extreme minimum, is seen as the most realistic. The other two clips display cautious behaviour and are seen as less realistic than the clip with normal behaviour.

7. Implementation

This section explains the implementation of our model and any alterations to the theory. The model was implemented using Unity3D, a game engine and game development platform [35]. Due to our choice of game engine some alterations have been made to the theory as explained in Section 4. As far as we know, these changes have yielded no apparent limitations. Moreover, the advantages of using this game engine far outweigh any alterations that have been made because of it.

7.1 Alterations to BDI

Our agent model roughly uses the interpreter loop as presented by Roa and Georgeff [28]. However, due to our choice of game engine, some changes in order have been made. Most important difference is that any new events spawned by the actions of the agent and/or the world are only viewed in the next iteration of the system. This forces the execute method to the end of the interpreter loop and the get-new-external-events method to the start of the loop. Furthermore, although the first alteration by Roa and Georgeff [28] mentions changing beliefs, they do not mention when these beliefs are updated. Therefore, the get-new-external-events method is changed to *updateBeliefs* to explicitly handle all changes in the world. This also drops the event-queue, which is not necessary anymore. This concurs with the suggestion of an extra procedure to delay events to the end of the interpreter loop. By handling all events at once, there is no need for an event-queue, although any internal events are still handled when prompted. The rest of the loop remains intact, although some names are changed as will be explained shortly.

Another alteration is that the role of plans has been made more prominent. Roa and Georgeff introduce plans as a means to both hold the information to achieve a new state, as well as the options available to the agent, effectively replacing intentions. In our method, plans are the only options available to the agent, which result in the option generator function to be the plan generator function, called *generateNewPlans*. This in turn has a consequence for the deliberate options function, which in our method deliberates over the newly generated plans. This function is therefore named *filterNewPlans* and also incorporates the functionality of the update-intentions procedure. Furthermore, it is obvious that most vehicles are made for transportation, to transport something or someone from A to B. Therefore, any agent in the model has only one desire at its roots: to reach its destination. All additional desires are an extension or addition to this; avoiding collisions, adhere traffic rules, or limit fuel consumption. This desire can be fulfilled by a single intention and thus plan: the description of a route to get from A to B, where B is the desired location, and A is the current or starting location. All other plans are

built upon this route and fulfill other desires: avoid collisions, or following traffic rules. For the interpreter, this results in an alteration of the drop-successful-attitude and drop-impossible-attitudes, which can be seen as a remove successful and failed plans. Having one desire and describing that desire as a single plan concurs with the intended usage of the agents: to act as background traffic.

7.2 Roadnet

In this section we give our implementation of the *RoadNet* as explained in Section 4.1. At the basis of the *RoadNet* lies the Lane, a data structure that acts as a description of the real world traffic lane. A single Lane has several vectors that describe its location. These vectors are: a start vector \vec{S} , end vector \vec{E} and a tangent vector \vec{T} for curved lanes. In the case of a curved lane, \vec{T} describes the single control point of a quadratic Bézier curve. Each curved lane is translated into several smaller straight lanes through interpolation with a size that can be determined by the user. This makes calculations with the route following easier since calculating intersections between a line and a circle is easier than between a circle and a Bézier curve. The driving direction is indicated by the directional vector $\vec{D} = \vec{E} - \vec{S}$.

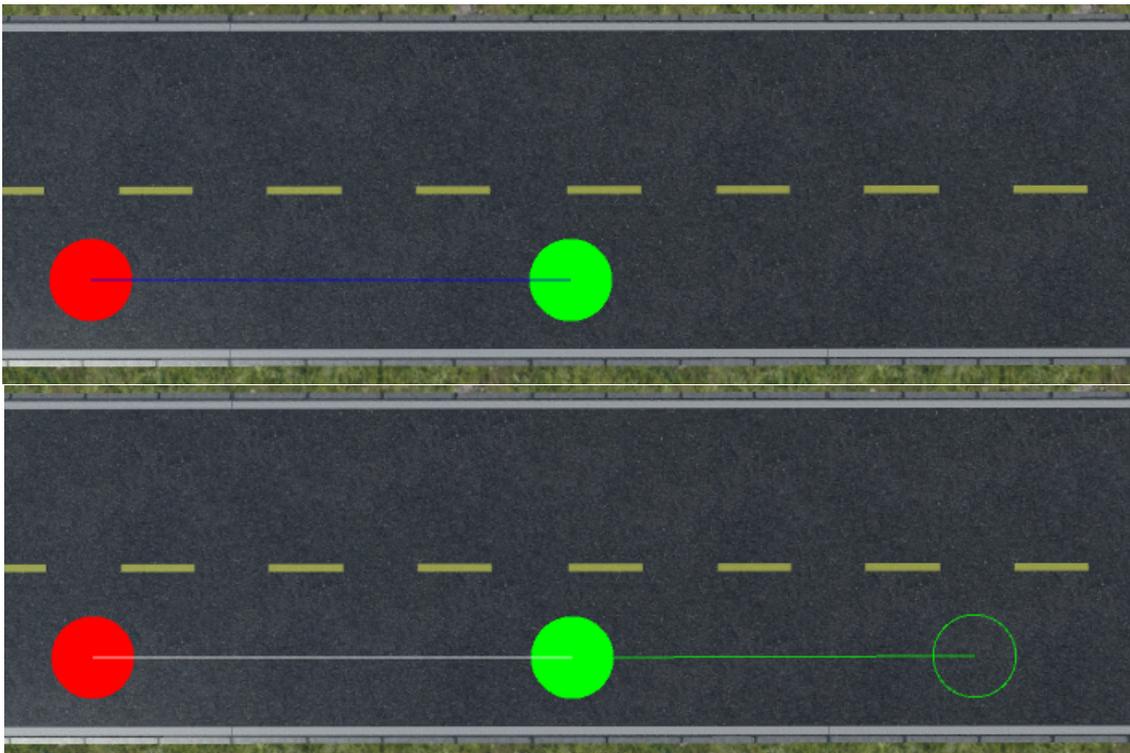


Figure 7.1: The creation of a Lane in the *RoadNet*, the red circle is the \vec{S} vector and the green circle the \vec{E} vector.

A Lane L can have none or several connecting lanes in either direction. The lanes in line with the driving direction, named next lanes, have the same coordinates for \vec{S} as L does for \vec{E} . The lanes in the opposite direction, named previous lanes, have the same coordinates for \vec{E} as L does for \vec{S} . Most lanes have only one lane in both

lists, multiple lanes only occur at intersections.

Furthermore each Lane has a speed limit, advisory speed limit, lane group number, list of allowed vehicle types and a list of traffic signs. Most of these variables reflect the traffic rules in the Netherlands [20]. Only the lane group number has no direct relation to traffic ruling.

Both speed limits are in km/h since that is the standard measurement for speed limits in the Netherlands. All lanes between intersections have the same lane group number to indicate that they are part of the same stretch of road. Agents can use that number to see if another agent is driving on the same road as they are. The list of allowed vehicles limits agents in the roads they can use for path planning. For example, an agent with the vehicle type ‘truck’ can only drive on lanes that allow trucks. Lastly, the list of traffic signs contains all traffic signs that apply to that lane. There are some restrictions and deviations for traffic signs that will be explained shortly. Furthermore, the *RoadNet* has no visible representation in the environment, traffic signs need to be placed by hand to alert the human player.

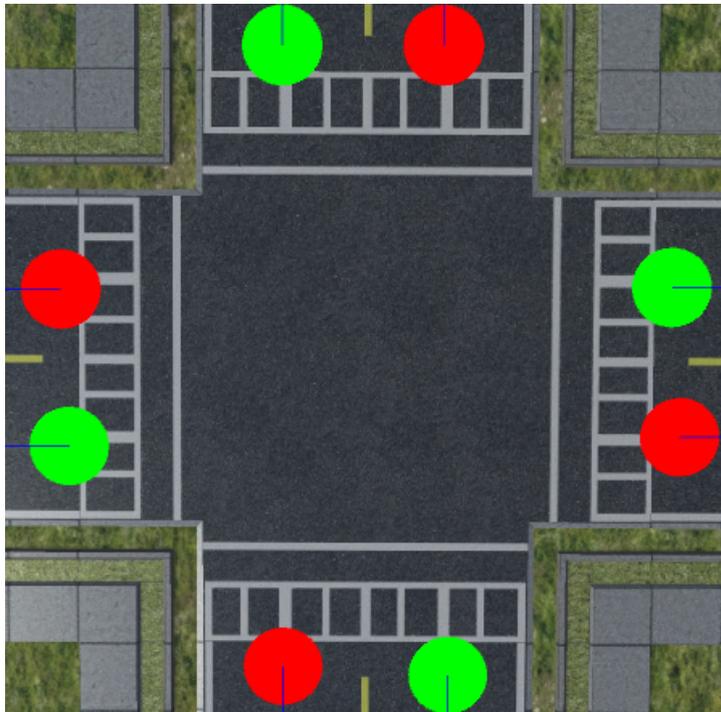


Figure 7.2: The first step in the creation of an Intersection in the *RoadNet*, the creation of the incoming (green) and outgoing (red) lanes.

Intersections are stored as a separate data structure. Each intersection has a center position vector \vec{CI} , a list of intersection lanes and a variable indicating the type of intersection. Types currently supported are equal intersections, biased intersections, traffic light intersections, and roundabouts. Intersections with traffic lights also contain a traffic light controller. \vec{CI} is the average of all \vec{S} and \vec{E} vectors of the outgoing and incoming lanes respectively. An intersection lane is a special kind of lane that is only used in intersections. It contains all information of a normal Lane plus the Intersection it falls under, the direction (straight, left, right or u-turn) and a center position vector \vec{CL} . \vec{CL} is the average of S and E for straight lanes, and

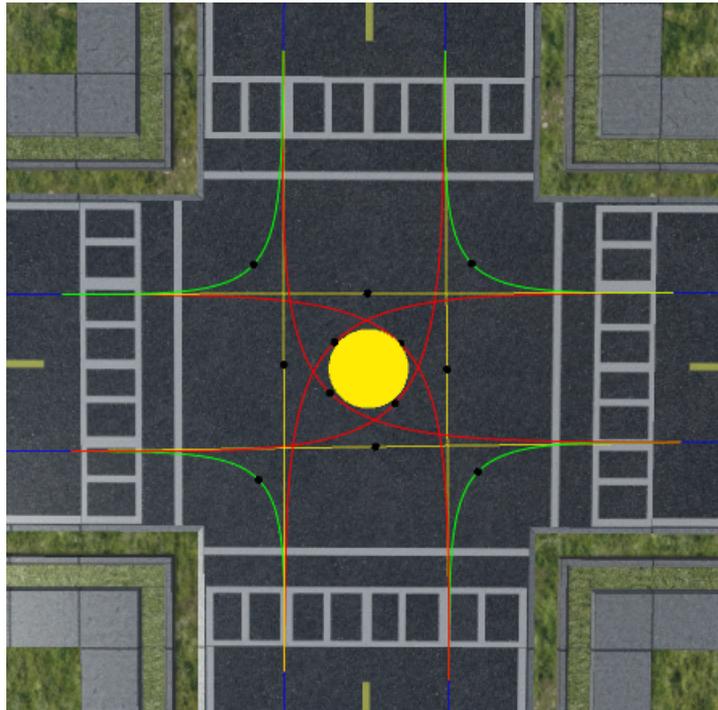


Figure 7.3: The result of an automatic generation of an Intersection in the *RoadNet*, the intersection lanes are colored according to their direction. The smaller black dots indicate the location of \vec{CL} and the yellow circle indicates \vec{CI} .

the center interpolation for the Bézier curves. All smaller straight lanes originating from a Bézier curve share the same \vec{CL} of that curve. \vec{CL} is used in collision checks on intersections. Intersection lanes start at the stop line, the white line that tells vehicles where to stop. The end position is parallel to the start position of the lane in opposite direction. Vector \vec{T} of curved lanes is the intersection between the incoming and outgoing lane.

7.3 Path planning and route following

Apart from holding information, the *RoadNet* is also used in path planning and route following. The actual path planning is done using A*. Each lane is handled as a node with \vec{E} as the position of the node. Neighbouring nodes are drawn from the next lanes list of each lane. When a correct path is found, A* returns a list of connecting lanes that act as the *Route* of the agent. Since the route is stored as a collection of lanes, all data from the *RoadNet* is stored within this collection. Therefore, many information requests can be handled by using the Route of the agent instead of using the far larger *RoadNet*, which increases performance.

Since path planning and route following is not the focus of our work we rely on the Indicative Route method as described in [22]. Given a lane L and a circle C around the agent's center position with radius r , the agent steers towards the intersection i of L and C that lies in the driving direction of the agent. i can be seen as the attraction point that draws the agent towards it. L is determined by first taking the current lane the agent is driving on and then continuing on the route till either an intersection is found or the distance from the agent to Vector \vec{E} of the lane is

greater than r . Should in the latter case no intersection be found, the agent steers towards the last known intersection. r is equal to the velocity of the agent and can be influenced to allow smoother or more abrupt steering. The minimum of r is half of the vehicles length plus one meter. This guarantees that the agent is always steering forward, even when standing still.

The steering output is the angle between the forward vector of the agent and i - the current location vector of the agent. That angle is translated to a value between -1 and 1, with -1 indicates steering to the far left and 1 to the far right.

7.4 Belief Base

In this section we present our implementation of the belief system including what variables we store, how we update them and what the most important variables are used for.

7.4.1 Vehicle

Vehicle contains all observable properties and can be categorized in fixed and dynamic variables. Fixed variables include vehicle type and vehicle dimensions. The dynamic properties which require constant updating are: current location, current Lane, heading, velocity, acceleration, minimum braking distance. Vehicle type is primarily used in path planning as each Lane contains a list of accepted vehicle types. An agent with a vehicle type not in that list will dismiss that lane as possible path. One of the most important properties is the current lane. As already mentioned in Section 4.2, it is too expensive to iterate over all lanes in the *RoadNet* to find the current lane the agent is driving on. Therefore we only search on the path the agent has planned for itself.

Minimum braking distance is the distance it takes for the vehicle to come full stop after braking at full strength at the current velocity. It is calculated using $v^2/(2 * tireFriction * 9.81)$ with v as the velocity and *tireFriction* being the friction coefficient between the tires and the road. This distance measure is often used to determine when to start braking to stop at a particular point. It is also very important to prevent collisions with other agents.

7.4.2 Driver

Driver contains the personality of the agent and other internal properties that are not observable by others. Similar to *Vehicle* there are fixed and dynamic properties. The fixed variables are related to personality while the dynamic variables are internal properties that are necessary for certain plans to function. The personality parameters are haste, aggression, preferred following distance, preferred minimum tail distance, a list of rules the agent obeys and lookahead distance. The dynamic variables are the the agent is taking and the intersection and intersection lane it is approaching. Although it is observable that an agent is approaching an intersection, it is not known which lane it is going to take.

Most of the personality variables have already been discussed in Section 5.3.4. One of the variable that has not been discussed is the list of rules the agent obeys. Although the list is currently still short it can be expanded to incorporate any traffic

rule. Plans that are designed to abide a certain traffic rule can then be ignored by the agent that is set to not follow that particular traffic rule. Lookahead distance is the range at which an agent perceives its surroundings and others. However, it is mostly used to prevent agents from unnecessary searches.

Of the dynamic variables the route the agent follows is the most important as it contains the path the agent is following. It is the same Route as explained in Section 7.3 and is actually a reference to the Route plan. What intersection and intersection lane the agent is approaching is determined by iterating over the path in Route. If an intersection and intersection lane is within a certain distance it is stored here. These variables are mostly used for the handle intersection plans, discussed in a later section.

7.4.3 External

External is the smallest of the belief bases and holds only two groups of entities: road data and other agents. Road data is a reference to the *RoadNet*, while other agents is a continuously updated list of any other agents in its vicinity. These other agents are found by checking if the vehicle's dimensions are within, or touching a sphere with a certain radius around the agent itself. This radius is determined by the *lookAhead* distance in *Driver*. The *Vehicle* belief base of the found agent is stored in a list if it does not already exist in that list. The list of found agents is also iterated over to find the closest vehicle directly in front of the agent.

7.5 Plans

Each plan is required to implement several functions and variables, to allow it to be integrated in the interpreter loop and used by the agent. In the plan generation phase the *MatchPlan* function is called of each plan to see if there is a possibility that the plan might match. In the filter phase the *EvaluatePlan* function of each matched plan is called to see if it is possible, safe, or necessary to execute the plan. This structure is similar to the graph from Hidas[21], where a lane changing manoeuvre is evaluated with the same tests. How the matching and evaluation are done is up to the user, however it is wise to avoid unnecessary calculations. Executing a plan is done through the *GetNextAction* function which returns a steering, throttle, and brake value. Each plan also has a *PlanComplete* and *PlanFailed* function which indicates if a plan was successfully executed or has failed. Furthermore, the fixed variables for the plans are: the layer the plan belongs to and the predicted speed when executing the plan. This variable is used when more than one plan passes the filter phase, on such occasions the plan with the higher predicted speed is chosen. For example, when an agent is approaching a truck on the highway it can choose to follow it, or try to overtake the truck. The predicted speed of the overtake plan is equal to the speed limit of that lane, while the predicted speed of the following plan is equal to the speed of the truck. Since the truck is driving under the speed limit the agent will overtake the truck. In the following subsections several important plans are explained. However, part of the strength of the model comes from the ability to write new plans for new behaviour, or altering existing plans.

7.5.1 Follow other

Follow other represents standard car-following behaviour: the agent will match its speed to that of the vehicle in front. The vehicle in front is the other agent from the *External* belief base, which acts as trigger in the generation phase. In the filter phase several checks and calculations are done to determine if it is necessary to follow the other agent. Should the other agents' speed be higher than, or equal to our own speed, it is not necessary to follow it. In the *External* belief base, no distinction is made as to where the other agent is, just that it is in front. Therefore we make sure that the other agent is on the same lane as we are, to prevent errors in curvy roads. The same applies for the on-intersection check, which prevents errors with regard to intersections. At execution, the agent will adjust its speed to that of the other agent as can be seen in Algorithm 2.

Algorithm 2 The *GetNextAction* function of the *FollowOther* plan.

Input: Belief base of the agent and the *Vehicle* of the other agent $Vehicle_{other}$.
Output: an *Action* with a value for steering s , throttle t and brake b . s remains unchanged.

```

 $d_{preferred} \leftarrow \text{getPreferredFollowingDistance}(\text{Belief})$ 
 $D_{other} \leftarrow \text{computeDistanceToOther}(Vehicle_{other})$ 
 $v \leftarrow \text{getVelocity}(\text{Belief})$ 
 $v_{other} \leftarrow \text{getVelocity}(Vehicle_{other})$ 
if  $D_{other} < d_{preferred}$  then
     $requiredDeacceleration \leftarrow (v_{other}^2 - v^2) / (2 * (D_{other} - d_{preferred}))$ 
     $b \leftarrow \lceil requiredDeacceleration / maximumDeacceleration \rceil$ 
else
     $t \leftarrow \text{MaintainSpeed}(v_{other})$ 
end if
return  $Action(s, t, b)$ 

```

7.5.2 Lane changing

Lane changing is made up of two plans, *ChangeLaneLeft* and *ChangeLaneRight*. Both plans are very similar in structure and are used for multi-lane roads like highways. The plans therefore follow the reasoning steps of a vehicle trying to overtake someone on the highway. However, in highway scenarios it is very likely that a vehicle can change from the first to the second lane, but the number of vehicles on the first lane make it impossible to change to the right again. Therefore the overtaking is split into two plans to encompass this uncertainty.

The *ChangeLaneLeft* plan is matched if there is a vehicle in front that is driving slower than the agent, and a parallel lane exists to the left of the agent. In the filter phase it is made sure that the parallel lane is clear and that the agent is close enough to the front vehicle to initiate the manoeuvre. At execution, the agent changes lanes by swapping the lane it is driving on with the parallel lane. To avoid sudden steering the clearance is increased for smoother steering.

The *ChangeLaneRight* plan follows a similar structure as the *ChangeLaneLeft* plan,

Algorithm 3 Algorithm to maintain a desired speed.

```
Input: Belief base of the agent and the desired speed  $v_{desired}$ .  
Output: A value for throttle  $t$ .  
 $v \leftarrow \text{getVelocity}(\text{Belief})$   
 $a \leftarrow \text{getAccelerationBelief}$   
 $a_{preferred} \leftarrow \text{getPreferredAcceleration}(\text{Belief})$   
 $t_{previous} \leftarrow \text{getPreviousAcceleration}(\text{Belief})$   
 $\Delta t \leftarrow \text{getTimeSinceLastFrame}()$   
if  $v < v_{desired}$  &&  $a < a_{preferred}$  then  
   $t \leftarrow t_{previous} + \Delta t * a_{preferred}$   
else  
  if  $v > v_{desired}$  then  
     $t \leftarrow t_{previous} - \Delta t * a_{preferred}$   
  else  
     $t \leftarrow t_{previous}$   
  end if  
end if  
return  $t$ 
```

the biggest difference is that there is no check whether a vehicle is in front the agent. The plan is therefore matched if a parallel lane to the right of the agent exists. This is in accordance with the Dutch traffic rule to always keep to the right lane when possible [20]. In the filter phase it is made sure that the manoeuvre is safe to perform by checking if the parallel lane is clear of any other agents. The execution is equal to the *ChangeLaneLeft* plan, only mirrored.

7.5.3 Handle intersections

Any driving instructor will give the same reasoning steps on how to handle an intersection: observe the situation, evaluate the situation with regard to safety, and traffic flow, decide on the best course of action, perform that action. The same line of reasoning can be applied to the handle intersection plans. Observing the situation is largely done beforehand by the *External* belief base. The agent knows where the other agents are and what type of intersection it is dealing with. For the evaluation all information from the belief base is processed to decide if it is safe to go or not. With the safety information the agent decides what the best course of action is and return that action to the interpreter, fulfilling the last two steps.

This structure is similar amongst all handle intersection plans, with the biggest difference being in the evaluation of the situation. Each intersection type has different right of way rules, which impacts whether it is safe for the agent to go or not, as explained in Algorithm 4. Therefore, we explain the handle equal intersection plan in extensive detail and give only the differences for the other plans. All plans start when the appropriate intersection is within viewing distance and end when the agent has left the intersection.

The handle equal intersection plan is matched if an intersection is on the route of the agent and within lookahead distance. Any intersections that lie beyond the first do not trigger a plan. For the filter phase, an extra check is made to confirm the

match. Agents must deal with intersections along their path, failure to follow the intersections' rules will result in collisions. Therefore the predicted speed of the plan is set to the maximum value, to assure that it is chosen above other plans. This has a consequence for what the plan must handle, since it overrides other plans that pass through the filter phase. However, for our model this is only one other plan, the following of other agents. Therefore the functionality of following other agents is also made possible for other plans. Algorithm 3 explains the *GetNextAction* function of the *HandleEqualIntersection* plan.

Algorithm 4 The *GetNextAction* function of the *HandleEqualIntersection* plan. The *doNothing()* function means that the agent will continue its path unaltered as dictated by the Route plan.

Input: Belief base of the agent, the *Vehicle* of the agent directly in front *Vehicle_{front}*, distance to the stopline *D_{stopline}* .

Output: an *Action* with a value for steering *s*, throttle *t* and brake *b*. Is handled by the returned functions in this algorithm.

D_{front} \leftarrow computeDistanceToOther(*Vehicle_{front}*)

```

if stopLineIsInFront() then
  if Dstopline < Dfront then
    if isIntersectionSafe() then
      return doNothing()
    else
      return brakeForStopLine(Dstopline)
    end if
  else
    if isIntersectionSafe() then
      return FollowOther(Vehiclefront)
    else
      return brakeForOther(Vehiclefront)
    end if
  end if
else
  if Exists(Vehiclefront) then
    return FollowOther(Vehiclefront)
  else
    return doNothing()
  end if
end if

```

A biased intersection has at least one incoming and outgoing lane marked as a major road. An agent that is not on the major road has to give way to agents that are on the major road, even if they are to the right of the agent. In all other cases right of way is handled equally to an equal intersection.

At intersection with traffic lights, the right of way is not handled by any traffic rules but by the color of the light, making it irrelevant who has right of way or not.

On a roundabout, all vehicles on the roundabout have right of way over other vehicles. Vehicles that want to get on the roundabout have to wait for the vehicles that

Algorithm 5 Algorithm to check if an intersection the agent is approaching is safe to cross.

Input: Belief base of the agent.
Output: *True* or *False*.
for all $Vehicle_{other}$ **in** `getOtherAgents(Belief)` **do**
 if `isApproachingOurIntersection($Vehicle_{other}$)` **then**
 if `otherHasRightOfWay($Vehicle_{other}$)` **then**
 return *False*
 end if
 else
 if `isOnOurIntersection($Vehicle_{other}$)` **then**
 if `isOnCollisionCourse($Vehicle_{other}$)` **then**
 return *False*
 end if
 end if
 end if
end for
return *True*

are driving on it.

8. Conclusion

In this section we give a conclusion and discussion on the experimental results. Next, we give the contribution of our method followed by the future work.

8.1 Discussion

From the results we can conclude that the clips show significantly different behaviour that both students and instructors can recognize. However, that behaviour is not always recognized as it was intended. Furthermore, although there are differences between the number of significant results between the tests with only students and the tests with both students and instructors, there is no significant effect of profession on the score of the clips. All five behavioural tests with only students are significant, compared to four significant behavioural tests amongst both students and instructors. For the realism tests, these differences are even greater; three scenario's show a significant result on realism amongst both students and instructors, while amongst only students none of the realism tests were significant. These differences make it difficult to draw further conclusions. However, since students are the primary users of the driving simulator, their results should be considered first. Therefore, we must conclude that there is no significant effect of clip on realism, and that most clips are more realistic than unrealistic due to the fact that most scores are negative. Whether our own agent model is overall more realistic than the current model remains inconclusive. However, when it comes to normal behaviour there usually is a significant difference for realism between the clips. Returning to the behavioural scores, most aggressive behaviour was not seen as aggressive, but as normal. This was especially clear in the leaving traffic light and highway right scenario's. For the preferred acceleration and tail distance, this means that they have to be respectively increased and lowered to achieve aggressive behaviour. Furthermore, the normal close clip in the biased intersection scenario was seen as cautious, where it was meant as normal. The reason as to why can be drawn from the open question results, which were included for exactly this purpose. A common complaint of the biased intersection was that the vehicles waited too long before accelerating when a car passed. The same applies for the leaving traffic light scenario where clips were often referred to as 'slow'.

8.2 Contribution

We have created a background traffic agent model for driving simulators using a BDI framework. To the best of our knowledge, this is the first time BDI has been used

for a background traffic agent of a driving simulator. Furthermore, the model can operate in any environment that has a RoadNet, creating enormous flexibility. Using BDI and the accompanied plans also makes the method very adaptable, making it easy to incorporate new behaviour.

8.3 Future work

Although our method surpasses existing background traffic models, it does not display the full range of traffic behaviour. This can be done by adding more plans to deal with more traffic situations. However, there are some situations that require more work. For example, the RoadNet does not make it possible to properly deal with multiple lanes at traffic lights. At such situations it is common that a single lane diverges into two lanes to cover multiple directions. The RoadNet is not yet equipped to deal with these situations. Furthermore the types of vehicles supported are still very limited compared to the broad range of traffic users. Although all four wheeled vehicles and higher are covered, anything with three or less wheels is not yet possible. The biggest addition would be bicycles which are an intricate part of urban traffic in the Netherlands. However, they require their own lanes and traffic rules which would require more work. When it comes to future experiments it would be interesting to test more different traffic situations. Furthermore, further tests are needed to discover students' attitude towards deviant and aggressive behaviour since that aspect was lacking in our own experiments. Moreover, the experiments were limited to clips due to time, but it would be interesting to test and validate the model with an actual driving simulator. Overall, our method provides a sound foundation for future expansion to ultimately include the whole spectrum of traffic behaviour.

A. Questionnaire Example

%VIDEO_PLACEHOLDER%

38.

Hoe zou u het gedrag van de auto/auto's beoordelen?

Ik vind het rijgedrag: Voorzichtig Normaal Agressief

39.

Waarom vindt u dat de auto/auto's dat rijgedrag heeft/hebben? (optioneel)

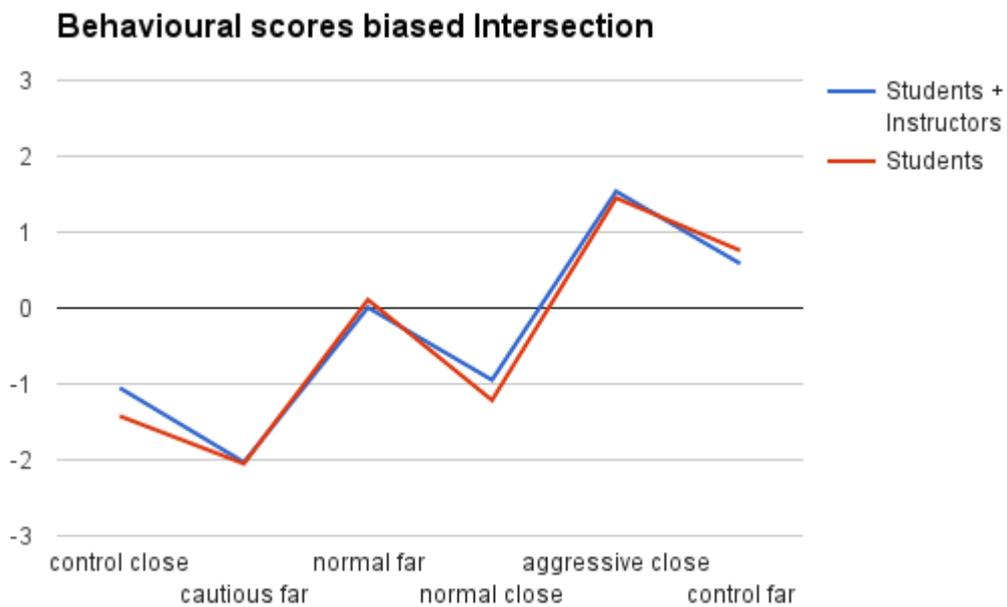
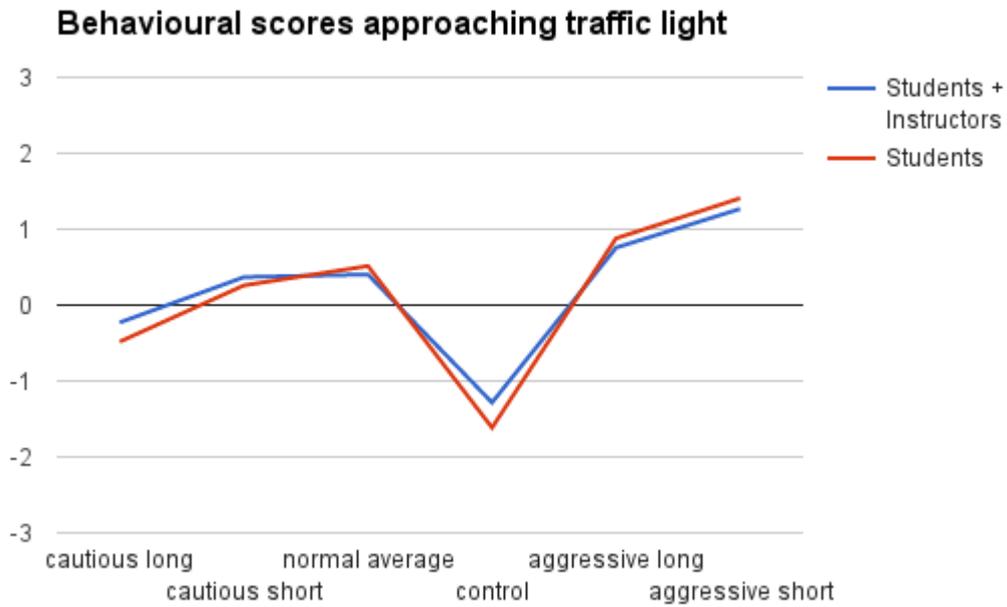
40.

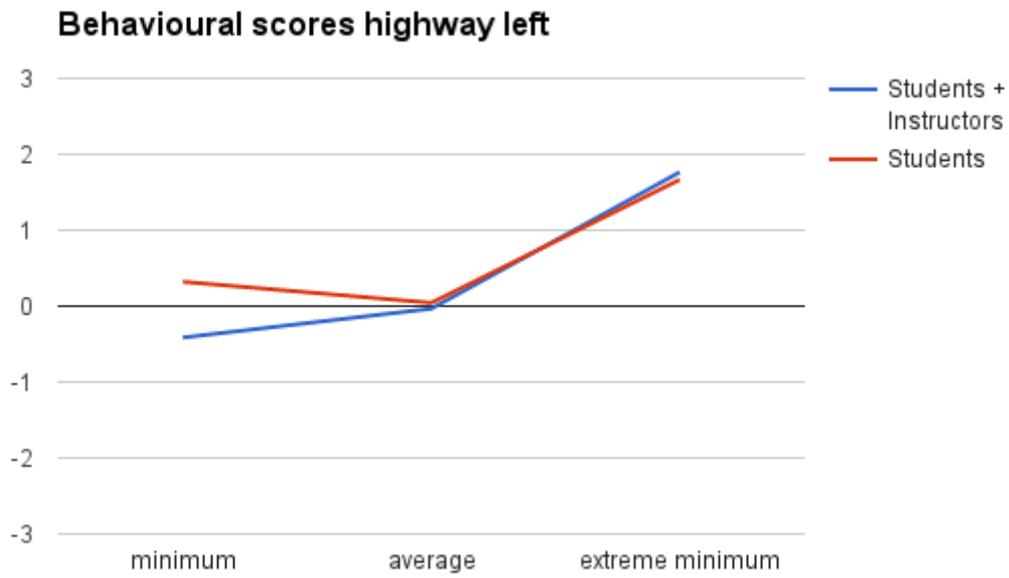
Stelling: Het rijgedrag van de auto/auto's is realistisch.

Ik ben het daar: zeer mee eens mee eens niet eens/niet oneens mee oneens zeer mee oneens

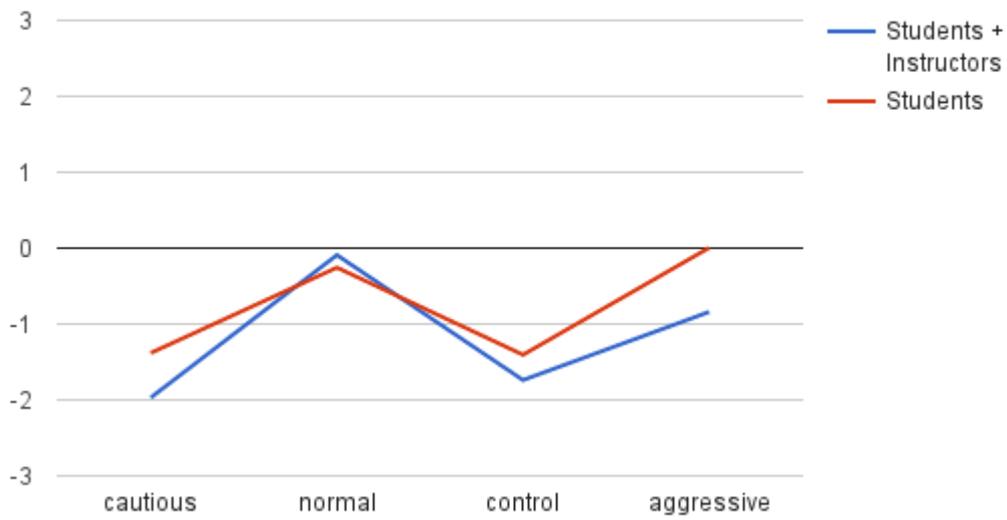
Volgende filmpje

B. Analysis Graphs

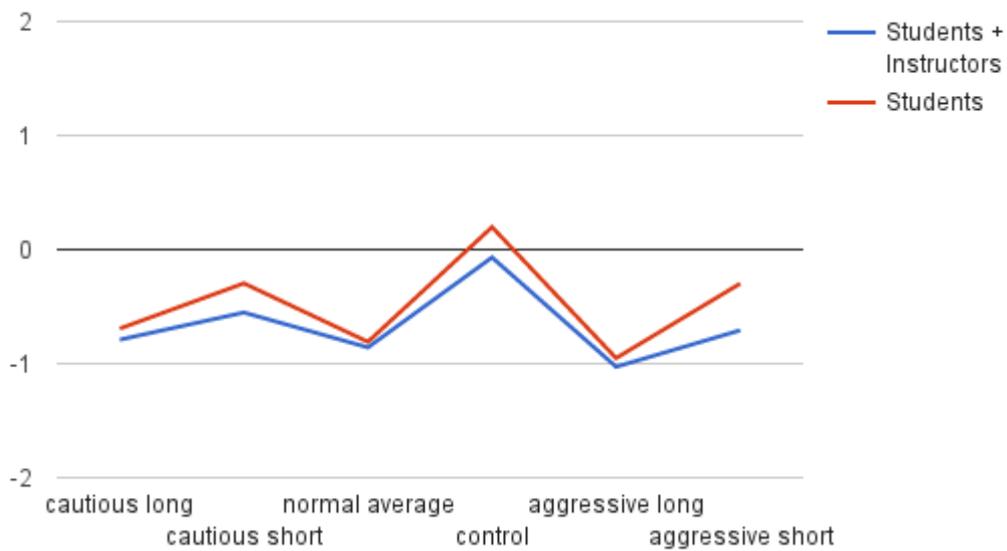


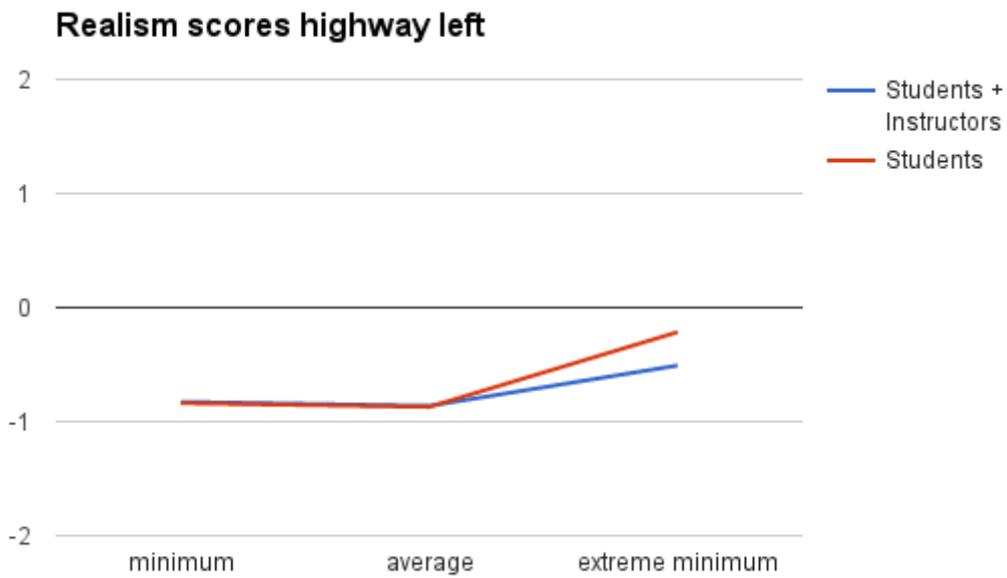


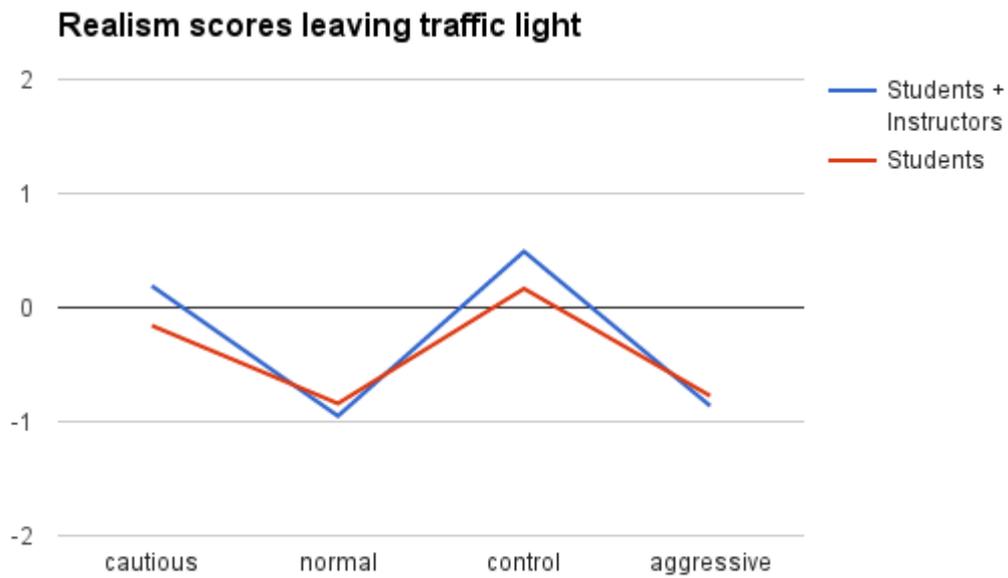
Behavioural scores leaving traffic light



Realism scores approaching traffic light







Bibliography

- [1] R Akçelik and M Besley. “Acceleration and deceleration models”. In: *23rd Conference of Australian Institutes of Transport Research (CAITR 2001)*, Monash University, Melbourne, Australia. 2001, pp. 10–12.
- [2] A L C Bazzan and F Klügl. “A review on agent-based technology for traffic and transportation”. In: *The Knowledge Engineering Review* 29.03 (2014), pp. 375–403.
- [3] V Bogdanović, N Ruškić, Z Papić, and M Simeunović. “The research of vehicle acceleration at signalized intersections”. In: *PROMET-Traffic&Transportation* 25.1 (2013), pp. 33–42.
- [4] M Bosnjak and B Batinic. “Understanding the Willingness to Participate in Online Surveys: the case of e-mail questionnaires”. In: *Online social sciences* 81 (2002), pp. 92–112.
- [5] M Brackstone, B Sultan, and M McDonald. “Motorway driver behaviour: studies on car following”. In: *Transportation Research Part F: Traffic Psychology and Behaviour* 5.1 (2002), pp. 31–46.
- [6] B Burmeister, A Haddadi, and G Matylis. “Application of multi-agent systems in traffic and transportation”. In: *Software Engineering. IEE Proceedings-[see also Software, IEE Proceedings]*. Vol. 144. 1. IET. 1997, pp. 51–60.
- [7] Green Dino BV. *Company website*. 2015. URL: <http://www.greendino.nl/>.
- [8] J Chaplier, T Nguyen That, and G Hewatt Mand Gallée. “Toward a standard: RoadXML, the road network database format”. In: *Actes INRETS* (2010), pp. 211–220.
- [9] B Chen and H H Cheng. “A review of the applications of agent technology in traffic and transportation systems”. In: *Intelligent Transportation Systems, IEEE Transactions on* 11.2 (2010), pp. 485–497.
- [10] M Dastani and B Testerink. “From Multi-Agent Programming to Object Oriented Design Patterns”. In: *Engineering Multi-Agent Systems*. Springer, 2014, pp. 204–226.
- [11] J C F De Winter, S De Groot, M Mulder, P A Wieringa, J Dankelman, and J A Mulder. “Relationships between driving simulator performance and driving test results”. In: *Ergonomics* 52.2 (2009), pp. 137–153.
- [12] H A Deery. “Hazard and risk perception among young novice drivers”. In: *Journal of safety research* 30.4 (2000), pp. 225–236.

- [13] M Demir and A Çavuşoğlu. “A new driver behavior model to create realistic urban traffic environment”. In: *Transportation research part F: traffic psychology and behaviour* 15.3 (2012), pp. 289–296.
- [14] M Dupuis and H Grezlikowski. “OpenDRIVE®-an open standard for the description of roads in driving simulations”. In: *Driving Simulation Conference*. 2006, pp. 25–36.
- [15] J.L. Ellis. *Statistiek voor de psychologie*.
- [16] B Filzek and B Breuer. “Distance behaviour on motorways with regard to active safety: A comparison between adaptive-cruise-control (ACC) and driver”. In: *Proceedings of the International Technical Conference on Enhanced Safety of Vehicles*. Vol. 201. 2001, pp. 1–8.
- [17] P G Gipps. “A behavioural car-following model for computer simulation”. In: *Transportation Research Part B: Methodological* 15.2 (1981), pp. 105–111.
- [18] N Goode, P M Salmon, and M G Lenné. “Simulation-based driver and vehicle crew training: applications, efficacy and future directions”. In: *Applied ergonomics* 44.3 (2013), pp. 435–444.
- [19] Dutch Government. *2seconden*. 2015. URL: <http://www.2seconden.nl/>.
- [20] Dutch Government. *verkeersborden-en-verkeersregels-in-nederland*. 2012. URL: <https://www.rijksoverheid.nl/documenten/brochures/2012/02/28/verkeersborden-en-verkeersregels-in-nederland>.
- [21] P Hidas. “Modelling lane changing and merging in microscopic traffic simulation”. In: *Transportation Research Part C: Emerging Technologies* 10.5 (2002), pp. 351–371.
- [22] I Karamouzas, R Geraerts, and M Overmars. “Indicative routes for path planning and crowd simulation”. In: *Proceedings of the 4th International Conference on Foundations of Digital Games*. ACM. 2009, pp. 113–120.
- [23] A Kesting, M Treiber, and D Helbing. “Agents for traffic simulation”. In: *arXiv preprint* (2008).
- [24] B Lacroix, P Mathieu, and A Kemeny. “Formalizing the construction of populations in multi-agent simulations”. In: *Engineering Applications of Artificial Intelligence* 26.1 (2013), pp. 211–226.
- [25] A K Maurya and P S Bokare. “Study of deceleration behaviour of different vehicle types”. In: *International Journal for Traffic and Transport Engineering* 2.3 (2012), pp. 253–270.
- [26] J J Olstam, J Lundgren, M Adlers, and P Matstoms. “A framework for simulation of surrounding vehicles in driving simulators”. In: *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 18.3 (2008), p. 9.
- [27] J Olstam, S Espié, S Måardh, J Jansson, and J Lundgren. “An algorithm for combining autonomous vehicles and controlled events in driving simulator experiments”. In: *Transportation research part C: emerging technologies* 19.6 (2011), pp. 1185–1201.
- [28] A S Rao and M P Georgeff. “BDI agents: From theory to practice.” In: *ICMAS*. Vol. 95. 1995, pp. 312–319.

- [29] A S Rao and M P Georgeff. “Modeling rational agents within a BDI-architecture”. In: *KR 91* (1991), pp. 473–484.
- [30] R J F Rossetti, R H Bordini, A L C Bazzan, S Bampi, R Liu, and D Van Vliet. “Using BDI agents to improve driver modelling in a commuter scenario”. In: *Transportation Research Part C: Emerging Technologies 10.5* (2002), pp. 373–398.
- [31] S Seele, R Herpers, and C Bauckhage. “Cognitive agents for microscopic traffic simulations in virtual environments”. In: *Entertainment Computing-ICEC 2012*. Springer, 2012, pp. 318–325.
- [32] T Al-Shihabi and R Mourant. “Toward more realistic driving behavior models for autonomous vehicles in driving simulators”. In: *Transportation Research Record: Journal of the Transportation Research Board 1843* (2003), pp. 41–49.
- [33] A Shoulson and N I Badler. “Event-centric control for background agents”. In: *Interactive Storytelling*. Springer, 2011, pp. 193–198.
- [34] M Si, S C Marsella, and D Pynadath. “Directorial control in a decision-theoretic framework for interactive narrative”. In: *Interactive Storytelling*. Springer, 2009, pp. 221–233.
- [35] Unity Technologies. *Company website*. 2015. URL: <https://unity3d.com/>.
- [36] I H C Wassink, E van Dijk, J Zwiers, A Nijholt, J Kuipers, and A Brugman. “Bringing hollywood to the driving school: Dynamic scenario generation in simulations and games”. In: *Intelligent Technologies for Interactive Entertainment*. Springer, 2005, pp. 288–292.
- [37] B G Witmer and M J Singer. “Measuring presence in virtual environments: A presence questionnaire”. In: *Presence: Teleoperators and virtual environments 7.3* (1998), pp. 225–240.
- [38] Z Yin, Y Lin, S Kamarthi, and R R Mourant. “Interactions Between Intelligent Autonomous Vehicles and Human Controlled Vehicles in 3D Simulation”. 2010.