



HAL
open science

Decision-making for automated vehicles at intersections adapting human-like behavior

Pierre de Beaucorps, Thomas Streubel, Anne Verroust-Blondet, Fawzi
Nashashibi, Benazouz Bradai, Paulo Resende

► **To cite this version:**

Pierre de Beaucorps, Thomas Streubel, Anne Verroust-Blondet, Fawzi Nashashibi, Benazouz Bradai, et al.. Decision-making for automated vehicles at intersections adapting human-like behavior. IV'17 - IEEE Intelligent Vehicles Symposium, IEEE, Jun 2017, Redondo Beach, United States. hal-01531516

HAL Id: hal-01531516

<https://inria.hal.science/hal-01531516>

Submitted on 2 Jun 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Decision-making for automated vehicles at intersections adapting human-like behavior

Pierre de Beaucorps* Thomas Streubel* Anne Verroust-Blondet* Fawzi Nashashibi*
Benazouz Bradai† Paulo Resende†

Abstract

Learning from human driver’s strategies for solving complex and potentially dangerous situations including interaction with other road users has the potential to improve decision-making methods for automated vehicles. In this paper, we focus on simple unsignalized intersections and roundabouts in presence of another vehicle. We propose a human-like decision-making algorithm for these scenarios built up from human drivers recordings. The algorithm includes a risk assessment to avoid collisions in the intersection area. Three road topologies with different interaction scenarios were presented to human participants on a previously developed simulation tool. The same scenarios have been used to validate our decision-making process. The algorithm showed promising results with no collisions in all setups and the ability to successfully determine to go before or after another vehicle.

1 Introduction

While the development of automated vehicles advances further, the focus in research moves towards more complex scenarios especially in urban traffic environments. Here, the high density of vehicles leads to numerous interactions that challenge the cognitive abilities of drivers and the development of automated vehicle systems alike.

The focus of this paper is the decision-making process for an automated vehicle at intersections and roundabouts in the presence of other vehicles. This problem is usually coupled with the introduction of a behavior estimation approach for manually driven vehicles as in [1]. The aim here is to investigate the human behavior in order to create a naturalistic decision process. The high error rate of drivers leading to accidents and casualties is a major argument for developing automated vehicles. Thus, our decision process is combined with an objective risk assessment to avoid collisions. However, the high adaptability and robust decision-making process of drivers should not be neglected. Therefore, our approach incorporates human driving behavior enhanced by a robust risk assessment.

We created a new multi-agent simulation environment for testing and validating the decision-making algorithm design. It enables the simultaneous driving of multiple agents (human or automaton). Such a tool is essential for the investigation of the interaction between human drivers and automated vehicles. It is also used for data acquisition that is required to tune the algorithm. The key aspect of our work is the introduction of a decision-making process capable of handling intersection and roundabout scenarios avoiding collisions.

1.1 Related Work

Different decision-making problems for automated driving have been recently tackled. Depending on the scenarios, some concentrate on lane-changing scenarios ([2, 3, 4]) while others focus on intersections. The relevant features at intersections are different than in lane change scenarios (see [5] for a survey). In [6], the post-encroachment time (PET) is mentioned as measurable parameter connected with risk in left turning scenarios. It refers to the area that is common to the paths of two vehicles coming from different directions. The PET is the time between one vehicle leaving and the other vehicle entering this zone. The study examined the driving behavior in a real driving scenario created on a test field. The situational conditions varied between comfort and hurried and the aim was to investigate the accepted comfort limits (expressed by PET and lateral acceleration) under both conditions. While in the comfort situation the PET was around 2 seconds, the lower limit is slightly above 1 second for the hurried condition. One of the challenges in such an experimental setup seems to be the realization of comparable situations for all test drivers since the approaching behavior at an intersection differs greatly between drivers.

Decision-making at intersections is addressed recently in [7] and [8]. Galceran et al. introduce a multi-policy approach, i.e. there is a previously defined discrete set of policies (maneuvers) estimated for vehicles including a Bayesian change-point detection where the intention of executing a certain maneuver switches. The maneuver set decreases the computation time in favor of a real-time application. The decision-making is based on a reward function taking several safety and comfort features into account and the optimal maneuver is found by forward simulation. Here, an interaction between vehicles on a maneuver level is taken into

*INRIA Paris, 2 rue Simone Iff 75012 Paris, France
{pierre.de-beaucorps, thomas.streubel, anne.verroust, fawzi.nashashibi}@inria.fr

†Valeo Driving Assistance Research - 34 rue St-André 93012 Bobigny, France

account. The decision-making is similarly approached in [8]. Although, the prediction is realized through a Hidden Markov Model (HMM) to determine the driving direction (left, right, straight) and the intention to pass (stop, yield, no yield). Further, several policies are created for the automated car and the observed vehicle by creating artificial speed profiles through discrete acceleration steps (step size 0.5 s) within a certain comfort limit for deceleration and acceleration and to comply with the speed limit. The prediction horizon is chosen very high with 8 seconds which is compensated with a wide step size. The decision of the best policy is made using a reward function that takes comfort and safety features into account. A potential conflict and the traffic law compliance are both handled as binary resulting in a high penalty for collision or breaking the law.

Driven speed profile extraction have been used recently for different purposes. This was used in [9] to model automated vehicle behavior by building a neural network that predicts the velocity declination in curvy roads and turns. In [10] a functional discretization of space for road intersection crossing has been designed considering speed profile extraction.

1.2 Concept

The process towards decision-making follows the scheme described in Figure 1. Concerning the **environment perception**, we suppose that the information is available and that sensing errors are neglected. The **behavior estimation** is realized by a simple extrapolation of the movement of other vehicles. The **trajectory planning** of the automated vehicle is split into two parts. First, the **path planning** is realized by defining a route corridor formed by a list of cubic Bézier curves maintaining G1 continuity. It is computed offline according to the road topology by an external path planning module. Further, the trajectory planning is completed by assigning a **speed profile** to the path. Speed profiles are retrieved through test drivers driving in our simulation environment to shape human-like behavior. This inherently assures a certain comfort boundary for lateral acceleration. The

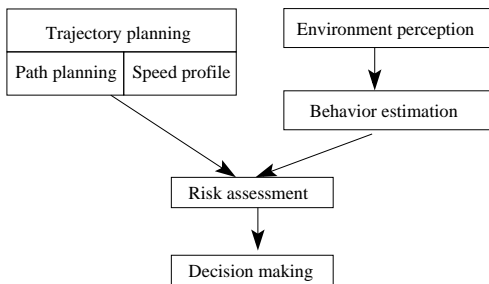


Figure 1: Decision-making framework

trajectory planning and estimated trajectories of other vehicles are used for a situational **risk assessment** to avoid collisions in the intersection area. The risk is

determined by the PET and other safety-related variables. The **decision-making** is finalized by choosing the lowest risk or, in case of no risk, the time-optimal speed profile available among the several speed profiles clustered from the retrieved driving data. So, this approach combines natural speed profiles extracted from the human driving data recorded on our simulation tool and objective risk features to adapt the velocity of an automated vehicle and avoid collision risks in the intersection.

This paper is organized as follows. In section 2, the simulation tool is introduced, its functionality is described and the data acquisition is outlined in section 3. Further, section 4 presents the decision-making algorithm. Finally, in section 5 the simulation results are shown and in section 6 the conclusion is drawn.

2 Simulation tool

For testing decision-making algorithms in automated vehicles, a simulation tool is indispensable. Even on a test field, the risk of collision and damaging the equipment is too high. Since a simulator fitting all requirements was unavailable, a new custom simulation tool was created. In the following, we present its main characteristics.

2.1 Requirements

To deal with interactive behavior in urban traffic scenarios, the following requirements were identified:

1. The simulator tool should be capable of processing human inputs in a controlled environment (driver-in-the-loop). This way different drivers can be tested in the exact same scenarios. The inputs as well as the vehicle dynamics are recorded for further analysis.
2. New algorithms can be tested for objective (minimal distance to the other vehicle, collision free) and subjective (risk score, driver acceptance) criteria, even in critical driving situations, which is not easily feasible with a hardware platform.
3. Multiple vehicles with decision-making algorithms can run against each other for validation or for iterative learning approaches such as self-play reinforcement learning.

Overall, a realistic, versatile, multi-agent and generic simulator is required. Realistic refers to a reasonable physics model to generate vehicles executing actually drivable trajectories. Versatile is meant mainly in terms of road topology. The tool should allow a customization of road designs (crossroads and roundabouts with different intersection angles, lane widths, traffic signs, etc.) to create new scenarios easily and quickly. For interaction, a multi-agent simulation allows to gather several human inputs in real time and place them in the same virtual environment. This enables to study the process of human decision making

especially regarding the choice to go before or after another vehicle. The generic attribute refers to the ability to run human inputs, recordings and automated vehicles in the same software at the same time.

2.2 Vehicle model

We adapted an implementation of the bullet physics engine [11] in order to run the simulation on a server which centralizes the physics computations for every client of the architecture. The vehicle is controlled by setting three parameters: steering wheel angle, gas, and brake. The physical engine then calculates the next state of the world according to the friction slip of the ground and the vehicle model [12]. The latter includes four independent wheels with their suspension systems and a body mass on top. It reproduces pitch and bounce degrees of freedom for the powertrain giving realistic visual feedback to the human driver for throttle changes. Each suspension system includes stiffness, compression, damping and maximum travel.

2.3 Visualization

An example of the 3D view of the simulation can be seen in Fig. 2. A special feature of the tool is to enable a rapid integration of new road topologies in the scenario design. The visual representation of the environment is generated from a two-dimensional bird-view image, a description of traffic signs (type, position, orientation) and a list of available paths for vehicles. A monitor mode enables to display the scene from a bird-view to watch the study, restart the scenario and assign predefined paths to available human drivers.




Figure 2: left: bird view of the scene (monitor mode), right: driver view

2.4 Control Systems

The intention is to gather meaningful behavioral data of human drivers decisions in complicated driving situations. A detailed trace of trajectories is not required. For this reason and because it is more convenient for the study, the vehicles are controlled only by a computer mouse. To keep trajectories and decisions realistic the two following controls are applied to assist the driver.

2.4.1 Automatic Lateral Control

The driver is told to follow a route that is highlighted in the 3D view. As the only interaction regarding other road users is either pass or give way, there is no reason to escape a predefined route corridor. As the steering task is neither required nor easy in this artificial environment, an automatic lateral control helps to keep the lane according to the predefined path. The human driver is only required to set the speed of the vehicle.

2.4.2 Semi-Automatic Longitudinal Control

According to the above, the only remaining critical task to define a driving behavior is to set the speed profile (i.e. the adjustment of the vehicle speed over time). The assistance for the longitudinal control acts like a Cruise Control system; the user sets a target speed with the mouse cursor, and the longitudinal control adjusts the brake/gas ratio to reach it in a smooth, natural way, as follows, K being defined empirically:

$$Acc = Acc_{Max} \left(1 - \left(\frac{Speed_{current}}{Speed_{target}} \right)^3 \right) \quad (1)$$

$$Acc > 0 \Rightarrow gas = Acc$$

$$Acc < 0 \Rightarrow brake = -K \cdot Acc$$

The speed is set regarding the height of the mouse cursor: the higher the cursor, the greater the acceleration. After several minutes in a test scenario, the drivers seemed sufficiently accustomed to the simulated driving task.

3 Data acquisition

The simulated environment is further used to gather data about human behavior in several simple intersections. The experiment is conducted with 10 participants, each one of them driving their vehicle (ego-vehicle) in 56 different setups involving interaction with an automated vehicle (dummy-vehicle).

3.1 Scenarios

For the current work, we focus on the following road topologies (cf. Figure 3):

- Two crossroads with four legs; the first one in which two opposing routes lead to a yield sign and the second one without traffic signs except those reminding to give priority to the right
- A roundabout with four exits and one inner lane, every incoming lane leading to a yield sign

Seven main scenarios are selected: six for the crossroads and one for the roundabout. For the crossroad, we distinguish between the maneuver performed by the ego-vehicle crossing the path of another vehicle (Left Turn Across Path (LTAP), Right Turn Across Path



Figure 3: Available road topologies (RTAP), Straight Across Path (SAP)) and the direction from which the other vehicle is coming (Left Direction (LD), Opposite Direction (OD), Right Direction (RD)). In all cases the other vehicle has the priority. Therefore, there are six scenarios for the crossroad topology: three for Left Turn Across Path (LTAP/LD, LTAP/OD, LTAP/RD), two for Straight Across Path (SAP/LD, SAP/RD) and one for Right Turn Across Path (RTAP/LD). These scenarios are repeated with different setups where the starting position and the cruise speed for the dummy-vehicle change, so the interaction between the two vehicles always varies.

3.2 The tracking algorithm

Since the human drivers do not have the same behavior in the scene, they may enter the intersection with different timings. As an attempt to standardize each scenario over every participant in the study, the dummy-vehicle is manipulated to reach a given position with a given speed when the ego-vehicle enters the intersection. Therefore, a dedicated algorithm has been designed to perform a real-time tracking that shapes the speed profile of the dummy-vehicle. Since interaction between the dummy-vehicle and drivers was the primary target of this manipulation, risky situations were created where the human drivers do not always succeed to avoid a collision.

3.3 Procedure

All participants were in possession of a driver's license and were familiar with driving a car. To get used to the simulation environment, there was a test loop presented first to gain experience with the controls of the vehicle for several minutes. The study includes 30min of driving through the 56 setups of the 7 scenarios, each of them being ended by a subjective risk assessment form. Neither any detail about the purpose of the study nor any specific instructions are given, except a request to observe traffic rules and to drive as in a real environment. The simulator records all relevant information about all vehicles every 100ms. Their position, orientation, speed, acceleration (gas or brake applied) and steering angle are stored in a dedicated JSON file.

4 Algorithm design

The trajectory planning is divided into a predefined path computed offline and the variable speed profile that determinates the behavior of the automated vehicle. The outcome of the decision-making process is an




Figure 4: The Post Encroachment Time (PET)

optimal speed to be reached by the control module of the automated vehicle. This target speed is updated at each time step.

A set of reference speed profiles is determined offline, extracted from human data (cf. section 4.1). The decision-making process iterates over them to make a prediction and decides then to follow the reference speed profile that fits the best to the present situation.

In this approach, the post-encroachment time (PET) is used as the objective risk feature. It corresponds to the time difference between the first exit of the collision zone and the last entrance in it (see Figure 4). As the paths are predefined, the collision zone is approximated by a rectangle in crossroads and a sector of an annulus in roundabouts.

4.1 The reference speed profiles

The reference speed profiles are established for the seven scenarios described in section 3.1. For a given scenario, we consider the corresponding speed profiles generated by the users in the simulation tool. The aim is to identify the most common shapes of speed profiles and average them to create a group of reference speed profiles. In order to compare different speed profiles generated by drivers, we define them with respect to the curvilinear abscissa of the vehicle's position which is projected on the predefined path. Since the lateral control of the vehicles ensures that the predefined path is followed, this alignment of speed profiles is reasonably accurate.

The generated speed profiles are divided into three groups (ego vehicle first, other vehicle first, risky situation) considering their PET values. The ego vehicle passed first if $PET > 0$, and gave way to other vehicle if $PET < 0$. Cases considered too risky (collisions included) and situations with a high score on the driver's subjective risk (> 8 on a scale from 1 to 10) are discarded. Further, a k-means clustering is performed for the remaining two groups, giving a set of centroids of speed profiles that average the observed behaviors (Figure 5-B). To apply the k-mean algorithm on sequential data, every profile is considered a vector having one element per abscissa step (i.e. 1m step on the predefined path). The recorded speed profiles rarely

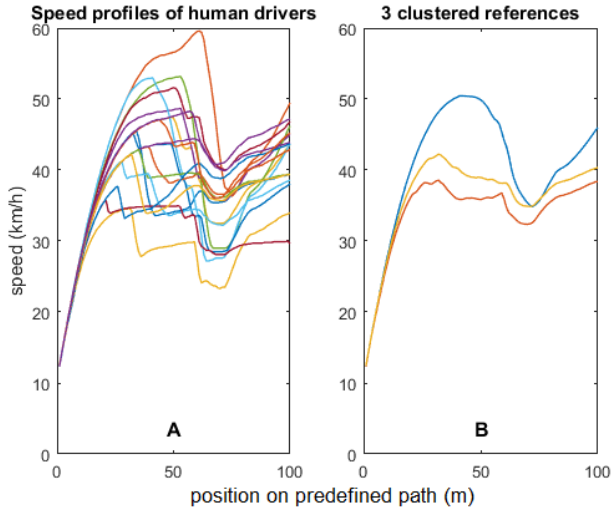


Figure 5: A: Recorded speed profiles in crossroad LTAP/RD scenario. Only "passing first" speed profiles ($PET > 0$) are considered. B: Three reference speed profiles after clustering

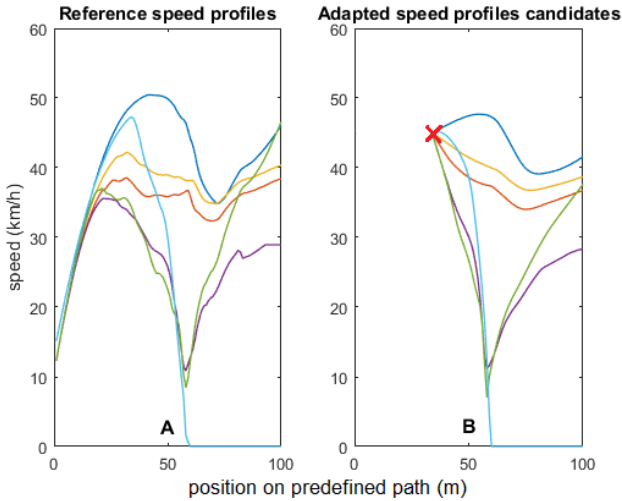


Figure 6: A: The six reference speed profiles in crossroad LTAP/RD scenario (3 passing, 2 yielding, 1 stopping). B: The speed profiles candidates, adapted to the current speed (red cross) with a control algorithm.

intersect over the abscissa, which enables the k-mean algorithm to perform a good separation of the speed profiles into clusters. On some scenarios, a first clustering builds singleton clusters with only one outsider speed profile inside. The outsider profiles are removed from data before performing a second clustering that is better balanced.

The reference speed profiles (see Figure 6-A for an example) consist of three speed profiles for entering the intersection fast (for situations where the ego vehicle passed first safely), two speed profiles for entering the intersection slower (ego yield to other vehicle), and one speed profile for stopping definitely before the intersection at the stop line.

4.2 Short term prediction

When the automated vehicle AV approaches an intersection while another vehicle veh2 is closing in from a different direction, a decision has to be made to go before or after veh2. For this purpose, a short term prediction is made at each time step (approx. 100ms), with a time horizon of 2s if the AV reaches the intersection within this horizon. The prediction of the AV utilizes the reference speed profiles while the maneuver of veh2 is assumed to be predicted correctly and the position on the path is an extrapolation with the current speed. About 30m before the intersection the AV enters the "decision zone", which is the region where the decision has to be made. The decision zone ends at the path of veh2.

To be adapted to any situation (i.e. to any driving speed), the reference speed profiles (Fig. 6-A) are considered as a target for the control algorithm used above in the simulated environment (cf. equation 1). The control algorithm starts at the current speed and then simulates the future longitudinal dynamics of the vehicle adapting to the given reference speed profile. The output gives actual executable speed profile candidates (Fig. 6-B).

Based on the predicted motion of both vehicles, the PET is computed for each of the six speed profiles candidates and is provided to the decision-making algorithm. The prediction iterates over every speed profile candidate for the automated vehicle, so different behaviors are tested against the situation. The prediction stops when the vehicle AV exits the "decision zone" and crosses the path of veh2. After this point, the decision process is set back to a default lane-following mode.

4.3 Decision-making

The decision-making process depends on the objective risk feature and the predicted speed of the vehicle. Estimated trajectories leading to a $|PET|$ below a critical threshold are rejected for being too risky. After the risk evaluation, the speed profile with the highest instantaneous speed is selected from the remaining candidates to contribute to time efficiency. In the cases where there is no appropriate candidate (i.e. emergency cases), the stopping speed profile is automatically selected. This leads either to a stop before crossing the path if possible or at least to a high reduction of speed. As previously mentioned, as soon as the path of veh2 is entered the AV changes into a lane-following mode since a stop in the path of the other vehicle is undesirable and potentially dangerous. Finally, the selected speed profile is sent to the control module.

5 Results

The algorithm is capable of passing the intersection safely. The driving strategy is reasonable and seems to replicate human-like behavior to the extent that such a property can be evaluated in the present simulated environment. The results are also presented in a video




Figure 7: Number of collisions

provided in [13]. The following will focus on an objective validation of our algorithm.

To enable an objective comparison, an automated vehicle with the described algorithm is tested in the same setups that were used to retrieve the speed profiles. The validation setups have been grouped into the seven scenarios described in 3.1. The ten human drivers have been clustered into two groups of five drivers according to their averaged speed inside the intersection over all scenarios. The group of faster drivers is called the "daring human drivers" and the other is the "cautious human drivers" group. Besides these two groups, the evaluation is performed with two versions of the decision-making algorithm differing in the *PET* threshold for taking a predicted speed profile into account. The automated vehicles 1 and 2 ensure respectively $|PET| > 1.5$ and $|PET| > 0.7$, which leads to two discriminable behaviors.

As presented before, the seven scenarios have eight variants leading to 56 different setups. As it is not trivial to condensate the complete course of a driving situation into one score, three features have been chosen for the evaluation of the driving performance.

- The number of collisions (Fig. 7). In case of the two groups of human drivers, this value is displayed with a min to max range referring to the collisions of an individual driver of the group and the mean indicator for the group.
- The number of daring situations, i.e. situations where the driver succeeded to pass before the other vehicle (Fig. 8). In case of the two groups of human drivers, this value is again displayed with a min to max range referring to the numbers of an individual driver of the group and the mean indicator for the group.
- The average speed (unless when stopped) inside the intersection (Fig. 9). For each scenario, the mean and standard deviation are computed over the 8 different setups (and over the 5 drivers in case of a group).

In all scenarios, automated vehicles 1 and 2 passed every setup without colliding with the other vehicle (see Fig. 7). However, the dummy car was implemented to generate comparable setups that are risky. So, its

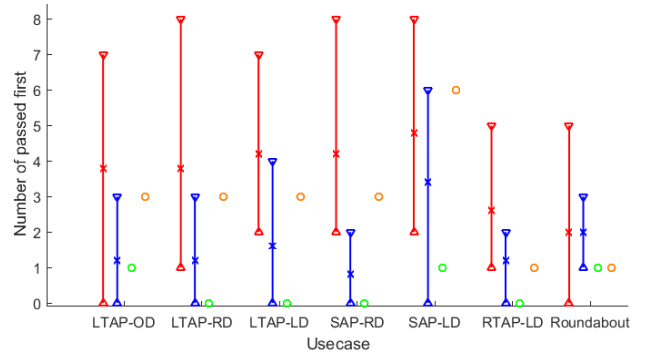


Figure 8: Number of situations where driver passed first

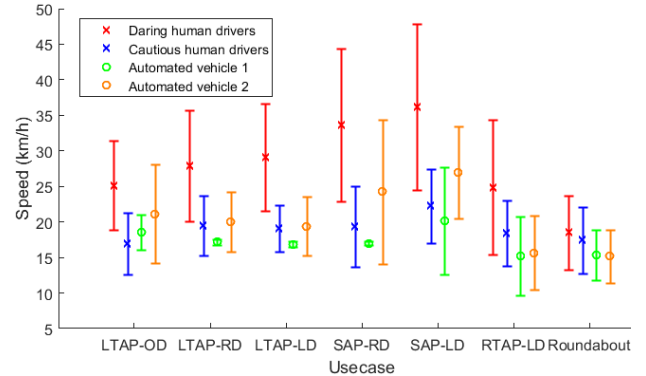


Figure 9: Average speed inside the intersection

behavior was highly artificial. This explains the high collision rate for the human drivers in both groups.

The daring human driver group enters the intersection first most frequently in average compared to the cautious human driver group (see Fig. 8). The only exception is in the roundabout scenario where the mean values are the same for both groups. The automated vehicles act more conservative to prevent collisions. While automated vehicle 2 with a higher risk acceptance is passing first multiple times in several scenarios, the automated vehicle 1 is almost always yielding to the incoming dummy vehicle. So, AV 2 is comparable with the human drivers in terms of using the opportunity to pass first and is often between both human driver groups. In the SAP-LD scenario, AV 2 manages to pass first even more often than the daring human driver group in average. Here, the human drivers seem to miss some opportunities to pass before the incoming vehicle. In roundabouts both AVs are rather conservative and pass only once in all the setups.

6 Conclusion

We presented a human-like decision-making algorithm for simple unsignalized intersections in presence of another vehicle. A dedicated simulation tool was used to record human speed profiles in especially risky situations where a conflict can arise if the driver dares to pass before the other vehicle. The aggregated data has been used to design an algorithm that enables human-like speed profile generation while ensuring safety with objective risk assessment. The post-encroachment time (PET) appeared to be a useful feature to estimate risk at unsignalized intersections. Therefore, it was used in

a short-term prediction that successfully avoided collisions in every tested scenario. The results show that our approach can handle standard intersection scenarios with another car approaching from a different direction without collisions. This is not only realizable by just avoiding the conflict at all and yielding in almost every setup (compare AV 1) but also by accepting some level of risk and passing first if it is possible (compare AV 2). This combines some level of human-like daring behavior with the objective risk evaluation of an automated system.

Work in progress extends this approach to more complex scenarios, especially with generic road topologies and more road users interacting simultaneously. A new method of validation is also investigated to estimate the acceptance of the automated vehicle behavior by human drivers in the scene.

References

- [1] T. Streubel, “Situation Assessment at Intersections for Driver Assistance and Automated Vehicle Control,” Dissertation, TU Chemnitz, 2016.
- [2] M. Bahram, A. Wolf, M. Aeberhard, and D. Wollherr, “A prediction-based reactive driving strategy for highly automated driving function on freeways,” in *Proc. of IEEE Intelligent Vehicles Symposium*, 2014, pp. 400–406.
- [3] J. Wei and J. M. Dolan, “A robust autonomous freeway driving algorithm,” in *Proc. of IEEE Intelligent Vehicles Symposium*, 2009, pp. 1015–1020.
- [4] J. Nilsson and J. Sjoberg, “Strategic decision making for automated driving on two-lane, one way roads using model predictive control,” in *Proc. of IEEE Intelligent Vehicles Symposium*, 2013, pp. 1253–1258.
- [5] M. S. Shirazi and B. T. Morris, “Looking at intersections: A survey of intersection monitoring, behavior and safety analysis of recent studies,” *IEEE Transactions on Intelligent Transportation Systems*, no. 99, pp. 1–21, 2016.
- [6] J. Bärghman, K. Smith, and J. Werneke, “Quantifying drivers’ comfort-zone and dread-zone boundaries in left turn across path/opposite direction (LTAP/OD) scenarios,” *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 35, pp. 170–184, 2015.
- [7] E. Galceran, A. G. Cunningham, R. M. Eustice, and E. Olson, “Multipolicy Decision-Making for Autonomous Driving via Change-point-based Behavior Prediction,” *Robotics: Science and Systems*, 2015.
- [8] W. Song, G. Xiong, and H. Chen, “Intention-Aware Autonomous Driving Decision-Making in an Uncontrolled Intersection,” *Mathematical Problems in Engineering*, 2016.
- [9] X. Geng, H. Liang, H. Xu, B. Yu, and M. Zhu, “Human-driver speed profile modeling for autonomous vehicle’s velocity strategy on curvy paths,” in *Proc. of IEEE Intelligent Vehicles Symposium*, June 2016, pp. 755–760.
- [10] M. Barbier, C. Laugier, O. Simonin, and J. Ibanez-Guzman, “Functional Discretization of Space Using Gaussian Processes for Road Intersection,” in *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC 2016)*, Rio de Janeiro, Brazil, Nov. 2016, p. 7.
- [11] “Bullet physics engine,” <http://bulletphysics.org>.
- [12] “Documentation about the raycast vehicle model in the bullet physics engine,” <http://tinyurl.com/ydfb71m>.
- [13] “Video demonstrating our approach running in intersection,” <https://team.inria.fr/rits/membres/pierre-de-beaucorps/demo-iv2017>.