



<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *ICRA International Conference on Robotics and Automation*.

Citation for the original published paper:

Karlsson, J., van Waveren, S., Pek, C., Torre, I., Leite, I. et al. (2021)
Encoding Human Driving Styles in Motion Planning for Autonomous Vehicles
In:

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-292195>

Encoding Human Driving Styles in Motion Planning for Autonomous Vehicles

Jesper Karlsson*, Sanne van Waveren*, Christian Pek, Ilaria Torre, Iolanda Leite and Jana Tumova

Division of Robotics, Perception and Learning,

KTH Royal Institute of Technology Stockholm, Sweden

Email: {jeskarl, sannevw, pek2, ilariat, iolanda, tumova}@kth.se

*Contributed Equally

Abstract—Driving styles play a major role in the acceptance and use of autonomous vehicles. Yet, existing motion planning techniques can often only incorporate simple driving styles that are modeled by the developers of the planner and not tailored to the passenger. We present a new approach to encode human driving styles through the use of signal temporal logic and its robustness metrics. Specifically, we use a penalty structure that can be used in many motion planning frameworks, and calibrate its parameters to model different automated driving styles. We combine this penalty structure with a set of signal temporal logic formula, based on the Responsibility-Sensitive Safety model, to generate trajectories that we expected to correlate with three different driving styles: aggressive, neutral, and defensive. An online study showed that people perceived different parameterizations of the motion planner as unique driving styles, and that most people tend to prefer a more defensive automated driving style, which correlated to their self-reported driving style.

Index Terms—Autonomous Vehicle Navigation, Formal Methods in Robotics and Automation, Human Factors and Human-in-the-Loop

I. INTRODUCTION

Driving styles are crucial for passenger trust and acceptance of autonomous vehicles (AVs) [1], [2]. While driving styles can be learned from human data [3]–[6], recent research showed that people tend to prefer a more defensive driving style for AVs over their own driving style [7]. Yet, many AVs tend to drive overly defensive, limiting their efficiency [5] and compromising passenger comfort [8].

Motion planners can incorporate different driving styles through adapting cost functions and a set of constraints. Often, developers carefully handcraft these cost functions to achieve higher passenger comfort, e.g. by favoring smooth accelerations. Modeling the complex aspects of driving styles in a single cost function is difficult and constraints are mostly used in a strict fashion (hard constraints) to ensure safety [9], road compliance [10] or the satisfaction of traffic rules [11]. Arguably, AVs may occasionally need to violate certain constraints (soft constraints) for driving progress, e.g., crossing lane markings when another vehicle blocks the current lane.

Constraints can be modeled by temporal logic, such as LTL [12] or STL [13], [14], which extends classical first order logic with the possibility to model temporal aspects. STL is particularly well-suited to model constraints in robotic

applications, as it operates over continuous signals (e.g., trajectories) and allows arbitrary evaluation functions. Contrary to other logics, STL’s inherent quantitative semantics allow for quality evaluation of a proposed solution, e.g., how much and how long a trajectory violates the specification.

In our prior work [15], [16], we use syntactically co-safe Linear Temporal Logic (scLTL) to encode traffic rules in various traffic scenarios. We can determine trajectories that minimally violate these rules, by assigning costs to potential violations of each rule. Such least-violating motion planning approaches attempt to find a motion plan that minimizes violation over the set of specifications, when not all specifications can be strictly guaranteed [17], [18].

To enrich the STL robustness metric, Lindemann et al. [19] propose the discrete average space robustness, which considers the average satisfaction of a formula, rather than the worst case scenario of the standard robustness measure. Barbosa et al. [20] use the STL robustness semantics from [21] for exploration tasks. We complement these works [21] [20], by showing that robustness metrics of STL formulas can be used to model human driving styles.

We previously proposed a cost function that encodes the importance of mission time against the satisfaction of spatial preferences [21]. In this work, we implement a motion planner that uses this cost function to support spatio-temporal constraints. For a set of scenarios, we run our motion planner with different calibrations to encode driving behavior from which an objective defensive score can be computed.

This objective defensiveness score is adopted from Basu et al [7], who define driving style in terms of *defensiveness* level, and compute an objective defensiveness score from a set of distinct driving features. Using the open-source driving simulator CARLA [22], we collect trajectories that each correspond to a distinct driving style (aggressive, neutral, and defensive).

In an online study, people evaluated the perceived level of defensiveness for different AV driving styles. We hypothesize that calibrating this cost function allows us to model human driving styles in motion planning. The main contributions of this work are as follows:

- 1) We provide STL formulas for the RSS model [23] as a specification to encode perceived driving styles;
- 2) We show that parameterizing the cost function results

in different driving behaviors that are perceived as distinct driving styles, and that people prefer different AV generated motion styles, which correlates to their self-reported driving style.

II. MOTION PLANNING PRELIMINARIES

We employ STL, which is defined recursively according to the grammar: $\psi ::= \mu \mid \neg\mu \mid \psi_1 \wedge \psi_2 \mid \psi_1 \vee \psi_2 \mid \psi_1 \mathbf{U}_{[a,b]} \psi_2$, where μ is an atomic predicate, s.t.

$$\mu = \begin{cases} \top & \iff g(\mathbf{x}) \geq 0 \\ \perp & \iff g(\mathbf{x}) < 0, \end{cases}$$

and $g(\mathbf{x}) : \mathbb{R}^m \rightarrow \mathbb{R}$ is an evaluation function of a signal $\mathbf{x} \in \mathbb{R}^m$, corresponding to the trajectory of a dynamical system.

The satisfaction of a formula ψ by a signal \mathbf{x} at time t , is defined by the satisfaction relation, $(\mathbf{x}, t) \models \psi$. Here $\psi_1 \mathbf{U}_{[a,b]} \psi_2$ denotes the *Until* operator, where $(\mathbf{x}, t) \models \psi_1 \mathbf{U}_{[a,b]} \psi_2$ specifies that ψ_1 must hold for signal \mathbf{x} at all times until a point in interval $[a, b]$ where ψ_2 holds. Other operators can be derived from the basic operators, such as *eventually*: $\mathbf{F}_{[a,b]} \psi \equiv \top \mathbf{U}_{[a,b]} \psi$ (ψ has to hold at some point in interval $[a, b]$) and *always*: $\mathbf{G}_{[a,b]} \psi \equiv \neg \mathbf{F}_{[a,b]} \neg \psi$ ($\neg \psi$ cannot hold at any time in interval $[a, b]$).

a) *Quantitative Semantics of STL formulas*: Here we will briefly introduce the notion of space- and time robustness, as proposed by Donzé et al. [13].

Let $\rho(\mu, \mathbf{x}, t)$ denote a real-valued function, such that $(\mathbf{x}, t) \models \psi \equiv \rho(\mu, \mathbf{x}, t) > 0$. It can be interpreted as the degree of satisfaction/violation of formula ψ by signal \mathbf{x} . This quantitative semantic is known as *space robustness*. Furthermore, let $\theta^-(\psi, \mathbf{x}, t)$ denote the *left-time robustness*, which can be interpreted as the duration leading up to time t in which ψ has been satisfied/violated, i.e.

$$\theta^-(\psi, \mathbf{x}, t) = \text{sign}(\rho(\psi, \mathbf{x}, t)) \max \{d \geq 0 \mid \forall t' \in [t-d, t] \text{ sign}(\rho(\psi, \mathbf{x}, t')) = \text{sign}(\rho(\psi, \mathbf{x}, t))\},$$

b) *Least-violating Motion planning with STL constraints*: Driving specifications can be satisfied or violated to a certain extent, e.g., driving at 100 km/h is worse than at 60 km/h on a road with 50 km/h speed limit. To measure the violation of such a specification ψ , we use the *modified left-time robustness* $\theta^*(\psi, \mathbf{x}, t)$ [21]:

$$\theta^*(\psi, \mathbf{x}, t) = \min(\theta^-(\psi, \mathbf{x}, t), 0). \quad (1)$$

Let $w(\alpha, A, \rho(\psi, \mathbf{x}, t))$ be a positive function that expresses the compromise between a time-efficient trajectory and one that prioritizes the spatial preferences, defined as follows:

$$w(\alpha, A, \rho(\psi, \mathbf{x}, t)) = \begin{cases} \infty, & \iff \rho(\psi, \mathbf{x}, t) < -\alpha \\ 0, & \iff \rho(\psi, \mathbf{x}, t) > 0 \\ -\frac{A}{\alpha} \rho(\psi, \mathbf{x}, t) & \iff \text{otherwise.} \end{cases} \quad (2)$$

where α and A are user-defined parameters that specify the lowest allowed spatial robustness, and the pace at which the trajectory is being penalized as it approaches the the lowest allowed spatial robustness.

The modified left-time robustness $\theta^*(\psi, \mathbf{x}, t)$ and the cost function $w(\alpha, A, \rho(\psi, \mathbf{x}, t))$ can together quantify spatio-temporal robustness in STL. The cost of a trajectory \mathbf{x} , with duration D , using said robustness is defined as:

$$P_\Psi(\mathbf{x}) = - \int_0^D \theta^*(\psi, \mathbf{x}, t) w(\rho(\psi, \mathbf{x}, t)) dt, \quad (3)$$

which can be viewed as the quantitative semantics of the operator $\widehat{\mathbf{G}}_{[0,D]}$ [21].

III. MOTION PLANNING PROBLEM

We consider an AV ν_{ego} operating in a traffic network governed by rules of the road, similar to Vasile et al. [15], and traffic behaviors derived from RSS [23]. The motion planning problem weights the duration of the task and the violation of the road-rules against one another to create a least-violating trajectory (Sec. III-C). The road rules and desired behaviors are specified as spatial preferences encoded in STL [13].

A. Vehicle Model

We define a vehicle ν_{ego} as a tuple $\nu_{ego} = (X, U, \mathfrak{R}, \text{Sense}, \Pi, \mathcal{L})$, where $X \subset \mathbb{R}^4$, $U \subset \mathbb{R}^n$ and $\mathfrak{R} \subset \mathbb{R}^2$ are the state space, control space and road network, respectively. For simplicity, the vehicle dynamics use a modified form of the Dubin's model, defined as follows:

$$\begin{aligned} \dot{p}_x &= V \cos \theta, & \dot{p}_y &= V \sin \theta, \\ \dot{\theta} &= u_1, & V &= u_2, \end{aligned} \quad (4)$$

where the state $x = (p_x, p_y, \theta, V)^\top$ corresponds to the position, yaw, and velocity of the vehicle. u_1 and u_2 are the control inputs, denoting turn rate and velocity, respectively. Let $\text{Sense} : X \rightarrow 2^{\mathfrak{R}}$ be the vehicle's limited sensing region, from which it detects road signs, obstacles and markings. Let Π denote the set of road signs and markings that annotate the road network. Further, let $h(\nu_i) : X \rightarrow \mathbb{R}^2$ denote the observation function that specifies the i th vehicle's position.

a) *Road Markings and Rules*: Let $\mathcal{L}(t)$ denote the road labeling map that assigns labels from Π to the sensing region $\text{Sense}(q)$. These labels determine which road rules should be adhered to in the road context defined by the sensing region.

The road rules are defined in the reactive form as $(\theta_j^a \xrightarrow{\text{Re}} \theta_j^g)$, where $\theta_j^a \in \Theta^a$, $\theta_j^g \in \Theta^g$, $\xrightarrow{\text{Re}}$, are assumption formulas, guarantee formulas and a reactive implication [15]. A traffic rule θ_j is active if the assumption formula θ_j^a is satisfied. Intuitively, the reactive implication can be understood as defining the restrictions based on the road context of the AV.

We write the full mission of the vehicle as:

$$\theta_j \mathbf{U} g \equiv (\theta_j^a \xrightarrow{\text{Re}} \theta_j^g) \mathbf{U} g, \quad (5)$$

where $g \subset \mathfrak{R}$ is a goal region to be reached in the road network. The cost of the mission is measured using the duration D , indicating the time it takes for a vehicle to reach the goal region.

TABLE I: Application-Specific Symbols

Variable	Description
ϵ	Safety footprint of obstacles in the environment, it describes the desired distance to all static obstacles.
μ	Safety threshold that defines how close to the center of the lane a vehicle should be.
l_c	Position of the center of lane parallel to the current position of the vehicle.
$a_{lat,max}$	Maximum lateral acceleration of a vehicle.
\mathbf{v}^{ν_i}	Velocity vector of the i th vehicle, contains lateral and longitudinal velocity

B. Spatial Preferences

We formalize common traffic rules from RSS [23] and driving practices, e.g., the *three-second rule* [24], using STL.

Rule 1 (*Safe longitudinal distance – “3-second rule”*) Let ν_1 and ν_2 denote two vehicles in a configuration s.t. ν_2 is ahead of ν_1 . The longitudinal distance d_{lon} between ν_1 and ν_2 is denoted as safe if the time t_{lon} it takes to traverse d_{lon} is at least three seconds.

The specification can be encoded in the following formula:

$$\Psi_{lon} = \left(\psi_{o,s}^a \stackrel{\text{Re}}{\Rightarrow} \widehat{\mathbf{G}}_{[0,D]} (|h_{lon}(\nu_2) - h_{lon}(\nu_1)| - d_{lon,min}) \right),$$

where $\psi_{o,s}^a = \widehat{\mathbf{G}}_{[0,D]} (v_{lon}^{\nu_1} - v_{lon}^{\nu_2})$.

The three-second rule recommends that the time-headway to the preceding car should not be smaller than 2 – 3 s. This can be compared to the brake reaction time, which has been shown to be 1 – 2 s [25]. The smallest allowed longitudinal distance between two cars, $d_{lon,min}$ is:

$$d_{lon,min} = (v_{lon}^{\nu_1} - v_{lon}^{\nu_2}) t_{BT}, \quad (6)$$

where t_{BT} denotes the brake reaction time; here, 3 s.

The intuition behind the activation function $\psi_{o,s}^a$ is that if the velocity of ν_2 is higher than that of ν_1 , then there is no need to consider a minimum distance, because the distance between the two vehicles will naturally increase over time.

Rule 2 (*Safe lateral distance*) The lateral distance d_{lat} between ν_1 and ν_2 is unsafe if the lateral velocities $v_{lat}^{\nu_1}$, $v_{lat}^{\nu_2}$ are pointed towards each other and the time it takes to traverse $d_{min,lat}$ (see Eq. 7) is within the interval $t \in [0, t_{PT}]$.

The specification can be encoded in the following formula:

$$\Psi_{lat} = \left(\psi_{lat}^a \stackrel{\text{Re}}{\Rightarrow} \widehat{\mathbf{G}}_{[0,D]} (|h_{lat}(\nu_2) - h_{lat}(\nu_1)| - d_{lat,min}) \right),$$

where

$$\psi_{lat}^a = \neg \Psi_{lon} \wedge \widehat{\mathbf{G}}_{[t, t+t_{PT}]} (-\mathbf{v}^{\nu_1} \cdot \mathbf{v}^{\nu_2}).$$

The smallest allowed lateral distance between two vehicles, $d_{min,lat}$ is:

$$d_{min,lat} = (v_{lat}^{\nu_1} + v_{lat}^{\nu_2}) t_{PT} + \frac{a_{lat,max} t_{PT}^2}{2}. \quad (7)$$

This corresponds to the minimum distance in which it is possible for a driver to react in case of unforeseen lane-changing or swerving from adjoining lane. The value of t_{PT} is based on the concept of perception time, which has shown to be 0.5–1 s [25] and $a_{lat,max}$ is the maximum acceleration in the lateral direction (see Table I).

The assumption formula ψ_{lat}^a states that Rule 2 is active whenever ν_{ego} is within a minimum safety distance to another vehicle, and the lateral velocities have different directions. The guarantee formula ensures that the vehicle lateral distance between the two vehicles is larger than $d_{lat,min}$ (Eq. 7).

Rule 3 (*Safe overtaking manoeuvre – single vehicle*) The overtaking manoeuvre \mathbf{x}^{ν_2} of vehicle ν_1 is denoted as unsafe, if at any point during the manoeuvre, ν_1 is to the right of ν_2 relative to the direction of travel.

The specification can be encoded in the following formula:

$$\Psi_{o,s} = \left(\psi_{o,s}^a \stackrel{\text{Re}}{\Rightarrow} \widehat{\mathbf{G}}_{[0,D]} (h_{lat}(\nu_1) - h_{lat}(\nu_2)) \right).$$

If ν_2 is travelling at a velocity that is higher than the velocity of ν_1 , then there is no need for an overtaking manoeuvre. The assumption formula $\psi_{o,s}^a$ ensures that no overtaking manoeuvres are attempted unless needed. The guarantee formula ensures that the overtaking manoeuvre is performed from the left of the vehicle, in the direction of travel.

Rule 4 (*Maintaining the speed limit*) A trajectory \mathbf{x} is deemed unsafe if at any point during the run, the velocity of the vehicle exceeds the speed limit of the road, denoted v_{max} .

The specification can be encoded in the following formula:

$$\Psi_v = \top \stackrel{\text{Re}}{\Rightarrow} \widehat{\mathbf{G}} (v_{max} - v^{\nu_{ego}}).$$

The formula ensures that the velocity of the ego vehicle should not exceed the speed limit of the road.

Rule 5 (*Safety distance – static obstacles*) Let *obs* denote an obstruction in the sensing region of ν_{ego} . The distance d_{safe} between ν_{ego} and *obs* is denoted as safe if $d_{safe} > \epsilon$.

The specification can be encoded in the following formula:

$$\Psi_{safe} = \left(\top \stackrel{\text{Re}}{\Rightarrow} \widehat{\mathbf{G}}_{[0,D]} (|h(\nu_{ego}) - h(obs)| - \epsilon) \right),$$

This road rule ensures that the vehicle should not violate the safety footprint of an obstacle, which is defined by the parameter ϵ (see Table I), e.g., half the width of the vehicle.

Rule 6 (*Maintaining center line*) A trajectory must consistently stay in the center of the lane, to avoid confusion to other driving participants and unsafe situations.

The STL formula encoding the specification is:

$$\Psi = \top \stackrel{\text{Re}}{\Rightarrow} \widehat{\mathbf{G}}_{[0,D]} (\mu - |h_{lat}(\nu) - l_c|).$$

Here μ is a user defined threshold specifying how closely a vehicle should stay to the center of its lane, l_c (see Table I). In this work, μ is defined as the width of the vehicle.

TABLE II: The features that operationalize the driving styles

Feature	Description
Mean Distance to Lead Car	Average distance to a vehicle ahead of (ν_{ego})
Mean Time Headway	Average time to another vehicle
Time Headway During Lane Change	Average time to another vehicle while performing a lane change
Distance Headway During Lane Change	Average distance to a vehicle when performing a lane change
Braking Distance from the Intersection	Stopping distance from an intersection
Time to Stop	Time it takes for a vehicle to go from the velocity at the point of deceleration to complete stop
Maximum Turn Speed	Maximum speed used by (ν_{ego}) while turning
Average Speed 20 Meters Before Intersection	Average speed of a vehicle when it is within 20m of an intersection
Average Speed 20 Meters Before Chicane	Average speed of a vehicle when it is within 20m of a chicane

C. Penalty Structure

The rules defined in Sec. III-B provide a set of complex spatio-temporal specifications, which the ego vehicle ν_{ego} needs to adhere to while driving. However, not all specifications can be guaranteed at all times. As such, we aim to provide a *least-violating trajectory*, which, in the case that some specification needs to be violated in order to fulfill the mission, will minimize the penalty:

$$P(\mathbf{x}) = D + P_{\Psi}(\mathbf{x}), \quad (8)$$

where $P_{\Psi}(\mathbf{x})$ is the cumulative violation of the set of spatio-temporal specifications Ψ .

IV. ONLINE STUDY

We studied whether people perceived the generated trajectories as intended and we obtained a preferred defensiveness value for each participant from their preference ratings of generated trajectories pairs.

A. Defensiveness Score

We computed the defensiveness score, $DS \in [0, 1]$, for our generated trajectories from eleven features (see Table II), ten were adopted from Basu et al. [7], and one was created to specify the average speed of a vehicle when it is within 20m (empirically decided) of a chicane. The chicane is one of the four scenarios we used in this work, Fig. 1 illustrates all four traffic scenarios. To simulate different driving behavior, we ran our motion planner 200 times per scenario with different calibrations. Rules 1, 2, 3 and 5 use the same parameter calibration, which is randomized from an empirically chosen parameter interval: $A = [0, 20]$, $\alpha = [0, 4]$. The parameter intervals of Rule 4 were $A = [0, 10]$, $\alpha = [0, 20]$, and the intervals of of Rule 6 were $A = [0, 20]$, $\alpha = [2, 10]$. For each run, we normalized each feature using min-max normalization, regardless of the scenario. Then, we calculated the defensiveness score by averaging all feature values for a run

and task to generate trajectories corresponding to a high defensiveness score (*defensive driving style*, $DS_D \in [0.8, 1.0]$), an intermediate defensiveness score (*neutral driving style*, $DS_N \in [0.4, 0.6]$), and a low defensiveness score (*aggressive driving style*, $DS_A \in [0.0, 0.2]$). The final trajectories are visualized in the media attachment of this paper.

B. Manipulated Variables

We manipulated the defensiveness score, scenario, and perspective. We compared two perspectives: frontal and top (see Fig. 2 for the frontal view and Fig. 1c for top view of the chicane scenario). To get perceived driving styles, we presented participants with 3 (level of defensiveness) x 2 (perspective) x 4 (scenarios) = 24 videos, in randomized order.


To collect participants' preferred defensiveness scores, we manipulated whether the participant started with a low-defensive video or with a high-defensive video (start score). Participants indicated their preference for pairs of our generated trajectories. Each pair was either low-defensive vs. neutral, or high-defensive vs. neutral. If they indicated on the first trial to prefer the low-defensive or high-defensive video, they would converge on that trial and would not be presented with the other comparison. If they picked neutral as preferred style for both trials, would converge on the neutral defensiveness score. As the first trial may influence people's response, each participant would converge on a score twice for the same set of trials, once starting from a low-defensive vs. neutral video pair and once starting from a high-defensive vs. neutral video pair. All factors were treated as within-subjects and counterbalanced to avoid order effects, resulting in a maximum of 4 (scenario) x 2 (start score) x 2 (perspective) x 2 (comparisons) = 32 video pairs, presented in randomized order.

C. Participants


35 participants (26 male and 9 female) between the ages of 24 and 66 (mean = 38.3, SD = 11.3) were recruited on Amazon Mechanical Turk. The nationalities of the participant include Americans, 1 German, 1 Brazilian, and 1 Romanian. Participants had a driver's license for 17.8 years on average, and 29 participants indicated to drive on average more than 2-4 times a week; as the current Covid-19 pandemic may affect people's current day-to-day life, we specifically asked for pre-pandemic driving habits. All participants indicated that they had heard of AVs and/or self-driving vehicles prior to participating in the study.

D. Procedure and Measurements


After filling in demographic information and indicating how often they would drive a car, participants are instructed to imagine they owned an AV and were inside it. To obtain convergence scores, we asked them to indicate which trajectory they preferred the AV to drive, using a two-alternative forced-choice, with open answer to motivate their choice. Participants were informed that they would see the car sometimes from the frontal perspective and sometimes




(a) The AV attempts to merge onto a roundabout with a vehicle approaching from the left.



(b) The AV approaches a junction, equipped with traffic lights, with the goal to turn left, with an oncoming vehicle approaching.



(c) The AV has to pass a chicane to reach the end of the lane, with an oncoming vehicle approaching.



(d) The AV is blocked by a slow moving car on a straight two-lane road, while a faster moving car is approaching in the oncoming lane.

Fig. 1: In top-down perspective, the four different scenarios used in the online study.




Fig. 2: The chicane scenario as seen from frontal perspective.

from a top perspective. Then, we administered additional questionnaires as described in the next paragraph and we collected perceived defensiveness for each of the videos on a 7-points scale (1 being *extremely defensive* and 7 being *extremely aggressive* (reverse scored in analysis). Finally, participants had the option to write final comments; they were compensated \$8USD.

The Ten-Item Personality-Inventory (TIPI) [26] was administered to assess personality traits, and to assess self-reported driving style, participants indicated on a 7-points scale whether (1) they identified as being an conservative or adventurous driver, (2) they liked joy of motion or comfort of steadiness, (3) they varied their driving by external conditions

(e.g., road conditions), and (4) their driving experience from somewhat experienced to very skillful, adopted from Basu et al. [7].

V. RESULTS

Firstly, we ran a mixed-effects linear model with perceived defensiveness rating as response variable, defensiveness level and perspective as predictors, and participant and scenario as random effects. As can be seen in Fig. 3, participants perceived the degree of defensiveness (high score means high defensiveness) of each video as intended, in that participants rated the high-defensive (HD) videos as more defensive than the neutral (N) ones, and the neutral ones more defensive than the low-defensive (LD) videos ($\chi^2(2) = 521.92, p < .001$; $M_{LD} = 2.46, M_N = 4.38, M_{HD} = 5.42$). Participants also mentioned defensiveness when motivating their choices: “While I can appreciate patience in a car (...) There is such a thing as too defensive when driving which can cause accidents much the same way being overly offensive can do (...)”

Secondly, we investigated whether participants’ preferred defensiveness score was influenced by their starting defensiveness score, and by the perspective (frontal and top). We fitted a cumulative link mixed model with convergence score as dependent variable, starting score and perspective as predictors, and participant id and scenario as random effects. We also added the individual differences measured after

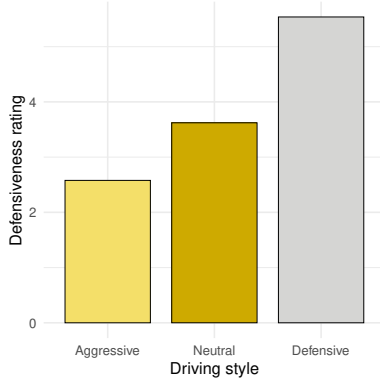


Fig. 3: Perceived defensiveness ratings of the driving styles.

the video evaluation phase (gender, age, Big-5 personality traits and self-reported driving behavior) to the model as covariates. We found a main effect of starting score ($\chi^2(1) = 44.86, p < .001$). People converged more towards a defensive style if they started from a high-defensive video; conversely, they converged more to an aggressive style if they started from a low-defensive video (see Fig. 4).

There was no main effect of perspective however ($\chi^2(1) = 0.21, p = .64$), and no interaction ($\chi^2(2) = 0.27, p = .87$), suggesting that participants converged towards a high-defensive or low-defensive video regardless of whether they saw the video from a frontal or top view. Regarding the individual differences, we only found a main effect of self-reported driving behavior, specifically on whether they considered themselves a conservative or adventurous driver ($\chi^2(1) = 8.26, p < .005$). Specifically, people who considered themselves more conservative converged towards a more defensive AV driving preference.

VI. DISCUSSION AND CONCLUSIONS

We extend prior work that proposed a penalty structure for STL formula [21] and a defensiveness measure [7] by studying whether human driving styles can be incorporated into a motion planner by calibration of this penalty structure. For several urban traffic scenarios, we generated trajectories with low, neutral, and high defensiveness score.

First, we hypothesized that people would perceive different parameterizations of our motion planner as a unique driving style. The results confirmed our hypothesis and showed that people’s defensiveness ratings corresponded with the defensiveness score of each calibration, which indicates that it is possible to use STL and its quantitative semantics to encode human driving styles. This opens up for several interesting venues of research in the field of formal methods and motion planning. For instance, parameter learning and synthesis of signal temporal logic specifications that leverage human driving styles. As the primary strength of formal method based approaches is that they provide verification of encoded specifications, it would be possible to guarantee that the set driving style holds for the entire lifetime of a system.




Fig. 4: Convergence scores divided by starting scores.

Second, we evaluated whether we can get a preferred value of defensiveness for people and relate this to their perceived driving style, so that in the future our system can adapt to that. Interestingly, the initial level of defensiveness that they would see influenced which level of defensiveness they would converge on. However, most people tended to converge to a defensive style, regardless of starting point. Participants motivated their choice for a defensive driving style with human attributes mentioning that the more aggressive car “(...) shows impatience.”, and “i like that the right [car] video is patient (...) that is respectful driving”. This suggests that additional factors should be considered when designing automated driving styles, such as driving etiquette.

We encoded only a subset of rules from the RSS model, providing a basic framework for AV behavior in any traffic scenario, but it will struggle when differentiating subtly different trajectories, as these differences might not be reflected in the specification. For instance, the current specification allows the agent to determine what speed limit it should adhere to, yet fails to consider jerk. Encoding a larger and more diverse set of rules from the RSS model, would allow the motion planner to represent a larger subset of common driving patterns.

Avenues of future research include a further investigation of individual characteristics that may influence people’s preferred driving style, a more fine-grained analysis of the individual and combined influence of specifications on driving style, and conducting a simulator study in which participants evaluate the automated driving styles as generated by our motion planner.

VII. ACKNOWLEDGMENTS

We thank Dmytro Kalpakchi, Bram Willemsen, and Patrik Jonell for their contributions to the crowdsourcing pipeline. We thank Alexis Linard for valuable discussions. This work was supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation and by the Swedish Research Council (grants 2017-05189 and 2017-05102). The authors are affiliated with Digital Futures.

REFERENCES

- [1] F. Ekman, M. Johansson, L.-O. Bligård, M. Karlsson, and H. Strömberg, "Exploring automated vehicle driving styles as a source of trust information," *Transportation research part F: traffic psychology and behaviour*, vol. 65, pp. 268–279, 2019.
- [2] C. Strauch, K. Mühl, K. Patro, C. Grabmaier, S. Reithinger, M. Baumann, and A. Huckauf, "Real autonomous driving from a passenger's perspective: Two experimental investigations using gaze behaviour and trust ratings in field and simulator," *Transportation research part F: traffic psychology and behaviour*, vol. 66, pp. 15–28, 2019.
- [3] P. Abbeel and A. Y. Ng, "Apprenticeship learning via inverse reinforcement learning," in *Proceedings of the twenty-first international conference on Machine learning*, 2004, pp. 1–8.
- [4] M. Kuderer, S. Gulati, and W. Burgard, "Learning driving styles for autonomous vehicles from demonstration," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 2641–2646.
- [5] D. Sadigh, S. Sastry, S. A. Seshia, and A. D. Dragan, "Planning for autonomous cars that leverage effects on human actions," in *Robotics: Science and Systems*, vol. 2. Ann Arbor, MI, USA, 2016, pp. 1–9.
- [6] D. Sadigh, N. Landolfi, S. S. Sastry, S. A. Seshia, and A. D. Dragan, "Planning for cars that coordinate with people: leveraging effects on human actions for planning and active information gathering over human internal state," *Autonomous Robots*, vol. 42, no. 7, pp. 1405–1426, 2018.
- [7] C. Basu, Q. Yang, D. Hungerman, M. Singhal, and A. D. Dragan, "Do you want your autonomous car to drive like you?" in *2017 12th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2017, pp. 417–425.
- [8] S. Scherer, A. Dettmann, F. Hartwich, T. Pech, A. C. Bullinger, and G. Wanielik, "How the driver wants to be driven-modelling driving styles in highly automated driving," in *7. Tagung Fahrerassistenzsysteme*, 2015, pp. 1–6.
- [9] C. Pek, P. Zahn, and M. Althoff, "Verifying the safety of lane change maneuvers of self-driving vehicles based on formalized traffic rules," in *IEEE Intelligent Vehicles Symposium (IV)*, 2017, pp. 1477–1483.
- [10] M. Werling, S. Kammel, J. Ziegler, and L. Gröll, "Optimal trajectories for time-critical street scenarios using discretized terminal manifolds," *The International Journal of Robotics Research*, vol. 31, no. 3, pp. 346–359, 2012.
- [11] A. Rizaldi and M. Althoff, "Formalising traffic rules for accountability of autonomous vehicles," in *2015 IEEE 18th International Conference on Intelligent Transportation Systems*, 2015, pp. 1658–1665.
- [12] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, "Temporal-logic-based reactive mission and motion planning," *IEEE transactions on robotics*, vol. 25, no. 6, pp. 1370–1381, 2009.
- [13] A. Donzé and O. Maler, "Robust satisfaction of temporal logic over real-valued signals," in *International Conference on Formal Modeling and Analysis of Timed Systems*. Springer, 2010, pp. 92–106.
- [14] V. Raman, A. Donzé, D. Sadigh, R. M. Murray, and S. A. Seshia, "Reactive synthesis from signal temporal logic specifications," in *Proceedings of the 18th international conference on hybrid systems: Computation and control*, 2015, pp. 239–248.
- [15] C.-I. Vasile, J. Tumova, S. Karaman, C. Belta, and D. Rus, "Minimum-violation scctl motion planning for mobility-on-demand," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 1481–1488.
- [16] J. Karlsson, C.-I. Vasile, J. Tumova, S. Karaman, and D. Rus, "Multi-vehicle motion planning for social optimal mobility-on-demand," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 7298–7305.
- [17] L. I. R. Castro, P. Chaudhari, J. Tümová, S. Karaman, E. Frazzoli, and D. Rus, "Incremental sampling-based algorithm for minimum-violation motion planning," in *52nd IEEE Conference on Decision and Control*, 2013, pp. 3217–3224.
- [18] J. Tumova, S. Karaman, C. Belta, and D. Rus, "Least-violating planning in road networks from temporal logic specifications," in *ACM/IEEE 7th International Conference on Cyber-Physical Systems (ICCCPS)*, 2016, pp. 1–9.
- [19] L. Lindemann and D. V. Dimarogonas, "Robust motion planning employing signal temporal logic," in *American Control Conference (ACC)*. IEEE, 2017, pp. 2950–2955.
- [20] F. S. Barbosa, D. Duberg, P. Jensfelt, and J. Tumova, "Guiding autonomous exploration with signal temporal logic," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3332–3339, 2019.
- [21] J. Karlsson, F. S. Barbosa, and J. Tumova, "Sampling-based motion planning with temporal logic missions and spatial preferences," *IFAC-PapersOnLine*, pp. 1–7, 2020. [Online]. Available: <https://kth.diva-portal.org/smash/record.jsf?pid=diva2%3A1538022&dsid=9784>
- [22] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," *arXiv preprint arXiv:1711.03938*, pp. 1–16, 2017.
- [23] S. Shalev-Shwartz, S. Shammah, and A. Shashua, "On a formal model of safe and scalable self-driving cars," *arXiv preprint arXiv:1708.06374*, pp. 1–37, 2017.
- [24] Missouri Department of Revenue, "A guide to understanding Missouri Motor Vehicle Laws and Licensing Requirements," 2019, [Online; accessed 2020-01-09].
- [25] H. Summala, "Brake reaction times and driver behavior analysis," *Transportation Human Factors*, vol. 2, pp. 217–226, 09 2000.
- [26] S. D. Gosling, P. J. Rentfrow, and W. B. Swann Jr, "A very brief measure of the big-five personality domains," *Journal of Research in personality*, vol. 37, no. 6, pp. 504–528, 2003.