Do You Want Your Autonomous Car to Drive Like You?

Chandrayee Basu
UC Merced
cbasu2@ucmerced.edu

Qian Yang
Carnegie Mellon Univ.
qyang1@cs.cmu.edu

David Hungerman
UC Merced
dmhumgerman@ucmerced.edu

Mukesh Singhal
UC Merced
msinghal@ucmerced.edu

Anca D. Dragan
UC Berkeley
anca@berkeley.edu

ABSTRACT

With progress in enabling autonomous cars to drive safely on the road, it is time to start asking *how* they should be driving. A common answer is that they should be adopting their users' driving style. This makes the assumption that users want their autonomous cars to drive like they drive – aggressive drivers want aggressive cars, defensive drivers want defensive cars. In this paper, we put that assumption to the test. We find that users tend to prefer a significantly more defensive driving style than their own. Interestingly, they prefer the style they *think* is their own, even though their actual driving style tends to be more aggressive. We also find that preferences do depend on the specific driving scenario, opening the door for new ways of learning driving style preference.

Keywords

autonomous cars; driving preferences; driving style

1. INTRODUCTION

In the age of autonomous driving, researchers and companies are getting ever-so-close to enabling cars to generate driving behavior that includes reaching the destination while satisfying safety constraints, like not colliding with other cars or pedestrians.

Once autonomous cars attain that level of capability, initially, they might be able to generate, for each driving situation, only one solution trajectory (or behavior) that satisfies these safety and feasibility constraints. But really, many solutions exist – there are many ways to drive. This depend on the individual trade-offs that each driver makes. We have an existence proof for that. Some of us are more *aggressive* drivers, valuing efficiency and being comfortable getting close to other cars on the road. Others are more *defensive*, a bit more conservative when it comes to safety, leaving a large distance to the next car for example, or quickly braking when someone attempts to merge in front.
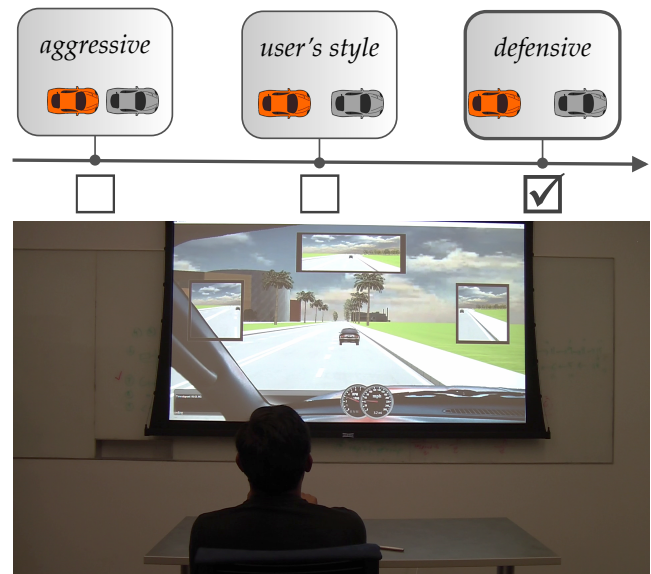
Figure 1: We first get data from user driving in different scenarios, and in a second session ask them to compare their own style (without knowing it is theirs), a more defensive style, and a more aggressive style. Participants tended to prefer a more defensive style than their own, but mistakenly thought they were actually picking their own.

Soon after we are able to generate *one* feasible behavior, we will be asking ourselves *which* behavior we should try to generate: what driving *style* should an autonomous car have? There is a natural answer to this question: cars should do what users want them to [9, 6, 22]. If the user drives aggressively, so should the car. The car should borrow the user's driving style (though not the imperfections). This is very apparent from the expression "back seat driving", which suggests that people want the driver to do what they would do.

Prior work has focused on identifying the user's driving style, via Inverse Reinforcement Learning [9, 14, 10]. Throughout it all, the underlying assumption is that we want cars to match our driving style: that we want them to drive like us.

In this paper, we challenge this assumption, and hypothesize that users want a driving style that is different from their own. We design and conduct a user study to start analyzing the potential differences between how users drive and how they want to be driven. Our study, conducted in a

driving simulator, has two parts: first, the users come in and demonstrate their driving in different environments; second, at a later date, the same users come in and test four driving styles: their own (though they do not know it is their own), an aggressive style, a defensive style, and another user's style. We measure their preference for these styles, as well as the perceived similarity to their own style.

Our results suggest that there is truth to both sides:

> Users do not actually want the car to drive like they drive. Instead, they want the car to drive like they *think* they drive.

We found a significant difference in user's own style and preferred style, with users typically preferring more defensive driving when they are passengers. However, we also found a strong correlation between the style that users preferred, and the style that users perceived as closest to their own. There was little correlation, however, between what they thought was their own style and what *actually* was their own style.

Overall, our work does not contradict the need for customization, but suggests that it might not be sufficient to learn how the user drives. Instead, we need to learn how the user actually wants to be driven. This raises challenges for learning, because we can no longer rely on demonstrations – users can easily demonstrate how they drive, but they might not be able to demonstrate the driving style they want. Instead, we need to rely on different kinds of input and guidance from users in the learning process.

Furthermore, there is a tension between what users think they want (their style) and what they actually want (a more defensive style). On the bright side, our results suggest that the learned output should be easily accepted by users: when the car drives in the preferred style, chances are users will perceive it as their own style anyway!

2. RELATED WORK

The typical behavioral patterns of a driver are usually referred to by the term *driving style*. This includes the choice of driving speed, headway, overtaking of other vehicles, or the tendency to commit traffic violations [20].

Defensiveness-aggressiveness is the most commonly used metric for defining driving style. Prior work refers to drivers as aggressive/assertive versus defensive [22]; or mild versus moderate versus aggressive [21]. In the Multidimensional Driving Style Inventory (MDSI), Taubman-Ben-Ari *et al.* identified four broad driving styles: (1) reckless and careless driving, characterized by, for example, higher speed; (2) anxious driving; (3) angry and hostile driving, characterized by more use of the horn and flash functionality; and (4) patient and careful driving [18]. Similarly, Huysduynen categorized driving style as angry driving, anxious driving, dissociative driving, distress-reduction driving and careful driving style [20]. Horswill *et al.* provided a valuable distinction between skill and style in the context of driving behaviors [8]. Hong *et al.* [7] differentiated styles in terms of defensiveness, as well as by propensity for violation of rules. Scherer defined driving style in terms of comfort [16]. Lee *et al.* [11] analyzed lane changes as a function of its severity (degree to which the vehicle in the destination lane was cut off), urgency (how soon the lane change was needed), and type classification for the full population of 8,667 lane changes.

We focus on driving style based on degree of defensiveness.

Driving style is a "humanized driving" quality [6]. Hence, most of the driving style literature relates to understanding and modeling human driver behavior, in very specific traffic situations or contexts, like lane changing [11, 15, 12], intersection crossing [7, 3, 5], car following [4], and in terms of driving actions specific to those contexts (e.g., throttle and braking level, turning) and features thereof (e.g. rate of acceleration, rate of deceleration, maximum speed in a time window). *We define driving defensiveness in our work as an aggregate of driving features in various driving scenarios.* Therefore, in our study, we present a combination of all of the aforementioned traffic conditions and scenarios to our participants.

Research on driving styles has been extended to autonomous cars in two forms. One body of work includes exploratory studies on understanding how explicitly-defined driving styles relate to comfort [16]. The second body of work encompasses research on ways to teach an autonomous car how to drive from human demonstrations [1, 23, 9, 17]. Both these groups assume that an autonomous car should learn their own user's driving style or driving behavior.

3. METHODS

3.1 Hypothesis

Because being a passenger is a different experience than being a driver, we hypothesize that:

H. *Users of autonomous cars prefer a driving style that is significantly different than their own.*

3.2 Procedure

In order to test our hypothesis, we leverage a driving simulator, and let users experience and evaluate autonomous cars with different driving styles, including their own style (without their knowledge).

We conducted a study in two parts. In the first part we collected driving data of participants in a simulation environment, so that we could let them experience their own style in the second part of the study.

3.3 Manipulated Variables

We manipulated the driving styles of autonomous cars at four levels of defensiveness: *aggressive*, *defensive*, *own style*, and a *distractor style* (a different participant's style). Users did not know if any of the styles were their own. Likewise, we also consciously avoided the use of the phrase "driving style" anytime during the studies, as well as, in the pre-study screening.

We define the *defensiveness* of the style objectively, as a function of several driving features (e.g. distance to other cars – the larger the distance, the more defensive the driving). We use features informed by existing literature. We describe them in Sec. 3.6.

We created the aggressive and the defensive styles of driving by demonstration, and then validated these styles using our driving features (see our manipulation check Sec. 3.7).

3.4 Simulator and Driving Tasks

We conducted both parts of the study in a simulation environment. Our simulation environment consisted of a standard classroom projection screen and table in front of the screen fitted with Logitech G920 steering wheel, brake, and gas pedal. We used the OpenDS driving simulation software [13] for running each of the driving simulations. The simulation platform was set up on a standard PC augmented with NVIDIA GeForce GTX 1070 and was hidden from the participants' view.

In the first part, the participants drove on a 9.6 mile long test track that consisted of 14 different driving tasks designed using the City Engine software (Fig.2).

We define a *driving task* as a sequence of driving maneuvers in response to specific traffic conditions. For each task there are two to three simulated traffic conditions that resemble everyday traffic, so as to elicit natural driving from the participant.

In the second part of the study, the participants experienced 6 of these 14 tasks, each performed by autonomous cars of four different styles.

3.5 Procedure

Before the driving session in part one of the study, we familiarized participants to the driving simulator. We asked each participant to practice on two different test tracks until they felt that they were driving as they would in their everyday driving. The first track had several traffic signals and turns, and second one was on congested city roads with several traffic cars. Their driving was assisted by a voice navigation. There were road signs for speed change zone, speed limit, sharp turns, entry to expressway and exit from expressway. We instructed the participants to drive as they would on actual roads and to treat the speed limits the way they would in their usual driving. This practice session lasted 5-10 minutes.

Participants then began the first part of the study, which consisted of 15-20 minutes of driving along the 14 tasks-test track, followed by a 10 minute interview.

In the second part of the study, the autonomous cars performed six tasks (combined into four test tasks) from this list with the participant as a passenger, shown in bold letters on the list in Fig.2. To simplify, we combined the second and the third tasks in the list i.e. lead car slow down forcing lane change and merge back to right lane into a single test task, which we refer to as Task 1 in the rest of the paper. Likewise, we combined the sixth and the seventh tasks into a single test task, called Task 2 in the rest of the paper. Thus, each autonomous car performed four test tasks in total. Two of the test tasks were on the expressway and lasted approximately 4 minutes for each style and the other two tasks on the inner city roads were shorter than 2 minutes.

After the participants had driven in an autonomous car of each driving style for each of the test tasks, we conducted a short interview-based survey with each participant.

3.6 Dependent Measures

Perceived similarity to real driving. In the first part of the study we conducted a post-driving open-ended interview with the participants to understand whether the manual driving in the simulation environment resembled their everyday driving. We asked three questions in this interview, each followed by a request for more elaboration. We

asked the following questions in the interview:

1. Did you enjoy the drive?

2. Are there any positive or negative aspects of the simulation environment, the driving controls and the traffic conditions that you would like to mention?

3. On a scale of +3 to -3 [2], please rate how similar or different is this experience from your daily driving?

Open-ended responses. In the second part we asked each participant to think aloud about their emotions and feelings as they were experiencing autonomous driving.

Main subjective measures: Preference and perceived similarity to own style. After a participant had experienced each autonomous style for a given task, we conducted an interview-based survey. We asked the participants to rate each style of driving for *comfort*, *safety*, *preference for everyday use*, and *similarity with their own driving* on 7 point Likert scale.

Main objective measures: Driving style features and overall defensiveness. We measured the user's style quantitatively using task specific driving features, derived from existing literature. We carefully considered the contexts and subject demographics of each of these existing studies to ensure as much similarity in the context as possible with our study.

For car following, lane changing, and return to preferred lane, we selected the features described by Lee *et al.* in "A Comprehensive Examination of Naturalistic Lane-Changes" [11]. This study analyzed the largest naturalistic lane change dataset and specifically labelled lane change data resulting from the slowing down of the leading car. The speed range of 45 mph to 55 mph matches our driving conditions. Their dataset consisted 8667 lane changes over 23,949 miles of driving from 16 commuters of age group 20 to 60. They studied car following, lane changing, and return of preferred lane in terms of distance, time to collision, and relative speed classified by severity and urgency of lane change.

The features for tasks like turning at the intersection with a green light or stop light were derived from our preliminary interview with the participants and from Hong *et al.* [7] and Banovic *et al.* [3].

Table 1 summarizes all the features for the four driving test tasks. We used *mean distance to lead car*, *mean time headway*, *time headway during lane change*, and *distance headway during lane change* as features for Task 1 and Task 2. Task 1 had an extra feature *distance headway merge back* for scoring the merge back behavior to the right lane.

Task 3 consisted of two sub-tasks (approaching intersection at a stop light and then making a left turn at green ball). We characterized this task with 5 features: *Braking Distance from the intersection*, *Average speed for 20 meters before intersection*, *Time To Stop*, *Speed at the intersection*, and *Maximum turn speed*.

Task 4 constituted approaching intersection at green ball and then turning right without stopping. The features for this tasks are *Speed at the intersection* and *Maximum Turn Speed*.

We objectively measured a participant's overall driving style in terms of a **Defensiveness Score**. We first normalized the feature values across participants for each feature irrespective of the task. We calculated a *Defensiveness Score* for

1. Merge onto expressway between two cars
2. **Lead car slow down forcing lane change**
3. **Merge back to right lane**
4. < 90 degrees Turn negotiation on expressway
5. Stop light at intersection on expressway
6. **Slow lead car in right lane forcing lane change when another car is approaching fast in the destination lane**
7. **Merge back into right lane with a continuous traffic moving at constant gap and constant speed**
8. Expressway exit
9. Sharp right turn no signal
10. Left turn on green light
11. Right turn at stop light
12. **Left turn at stoplight and yield at green**
13. **Right turn on green light**
14. Park along curb

Figure 2: Designed track: Tasks (shown in the list below the figure) are indicated in square brackets. Total road stretch is 9.6 miles.

each participant and for each task as the average over all the normalized feature values for that participant and task. We then computed an *Aggregate Defensiveness Score* for each participant by averaging their scores across the four test tasks.

3.7 Manipulation Check

We performed a manipulation check on our aggressive and defensive driving styles. We measured the aggregate defensiveness score for each style, plotted on the bottom right of Fig.3. We found that indeed, the aggressive style was less defensive than the defensive style (lower defensiveness score). We found that 86.67 % of the users' styles scored higher than the aggressive style, and lower than the defensive style. This suggests that the two reference driving styles created by demonstrations resulted in meaningful representations of aggressive and defensive driving.

3.8 Participants

Subject Allocation. We opted for a within-subjects allocation because the participants needed to choose a preferred style out of the set of available ones. We randomized the order of the conditions.

Demographics. We recruited 15 participants consisting of a mix of graduate students and undergraduate students. Before the study we sent out a screening form to each participant in order to ensure a wide distribution of demographics, driving experience and perceived driving behaviors of the participants. We also checked for a valid driving license. 3 of our participants were 30 to 31 years old, the rest of the participants were 18-24 years old.

The mean driving experience of the participants was 5.46 years with a standard deviation of 4.5 years. Participants had driven an average 214 miles with a standard deviation of 188 miles on the week before they filled out the screening form.

We asked the participants to give us some information about their perceived driving behavior using the following questions: 1. Please rate if you consider yourself a conservative or an adventurous driver on a 7-point scale, 1 being *conservative* and 7 being *adventurous*. 2. Please rate on a 7-point scale what you like about driving, 1 being *joy of motion* (like feeling the force as you accelerate) and 7 being *comfort of steadiness*. You may like some of both. 3. Rate on a 7-point scale if you think you vary your driving by road conditions, traffic and time availability, 1 being *vary always* and 7 being *I don't vary at all*. 4. Please rate your driving experience from somewhat experienced to very skillful. The purpose of these questions was to acquire some information about the participants' driving styles without explicitly using the term style or in other words give away the original goal of the study.

Approximately 46 % of the participants considered themselves well experienced in driving, and 20 % considered themselves experienced. The rest were equally distributed between somewhat experienced to very skillful. The mean score for perceived conservative-adventurous driving behavior was 3.6. Most of the participants considered themselves to be in the middle of the spectrum. Only one participant considered himself to be conservative. More participants preferred comfort and steadiness over joy of driving, the average rating being 4.46. The mean rating for variation of driving style in response to environment and traffic was 3, which means most participants believed that they alter their driving behavior according to traffic.

3.9 External Validity and Controlling for Confounds

Driving environment. We used a simulator and not real autonomous cars. However, we designed a simulation track and traffic conditions so as to elicit natural driving responses. We also collected participant feedback in the first part of the study on the simulation environment and how their driving behavior in the simulated track related to their actual driving behavior.

Features	Definitions
Mean Distance to Lead car	During car following (with 200 meters distance) the average distance between middle of the driver car and the lead car.
Mean Time Headway	During car following (with 200 meters distance) average time headway, defined as ratio of Distance headway and speed of the driver car.
Time Headway during Lane change	Distance headway divided by the speed of the driver car during lane change.
Distance Headway during Lane change	Distance between the middle of the driver car and the lead car during lane change
Distance Headway Merge Back	This is the same as Distance Headway during lane change except measured in between driver car and the following car in the destination lane.
Braking Distance from the Intersection	The distance from the intersection at which a person starts applying brakes.
Time To Stop	Braking distance divided by the speed of the car right before brake is applied.
Maximum Turn Speed	Maximum speed of the driver car over a time window during a left turn or a right turn.
Speed at the Intersection	Instantaneous speed at the intersection.
Average Speed for 20 meters before Intersection	This is the speed of the driver car averaged over a distance range of 20 meters from the intersection.

Table 1: Features for style classification

Masking own style. One of the major challenges of this work was to ensure that a participant could not recognize his or her driving style from simulation peculiarities like scenes, traffic and controls. We wanted the participants to only recognize their driving style based on their traffic maneuvers and actions. We took several steps to camouflage the driving data of a participant in the second part of the study:

- We retained the traffic conditions and route from the first part of the study while changing the surrounding scenes and traffic cars, such that we can replicate the user's driving while removing the bias of familiar environment.

- We let the participant perform approximately 14 driving actions in the first part of the study and picked only some of these tasks for the second part of the study.

- During the second part of the study we presented the tasks in an order different from how they occurred in the manual mode. For example: In the first part, the participants first entered the expressway and performed some driving actions on the expressway and then exited the expressway and performed some more driving maneuvers on the city roads. During the second part, we presented one city road task and one expressway task in an alternate order.

- We presented the four styles for each of the test driving tasks in a randomized order, which made it more difficult to consistently recognize one style.

- We post-processed the users' driving to remove peculiarities, which we explain below.

During our pilot studies we found that due to some peculiarities of simulation environment (over-sensitive steering, less sensitive braking) and the resultant jitter in the driving data, some participants were able to recognize their own driving. For example, a participant mentioned: *"This looks like how I was driving. I had to stop at the intersection because I pressed the brake too early. The brake was tight."* Idiosyncrasies of the simulator led pilot participants to identify their driving behavior in the second part of the study. In order to eliminate these peculiarities of the simulation environment, we changed the brake stiffness and steering sensitivity and presented participants with smoothed version of their data in the second part of the study.

3.10 Trajectory Smoothing

We filtered the driving trajectories to eliminate idiosyncrasies that make the trajectory instantly recognizable.

We applied a Bilateral Filter [19] to reduce the lateral variance (or equivalently, the variance of the lateral displacements from the center of the lane) of the trajectories. By affecting only the lateral components of the trajectory, this filtering preserves distance between the cars. We applied filtering only to the stretches of the trajectory on the expressway.

Fig.4 shows a smoothed trajectory for one participant. It has 15 % lower lateral variance than the original trajectory.

4. RESULTS

4.1 Simulation Realism

In the first part of the study, in addition to collecting user driving data, we also wanted to ensure that this driving data corresponded to participants' everyday driving as much as possible. We conducted a post-driving interview, as described in the Dependent Measures subsection (Sec. 3.6). Here we present the results of the interview.

The rating mode for similarity between driving on road and driving in our study simulator was +1 on -3 to +3 scale. Four participants gave a rating of +2. Some of their positive comments were: "Not considering the room environment and just looking at the simulation graphics and the car it was pretty much the same environment as real. I would give +3
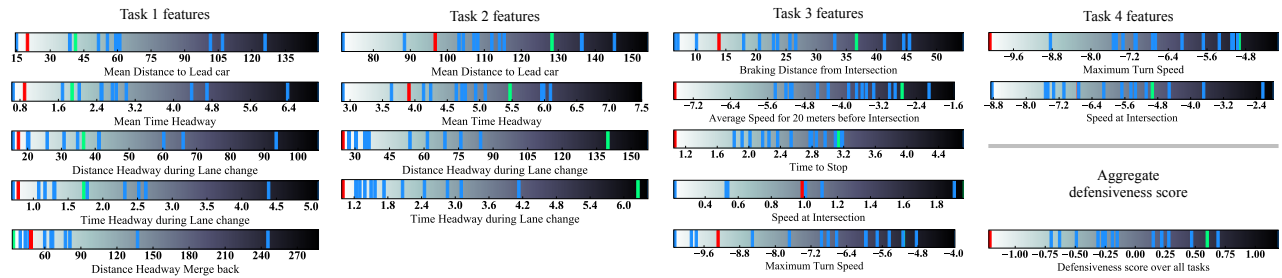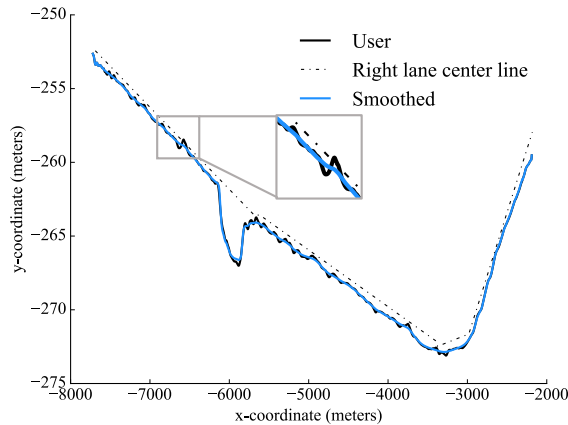
Figure 3: Participants' feature distribution

Figure 4: Smoothed trajectory compared to original trajectory of task 1 of one participant at 15 % smoothing

for surrounding traffic conditions". Other participants said that they felt relaxed in the simulator environment and that they could drive cautiously as they would in real traffic.

One participant who rated the driving experience similarity -2 complained about the lack of motion feedback in the system. This is the same participant who gave high rating for *joy of motion* in the screening question. However, no other participant had the same concern and got well-adjusted to the simulation environment.

Most of the participants who rated +1 to -1 found steering re-centering or brake insensitivity difficult. We also received quite opposite feedback from two participants when they compared their everyday driving to the simulator driving. For example, one participant mentioned "It felt real. It was something I could get used to after driving a while. The gas and brakes were more sensitive than my car". Another participant felt that the brakes were excellent, different from regular car.

One participant reported that she was so immersed after driving for a while, that she caught herself turning her head back to check for oncoming traffic in the destination lane. We found that participants with one or less years of driving experience could not use the simulation environment properly.

Overall, the ratings and the comments supported that the simulator conditions are not *too* far from real conditions.

4.2 Feature Distribution for Participant Styles

We define driving style in terms of features mentioned in

the Sec. 3.

Fig.3 shows, for each task, feature, and participant, what the participant's feature value was for that task (blue marks). The figure also shows the aggressive style values in red and defensive style values in green.

Higher negative values correspond to more aggressive behavior. All the feature values are arranged from aggressive on the left to defensive on the right. However, for features like speed where lower values mean more defensive we show and use the negation of these features.

The bottommost plot to the right shows the aggregate defensiveness score. This score is derived from the normalized feature values. 60 % of the participants are within 0.75 standard deviation aggressive and 40 % within 0.75 standard deviation defensive. Only two of the participants were more defensive than the autonomous defensive car, one of them being very close to the defensive car in the score.

> *When looking at the aggregate defensiveness, most participants lie between the aggressive and defensive styles.*

There are, however, exceptions, but for particular features in particular tasks. For task 1, several participants were more defensive than the defensive autonomous car. For the last feature of task 1, Distance Headway Merge Back, the aggressive car was not as aggressive as several participants and even our defensive car. In task 2, the aggressive and the defensive autonomous cars enclosed a middle section of the spectrum for Mean Distance Headway and Mean Time Headway. In other words, several participants were more aggressive and more defensive than the aggressive and defensive autonomous cars respectively. This is because these features were measured during car following over a long time span and are expected to have wider distributions than features characterizing instantaneous actions.

4.3 Preferred Style in Relation to Own Style

We asked participants to rate how much they would prefer driving with each style, for each task. We refer to the highest rated style(s) as the participant's preferred style(s).

Our main finding is that overall, users preferred a different style than their own. A total of 9 out of 15 participants preferred a different style than their own on at least one of the tasks. A matched pairs *t*-test comparing actual and preferred defensiveness score showed a significant difference ($t(1, 60) = -2.58$, $p = .0121$), supporting our hypothesis. Here, whenever a user's highest rating was for multiple styles as opposed to a single one, we included each
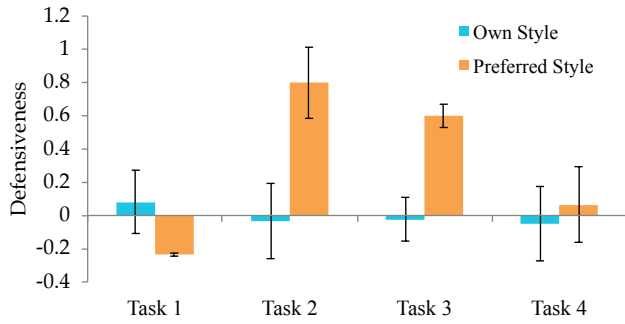
Figure 5: Mean Defensiveness Score Across Participants. The corresponding scores of aggressive and defensive autonomous cars are Task 1: (-0.768, -0.222) Task 2: (-0.885, 1.325), Task 3: (-1.82, 0.766) and Task4: (-1.49,0.72).

Figure 6: Scatter plot showing correlation between the style that users *thought* was their own and the style that they chose as their preferred.

Figure 7: Scatter plot showing little correlation between own style and perceived own style: users did not tend to identify their own style correctly. as evidenced by the off-diagonal points.

preferred style as a data point.

> *Overall, people prefer a significantly more defensive style than their own.*

We also investigated how this breaks down by task, and only found significant effects on the 2^{nd} and 3^{rd} tasks. See Fig.5 for comparison between average preferred style and own style of our participants for each of the four tasks. For task 1 we note that several participants were more defensive than other autonomous styles presented to them. However, they still preferred our defensive style, which explains that the average choice was more aggressive than the participants' own style. Interestingly, some participants did not perceive the extra defensive nature of their own style in task 1 positively. One participant mentioned about their own style that "In this one I felt like we gave a lot of room, more than I would have probably." (ironically, since they did *exactly* that). Two other participants made similar comments about their own lane changing behavior. Besides, a few participants also considered driving features beyond the ones we accounted for.

For task 2 and task 3 the defensive autonomous car was more aggressive than only 0 - 3 participants across all features and it was more defensive than the rest of the population by a major margin, in features like Distance Headway and Time Headway During Lane Change.

The task had a significant effect on the difference ($F(3, 58) = 4.13$, $p = .0101$), suggesting that people's preferences for a driving style are not consistent, but rather *change based on the context.* This motivates future research on predicting the desired driving style not just based on the individual, but also based on the current driving context.

4.4 Perceived Own Style in Relation to Actual Own Style

We also asked participants to rate each style in terms of similarity to their own. From this, we learned what participants *perceived* their own style to be.

We found that even though participants did not pick their *actual* style as their preferred (Sec. 4.3), participants did tend to prefer their *perceived* style. On each task, between 80 and 93% of participants opted for the same style as the one they *thought* was the closest to their own (and sometimes rated other styles as well as equally good). We found a significant
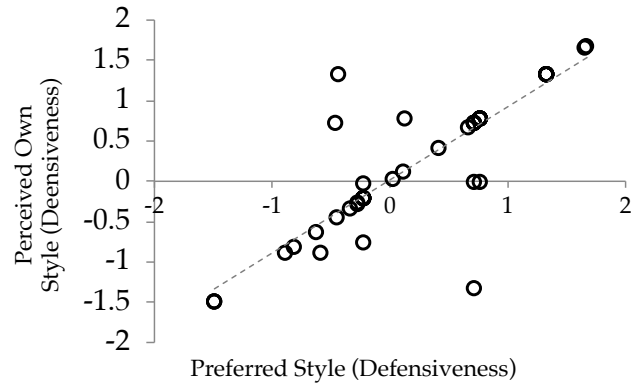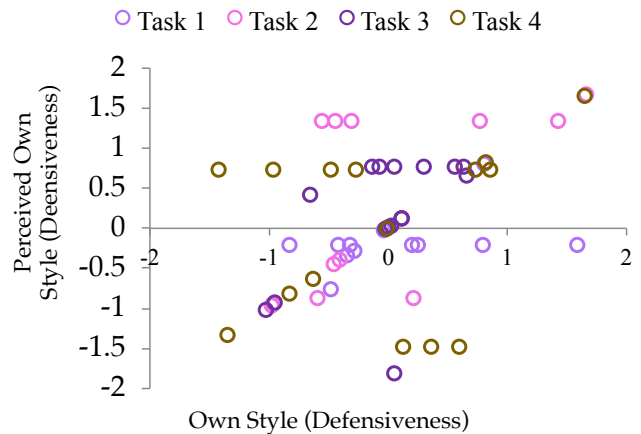
correlation between the perceived own and preferred styles, $r(58) = .86$, $p < .0001$. Fig.6 shows a scatter plot of preferred style by perceived style, with many points on the diagonal representing users who preferred driving in the style they thought was (closest to) their own.

However, even though the majority participants thought that they were picking their own style, they really were not. A total of 46 to 67% participants on each task did not correctly identify their actual own style, and the correlation between perceived and actual defensiveness score was only $r(56) = .40$ across tasks. Fig.7 paints a different picture from Fig.6: it plots the perceived style against the *actual* own style, showing many off-diagonal points, representing users who did not correctly identify their style.

In task 1 we see that several participants perceived themselves to be slightly more aggressive irrespective of their actual style. Likewise, both for task 2 and task 3 several participants perceived themselves to be more defensive irrespective of their actual style.

Participants tended to prefer the style that they thought was their own, but in fact that style had little correlation to their actual own style.

5. DISCUSSION

Summary. We hypothesized that users of future autonomous cars would prefer a driving style that is significantly different than their own. We conducted a user study in a driving simulator to test our hypothesis. We found that users preferred a more defensive style than their own. This echoes the finding from prior work [8] that when people are not in control of the driving they prefer lower speeds – autonomous cars are one instantiation of not being in control of the driving.

Interestingly, over 80% of users preferred the style that they *thought* was their own, but many times they were incorrect in identifying their own style. These results open the door for learning what the user's preferred style will be, but caution against getting driving demonstration from the user, since people can drive like they do, not like they want to be driven.

Limitations and Future Work Our work is limited in many ways:

- **Limited driving style features**. Following the most common conventions, we have only characterized style in terms of defensiveness. We also inherited from previous studies the feature choices in defining driving styles.

- **Limited driving style choices**. We presented participants with limited options along the spectrum of defensiveness and found that they preferred a style more defensive than their own. However, we did not learn the style they actually desired, only the best out of our few options.

- **Limited fidelity of simulation environment**. Our simulation environment does not provide motion feedback, which may limit the users' perception of speed. Although the interview results validated that participants' perception of the driving styles are sufficient, experiment results in a higher fidelity simulation environment might be more accurate.

Given the encouraging results from the our findings presented here, we believe that it is worthwhile to test more diverse feature choices and driving style representations in a higher fidelity setting. These experiments will provide a comprehensive evaluation of the study presented in this paper.

Going further, we are excited to investigate how we might learn a deviation from the user's driving style that is predictive of how they actually want to be driven, and explore learning techniques that augment user demonstrations with other types of user input and guidance.

6. ACKNOWLEDGEMENTS

7. REFERENCES

[1] P. Abbeel and A. Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1. ACM, 2004.

[2] ASHRAE. Standard 55 - thermal environmental conditions for human occupancy, 2010.

[3] N. Banovic, T. Buzali, F. Chevalier, J. Mankoff, and A. K. Dey. Modeling and understanding human routine behavior. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 248–260. ACM, 2016.

[4] M. Brackstone and M. McDonald. Car-following: a historical review. *Transportation Research Part F: Traffic Psychology and Behaviour*, 2(4):181–196, 1999.

[5] M. Elhenawy, A. Jahangiri, H. A. Rakha, and I. El-Shawarby. Modeling driver stop/run behavior at the onset of a yellow indication considering driver run tendency and roadway surface conditions. *Accident Analysis & Prevention*, 83:90–100, 2015.

[6] HERE360. How to humanize the autonomous car. http://360.here.com/2015/04/23/humanized-driving/, 2015.

[7] J.-H. Hong, B. Margines, and A. K. Dey. A smartphone-based sensing platform to model aggressive driving behaviors. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pages 4047–4056. ACM, 2014.

[8] M. S. Horswill and F. P. McKenna. The effect of perceived control on risk taking1. *Journal of Applied Social Psychology*, 29(2):377–391, 1999.

[9] M. Kuderer, S. Gulati, and W. Burgard. Learning driving styles for autonomous vehicles from demonstration. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2641–2646. IEEE, 2015.

[10] C.-P. Lam, A. Y. Yang, and S. S. Sastry. An efficient algorithm for discrete-time hidden mode stochastic hybrid systems. In *Control Conference (ECC), 2015 European*, pages 1212–1218. IEEE, 2015.

[11] S. E. Lee, E. C. Olsen, and W. W. Wierwille. A comprehensive examination of naturalistic lane-changes. Technical report, 2004.

[12] H. M. Mandalia and M. D. D. Salvucci. Using support vector machines for lane-change detection. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 49, pages 1965–1969. SAGE Publications, 2005.

[13] OpenDS. Welcome to opends 4.0! https://www.opends.eu/home, 2016.

[14] D. Sadigh, S. S. Sastry, S. A. Seshia, and A. Dragan. Information gathering actions over human internal state.

[15] D. D. Salvucci. Inferring driver intent: A case study in lane-change detection. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 48, pages 2228–2231. SAGE Publications, 2004.

[16] S. Scherer, A. Dettmann, F. Hartwich, T. Pech, A. C. Bullinger, J. F. Krems, and G. Wanielik. How the driver wants to be driven-modelling driving styles in highly automated driving. *Tagungsband*, 7:2015–26, 2015.

[17] D. Silver, J. A. Bagnell, and A. Stentz. Learning from demonstration for autonomous navigation in complex unstructured terrain. *The International Journal of Robotics Research*, 2010.

[18] O. Taubman-Ben-Ari, M. Mikulincer, and O. Gillath. The multidimensional driving style inventoryâĂŤscale construct and validation. *Accident Analysis & Prevention*, 36(3):323–332, 2004.

[19] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Computer Vision, 1998. Sixth International Conference on*, pages 839–846. IEEE, 1998.

[20] H. H. van Huysduynen, J. Terken, J.-B. Martens, and B. Eggen. Measuring driving styles: a validation of the multidimensional driving style inventory. In *Proceedings of the 7th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, pages 257–264. ACM, 2015.

[21] L. Xu, J. Hu, H. Jiang, and W. Meng. Establishing style-oriented driver models by imitating human driving behaviors. *IEEE Transactions on Intelligent Transportation Systems*, 16(5):2522–2530, 2015.

[22] N. M. Yusof and J. Karjanto. Comfort determination in autonomous driving style. In *3rd Workshop on User Experience of Autonomous Vehicles at AutoUIâĂŹ15*, pages 257–264. ACM, 2015.

[23] B. D. Ziebart, A. L. Maas, A. K. Dey, and J. A. Bagnell. Navigate like a cabbie: Probabilistic reasoning from observed context-aware behavior. In *Proceedings of the 10th international conference on Ubiquitous computing*, pages 322–331. ACM, 2008.