Communicating Intent on the Road Through Human-Inspired Control Schemes

Katherine Driggs-Campbell and Ruzena Bajcsy

***Abstract*— Given the current capabilities of autonomous vehicles, one can easily imagine autonomy released on the road in the near future. However, it can be assumed that the transition will not be instantaneous, meaning they will have to be capable of driving well in a mixed environment, with both humans and other autonomous vehicles on the road. This leaves a number of concerns for autonomous vehicles in terms of dealing with human uncertainty and understanding of cooperation on the road. This work demonstrates the need for focusing on communication and collaboration between autonomy and human drivers. After analyzing how drivers perform cooperative maneuvers (e.g. lane changing), key cues were identified for conveying intent through nonverbal communication. It was found that human observers can predict lane changes with over two seconds in prior to the lane departure, without use of a turning signal. Building on this concept, an autonomous control scheme is proposed that aims to capture these subtle motions before executing a lane change. To compare the proposed human-inspired methods, three possible control schemes for autonomous vehicles are implemented for a validation study on human subjects to provide feedback on their experience. By properly conveying intent through nuanced trajectory planning, we show that drivers can predict the autonomous vehicle's actions with 40% increase in prediction time when compared to traditional control methods, both as a passenger and while observing the autonomous vehicle.**

I. Introduction

Recently, there have been a number of breakthroughs in autonomous driving [2], [16], [17]. Although there have been many great advancements in sensing and control that have lead to many successful implementations [18], one major open area of research is understanding human drivers. While it is generally assumed that autonomy will be publicly available in the near future, it can be assumed that the transition will not be instantaneous [12]. This means autonomous vehicles must be able to drive well in a mixed environment, with both humans and other autonomous vehicles on the road. This leaves a number of concerns for autonomous vehicles in terms of dealing with understanding and predicting human drivers as well as interacting with them. It's also important for the passengers understand the intent of the autonomous vehicle, which has been shown to improve the acceptance of autonomy [20].

This material is based upon work supported by the Office of Naval Research MURI Award ONR-N000141310341 and DURIP Award N000141310679. The study protocol was approved by the University of California Institutional Review Board for human protection and privacy, under Protocol ID 2013-07-5459. Each subject was first informed on the experimental procedure and written informed consent was obtained.

Katherine Driggs-Campbell and Ruzena Bajcsy are with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720 USA (e-mail: {krdc,bajcsy}@eecs.berkeley.edu).

There has been a great deal of research in looking into driver modeling looking within the ego vehicle [13], [19], as well as studying driver perception [10]. However, there are still many fundamental problems that have not been addressed, including predicting driver's behaviors and how to interact and communicate with human drivers on the road. As humans drive, we are able to assess what other vehicles are likely to do and how they will likely respond to various actions. For example, when driving on the freeway, many drivers can estimate when nearby drivers want to change lanes, even without a turning signal, which will be shown in Section IV. Through experience, we learn cues from driver's motions and gain intuition of how our actions will influence other's behaviors. If an autonomous system is going to be introduced into this nuanced social system, we hope that the system will have similar intuition and understanding, and behave in an expected way.

Many researchers have attempted to address this by estimating driver intent, ranging from driving monitoring techniques to model based approaches [5], [11]. In [8], we attempted to estimate the driver intent by building a dataset of lane changes, where the driver subjects actively labeled their mode of intent as they were driving. Using these labels, it was possible to accurately predict the driver intent based on the observable states of surrounding vehicles. Meaning that by observing the vehicle and the states of nearby vehicle, it is possible to identify the driver state.

One of the key findings of this study was that human drivers convey their intent through their motions. Seconds before a driver begins executing a lane change maneuver, the driver will edge towards the next lane until there is sufficient space to safety merge. This is intuitive, as this preparation motion can be thought of as when a driver turns on her turning signal and will communicate their intent to surrounding drivers. As was mentioned, as a seasoned driver drives, she can predict that a vehicle in the next lane wants to change lanes even when no turning signal is used.

This concept of conveying intent through motion is well supported by studies in neuroscience and in Human-Robot Interaction. In [3], it was found that intent through motion is incredibly important in social settings when humans are interacting. Similarly, it was shown in [6] that by motion planning with intent in mind will lead to more understandable interactions between humans and robots.

Here, we wish to examine what effect of communicating intent through lane positioning has on predictability for passengers within an autonomous vehicle as well as surrounding vehicles. Incorporating nuanced motions will hopefully lead to better social acceptance and understanding when released on the roads with other human drivers. This is similar to the work presented in [1], where qualitative driver behaviors were learned using inverse reinforcement learning. The work presented here attempts to capture behaviors in the continuous space by integrating aggregated driver data into a control scheme, rather than learn discrete actions.

By using the data collected in [8], we wish to mimic the human driver's motions to capture the subtle communication that occurs in cooperative and collaborative manuevers. We do this by defining a human-inspired control scheme that controls the vehicle in a similar manner. Then, we validate the advantage of this control scheme over traditional controllers and human controlled vehicles through user studies.

This paper is organized as follows. Section II presents the formalization of this work and Section III describes how the control scheme was derived from human driving data. The validation study results are presented in Section IV. Finally, the discussion and future work is presented in Section V.

II. Methodology

In this section, the formulation of the problem will be presented, as well as the data used to understand the driver.

A. Modeling Driving as a Hybrid System

In this formulation, we envision the human driver (or autonomous vehicle) as a hybrid system that controls the vehicle differently depending on the mode of intent. Suppose we are given a vehicle with some dynamics:

$$x_{k+1} = f(x_k, u_k) \tag{1}$$

where $x_k \in \mathbb{R}^n$ is the state of the vehicle at time step $k \in \mathbb{N}$ and $u_k \in U$ is our input from our input space U, which we assume is given as a close compact set containing the origin.

We also assume we are given the set of constraints or safe regions in the environment, denoted $\mathcal{C}$, and some function that will tell us whether or not the safety constraints are satisfied:

$$\psi(x_k, \mathcal{C}_k) = \begin{cases} 1, & \text{if } x_k \in \mathcal{C}_k \\ 0, & \text{otherwise} \end{cases} \tag{2}$$

It is assumed that the states of surrounding vehicles are included in the constraint set, as well as a prediction of what the vehicles will do. Additionally, current and future road information is considered given.

When we consider how the vehicle is controlled by humans, we suppose that the control law for a given vehicle changes depending on the mode of intent. This can be thought of as a high level decision making function, that determines what the best course of action is given the scenario. This implies that depending on what high level action the driver wants to execute, the control law will change. This is shown in Algorithm 1.

Formally, we can assume that the input is defined by some algorithm that is dependent upon the discrete mode $q \in Q$, which is assumed to be known (or estimated) from a defined, finite set of modes of intent:

$$u_k \leftarrow \mathcal{A}_q(x_k, \mathcal{C}_k) \tag{3}$$

where all variables are as previously defined. It is assumed that this algorithm will find the optimal solution for the input to the vehicle, or return no solution.

The modes of intent used in this work build off of data collected in [8] that aimed to estimate driver intent.

Algorithm 1 Autonomous Control Scheme

Initialize Variables
for each time step, k **do**
  $x_k \leftarrow$ update_vehicle_state()
  $\mathcal{C}_k \leftarrow$ update_constraints()
  $q \leftarrow$ determine_mode$(x_k, \mathcal{C}_k)$
  $u_k \leftarrow \mathcal{A}_q(x_k, \mathcal{C}_k)$ ▷ Compute Optimal Control
  execute_control(u_k)
end for

B. Identifying Driver Modes

This section presents the driver model which models the discrete modes for the lane changing manuever. In [8], we presented a driver model that is able to identify the following modes of behavior: *lane keeping*, *preparing to change lanes*, and *lane changing* (see Fig. 1). This was executed using observable features in the environment and human labeled data to classify what mode the driver was in.

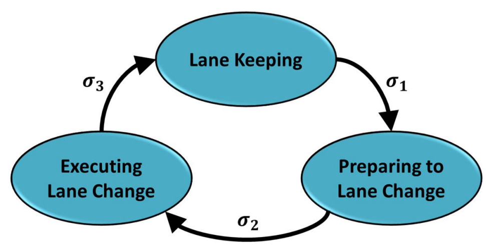

Fig. 1: Illustration of Driver Modes, as presented in [8].

The key here was using labeled data, meaning that the driver indicated which mode they were in as they were driving in order to capture each driver's decision making process. Using this dataset, we were able to effectively detect when drivers changed modes. The detection is based on environmental cues, designed to give us insight to the decision making process of the human driver. The resulting identification algorithm attempted to be as flexible and as portable as possible, meaning that it didn't rely directly on the control actions or state of the driver.

In this dataset, ten subjects were asked to execute lane changes, resulting in over 200 example lane changes per driver. The following features were collected, which we will use to understand the driver behavior in each mode: (1) ego vehicle information, including vehicle states and inputs; and (2) environment constraints, including road boundaries and observable, relative states of surrounding vehicles.

One of the key findings of this initial study was that driver's convey their intent through motion. It was observed that prior to executing a lane change, humans will edge over to the next lane, signaling to surrounding drivers their intent to change lanes. As shown in Figure 2, the distributions associated with these two modes are distinct.

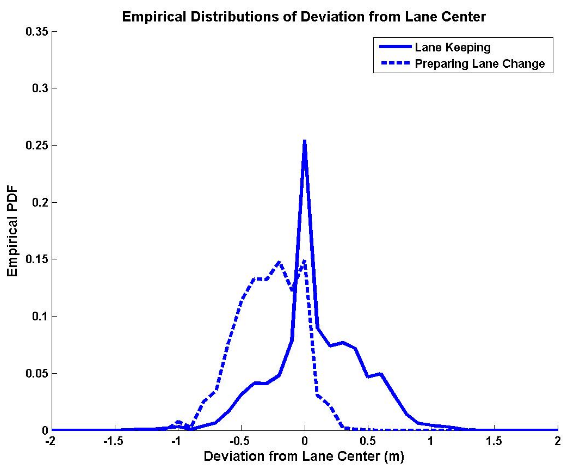

Fig. 2: Distribution of position in lane by mode, as presented in [9]. It can be noted that the driver edges toward the next lane while preparing to change lanes.

Using this data, our goal is to understand how the control law (or cost function) of the vehicle changes with respect to driver mode as well as the changes in the environment (e.g. distance to the lead vehicle), to better understand the communication and negotiation that occurs before a lane change. Thus, we further consider the subtle behavior in these *lane keeping* and *preparing to lane change* modes and how to easily incorporate this into control frameworks.

III. Resulting Controllers for Driver Modes

Given that the given dataset consists of multiple drivers with unknown distributions, we wish to further analyze how the drivers behave in the different modes. To do this, we utilize a concept similar to a cost map, which is built by discretizing the space and looking at how frequently the vehicle passed through a particular location. When staying in lane, this is defined by the position in the lane and the distance to the lead vehicle. This can be thought of as the spatial, empirical distributions associated with each driver mode. By looking at the regions that the drivers frequently inhabit, we can analyze how they generally behave. These empirical cost maps for *lane keeping* and *preparing to change lanes* are shown in Figure 3. It can be noted that the although the modes are similar, drivers tend to communicate their intent by edging toward the next line prior to indicating that they are changing lanes.

By analyzing this cost map, we see that we can simplify the problem by assuming that the drivers wish to follow some nominal trajectory, given by the empirical distribution on the cost maps. This is identified by finding the expected lateral position, associated with a longitudinal coordinate (i.e. expected lane position given a distance to the lead vehicle). The smoothed nominal trajectory is shown as the pink line in the distributions in Figures 3.

A. Control Implementation

In this paper, it is assumed that this control algorithm takes the following form, where $u_{k,\dots,k+\bar{N}}$ is the output of an optimization program:

$$\begin{array}{ll} \underset{x,u}{\text{argmin}} & J_q(x,u) \\ \text{subject to} & x_i = f(x_{i-1}, u_{i-1}) \\ & \psi(x_i, \mathcal{C}_i) > 0 \\ & u_i \in U \\ & \qquad \forall i = \{1,\dots,\bar{N}\} \end{array} \tag{4}$$

where $J_q(x,u)$ is the cost function that is defined for each mode q, x and u are a concatenated vectors of states and inputs from time step 1 to $\bar{N}$, which is the pre-defined time horizon, x_0 and u_0 are assumed to be given, and all other variables are as previously defined. In essence, this finds the optimal control over the next $\bar{N}$ time steps, given our safety constraints, input limits, and initial conditions. This implementation is can be thought of as similar to a Model Predictive Control framework[1][4].

Given that we can effectively identify the mode of intent, the underlying cost function or control scheme must be identified. There are many advanced techniques for identifying the cost function of a system, but many become infeasible when dealing with highly noisy data [15]. We note that there are extensions to many learning methods that include noisy models, but often a distribution must be assumed. From this dataset, it can be shown that driver's do not always follow known distributions, and particularly when looking at a collection of different drivers.

Thus, for simplicity, our cost function is assumed to be of the form:

$$J_q(x,u) = (x - x_q)^\top P(x - x_q) + u^\top Q u \tag{5}$$

where $x_q \in \mathbb{R}^n$ is our desired nominal position trajectory associated with mode q that is zero padded to account for the vehicle states other than position and velocity, P and Q are weighting matrices to tune the costs on the states and inputs respectively (which for simplicity are set to identity), and all other variables are as previously described. The velocity of the vehicle is set to be 15 m/s, to match the conditions set in the original data collection process.

The resulting scheme is a trajectory following framework, where the most significant change between modes is the x and y position. Given the cost map above, we can compute the expected lateral lane position given a distance to the lead vehicle, which we use as a nominal trajectory to follow in

[1] While it must be noted that Model Predictive Control and optimization based control frameworks can be difficult to implement, such optimization programs can be solved quite efficiently in practice, even in the presents of nonlinear dynamics [16]. Since the work presented here is a pilot study for assessing interaction and understanding on the road, the inputs are computed offline for ease of implementation, and executed for the subjects in real-time.

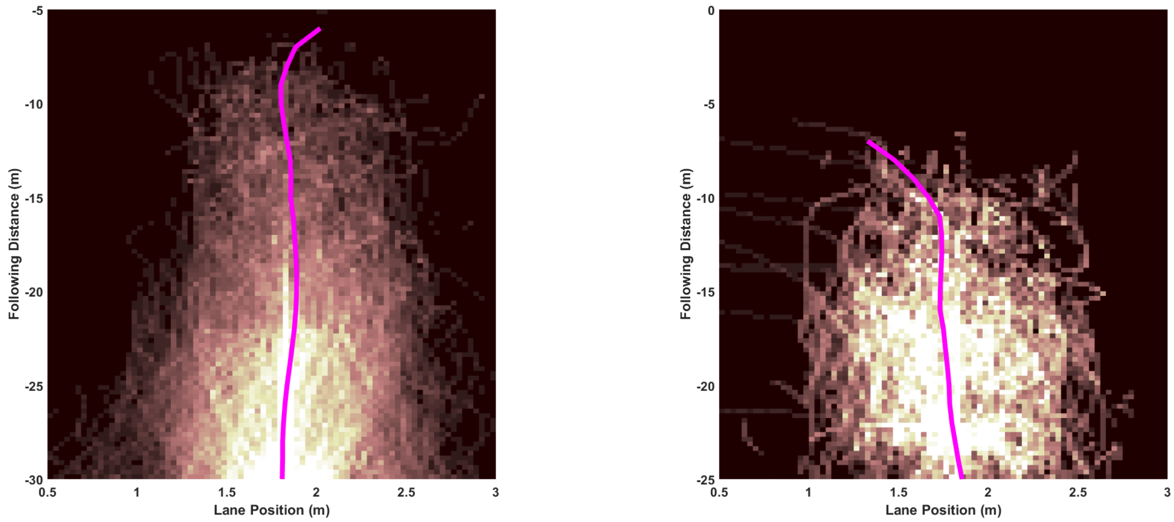

(a) Lane Keeping Mode

(b) Preparing to Lane Change Mode

Fig. 3: In these two figures, the cost maps for the two similar modes are visualized. The dark areas show locations with low probability, and light areas show locations with high probability. The pink line shows the mean position within the lane given the distance to the lead vehicle.

the framework presented. This allows us to mimic the driver behaviors and hopefully capture the subtle communication that occurs in this social scene.

IV. Validation Study

To show that by mimicking the driver data we are effectively communicating to the other drivers, a pilot validation study was completed as a proof of concept. The goal of this study was to verify two concepts: (1) that humans understand other human motions more effectively than generic autonomous motions and (2) that our control scheme using nominal trajectories is effectively capturing the non-verbal communication between drivers.

A. Comparing Control Schemes

In this study, we compared three nominal trajectory methods for identifying x_q in Equation 5:

1) *Controller design using standard methods*, where the vehicle uses generic trajectory templates for lane keeping and lane changing: $x_q = x_c$, where x_c denotes the lane center. The decision making still aims to mimic the human decision making process to match decision timing, but the trajectory the controller follows standard methods that aim to minimize deviation from the center of the lane.
2) *Controller design using the human inspired methods*, where the desired lane position is given by: $x_q = E[X_q|d]$, where we compute the expected lane position of the data associated with the current mode, X_q, given the distance to the lead vehicle, d. Not only does the decision making process mimic the human, but the trajectories are derived by using the templates found using the cost maps shown in Figure 3.
3) *Human controlled baseline*, where the inputs from a human driver are used to act as a baseline for understanding, meaning that the optimization control framework is not used. The input from a human driver is replayed, so the subject experiences the control scheme as if it where autonomous.

In this study, these control schemes were implemented in a scenario where the autonomous vehicle merges in front of a vehicle in the next lane, with the presence of a lead vehicle. In addition, these control schemes are examined from two different perspectives: (1) when the driver is experiencing the autonomy as a passenger in the driver's seat and (2) as another vehicle on the road. The scenario and the two viewpoints are visualized in Figure 4.

Nine subjects were asked to experience these autonomous (or seemingly autonomous) control schemes in a random order, riding in a motion platform vehicle simulator (Fig. 4). This simulator aims to recreate a safe, realistic environment for conducting human subject studies when developing controllers, using PreScan simulation software [7].

To understand how effective the communication of intent was for each method, the subject was asked to indicate in real-time when they thought the autonomous vehicle was about to change lanes, similar to when they believed the vehicle might turn its turning signal on. We note that no blinkers were used to verify that the subjects could predict the lane change using just motion cues. In addition, feedback from the subjects were obtained through a survey, targeting the understandability of the autonomous vehicle as well as user experience during the interaction.

For the three tested control schemes, the subjects were asked to experience the autonomous vehicle from view 1 and

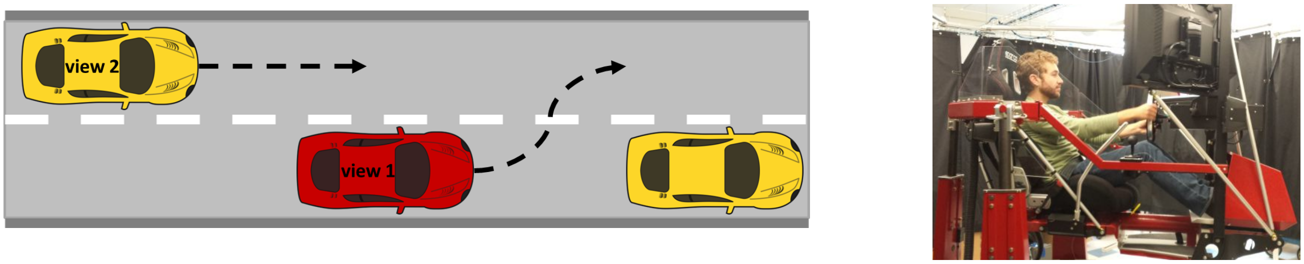

Fig. 4: *(Left)* The validation study asked the driver to indicate when the autonomous vehicle (shown in red) was about to change lanes from view 1, as a passenger in the driver seat, and from view 2, as another vehicle on the road. *(Right)* Motion simulator originally used to collect the dataset, and used to validate the different control schemes.

press a button when they believed the autonomous vehicle was about to change lanes. Then, the subject answered survey questions to obtain feedback on the three different control schemes. This was then repeated from view 2 to obtain feedback from a different perspective.

B. Predictability Results

To gauge how predictable the autonomous vehicle was, we compare the time that the subject indicated when the autonomous vehicle was about to change lanes relative to the time the vehicle exits the lane, crossing over into the next lane. For each method, this is defined as:

$$t_P = t_{\text{LaneExit}} - t_{\text{Human}} \tag{6}$$

where t_P is the prediction time (i.e. the time horizon prior to the lane change), t_{LaneExit} is the time at which the autonomous vehicle exits the lane, and t_{Human} is the time indicated by the human to let us know she believes the autonomous vehicle is about to change lanes. Ideally, the subject will be able to predict that the vehicle is intending on changing lanes well before it actually happens.

The average timing responses are provided in Table I and a visualization of the improvement in prediction time compared to standard control methods is shown in Figure 5. As shown, the prediction time is increased to more than a one second time horizon, which is significant given the limitations of human reaction time. By giving the driver extra time to react, smoother responses and improved acceptance can be expected.

From this study of predictability, the following observations can be made about how humans communicate on the road and the different control schemes:

1) *Effect of Lane Position vs. Heading Angle:* A common technique for predicting lane departures is to look at the distance to the lane marker and heading angle and compute the time to lane exit based on the current speed [14]. This means that if the heading of the vehicle is pointing toward the next lane, a time prediction can be calculated for the lane change. To see if the human prediction was similar to this model based method, we counted the instances when subjects indicated the lane change versus instances when there was a lane change predicted using this model based method within the next two seconds. We found that these two predictions were only in agreement approximately 40% of the time, indicating that humans are using cues other than heading angle to predict lane changes.
2) *Standard Methods:* From both viewpoints, the user is generally able to predict when the lane change is about to occur with approximately one second prediction time. It was also noted that this was prediction time was more consistent between subjects than the other methods, in terms of the variance. It was noted by subjects that timing of this lane change was highly predictable, due to the fact that the decision making and timing came across as human-like.
3) *Human-Inspired Methods:* As we can see, this shows a significant improvement of predictability from subjects within the vehicle. We see approximately a 40% increase in the time prediction, implying that the lane positioning is a key component of communicating on the road.
4) *Human Controlled Baseline:* This method provides the best predictability, verifying the hypothesis that humans understand their behaviors even without the traditional visual cues (like turning signals). This also validates the claim that humans communicate through

TABLE I: Average prediction time for each method in seconds.

Method	Standard	Human-Inspired	Baseline
View 1	0.958	1.462	2.307
View 2	1.110	1.452	2.102

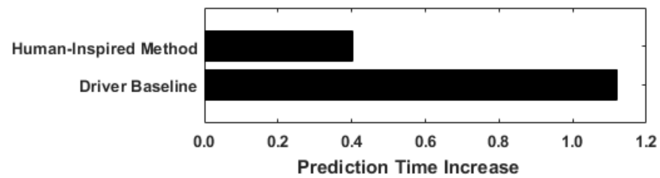

Fig. 5: Visualization of increased time prediction performance compared to the standard methodology, calculated as $(T_P - T_S)/T_S$, where T_P is the expected prediction time with the associated each method and T_S is the expected prediction time associated with the standard control design method, across both view points.

motion while driving to convey their intent, which is well understood by other drivers. We note that although this scheme gave a high predictability measure, the majority of subjects preferred the other control schemes, describing this its behavior as “erratic”.

C. Subject Feedback

Interesting feedback and comments were obtained through survey and comments that shed light on the user experience. When riding in the autonomous vehicle, just over half of the subjects preferred the Standard Method, stating that it felt smoother than the other methods. About half of these subjects indicated that it was also more predictable than other methods. This is somewhat counter-intuitive, as we can see in Figure 5, that this is not necessarily the case given the subjects’ prediction time.

Majority subjects also commented that the Standard Control Method executed a smooth and safe lane change. Meanwhile, comments and feedback on the Human-Inspired Method revealed that it came across as being less aggressive, which may have impacted the subjects perception on the automation’s competence. For the human controlled vehicle, the drivers stated that the controller seemed more erratic than other methods and indicated that this was the least trusted control scheme.

V. Discussion

In conclusion, we present the findings of a pilot study on human-inspired control schemes that could safely communicate through motion to surrounding drivers. This was completed by using human-inspired nominal trajectories for different driver modes that have been identified using realistic driver data from multiple drivers. The following ideas were confirmed: (1) humans communicate through motion while driving and (2) the presented control scheme was able to capture this and convey its intent to surrounding drivers.

Since this study was a proof of concept, there is a great deal of future work to be completed. More advanced methods for identifying the nominal trajectories and for controlling the vehicle must be explored to improve the feel of the autonomous system, and expand the framework to include a wider variety of scenarios and driver modes.

Another extension of this work would be generalizing and expanding this formulation to other spaces, beyond autonomous vehicles. This is crucially important as more and more robots are entering human dominated spaces, each with their own social cues and interactions. Questions arise as to how to implement the methods considered here in a more widely applicable form, particularly in areas and fields where there does not exist detailed datasets, and are left for future studies.

References

[1] P. Abbeel and A. Y. Ng. Apprenticeship learning via inverse reinforcement learning. *In Proceedings of ICML*, 2004.

[2] J. Bebel, N. Howard, and T. Patel. An autonomous system used in the darpa grand challenge. In *IEEE Intelligent Transportation Systems*, pages 487–490, 2004.

[3] C. Becchio, V. Manera, L. Sartori, A. Cavallo, and U. Castiello. Grasping intentions: from thought experiments to empirical evidence. *Fronters in Neuroscience*, 6(117), 2012.

[4] F. Borrelli, M. Baotic, A. Bemporad, and M. Morari. Efficient online computation of constrained optimal control. In *Proceedings of the 40th IEEE Conference on Decision and Control*, volume 2, pages 1187–1192 vol.2, 2001.

[5] A. Doshi, B. T. Morris, and M. M. Trivedi. On-Road Prediction of Driver’s Intent with Multimodal Sensory Cues. *IEEE Pervasive Computing*, 10(3):22 – 34, Sept. 2011.

[6] A. Dragan, K. Lee, and S. Srinivasa. Legibility and predictability of robot motion. In *Human-Robot Interaction*, March 2013.

[7] K. Driggs-Campbel. Experimental design for human-in-the-loop driving simulations. Technical Report UCB/EECS-2015-59, EECS Department, University of California, Berkeley, May 2015.

[8] K. Driggs-Campbell and R. Bajcsy. Identifying modes of intent from driver behaviors in dynamic environments. In *Intelligent Transportation Systems (ITSC), 2015 IEEE 18th International Conference on*, pages 739–744, Sept 2015.

[9] K. Driggs-Campbell, V. Shia, and R. Bajcsy. Improved driver modeling for human-in-the-loop control. In *2015 IEEE International Conference on Robotics and Automation*, May 2015.

[10] T. Kondoh, T. Yamamura, S. Kitazaki, N. Kuge, and E. R. Boer. Identification of visual cues and quantification of drivers’ perception of proximity risk to the lead vehicle in car-following situations. *Journal of Mechanical Systems for Transportation and Logistics*, 1(2):170–180, 2008.

[11] C.-F. Lin and A. G. Ulsoy. Calculation of the time to lane crossing and analysis of its frequency distribution. In *American Control Conference, Proceedings of the 1995*, volume 5, pages 3571–3575 vol.5, Jun 1995.

[12] T. Litman. Autonomous Vehicle Implementation Predictions: Implications for Transport Planning. `http://www.vtpi.org/avip.pdf`, 2015. Accessed: 2015-03-19.

[13] C. Little. The intelligent vehicle initiative: Advancing ”human-centered” smart vehicles.

[14] S. Mammar, S. Glaser, and M. Netto. Time to line crossing for lane departure avoidance: A theoretical study and an experimental setting. *IEEE Transactions on Intelligent Transportation Systems*, 7(2):226–241, 2006.

[15] A. Y. Ng and S. Russell. Algorithms for inverse reinforcement learning. In *in Proc. 17th International Conf. on Machine Learning*, pages 663–670. Morgan Kaufmann, 2000.

[16] V. Shia, Y. Gao, R. Vasudevan, K. Campbell, T. Lin, F. Borrelli, and R. Bajcsy. Semiautonomous vehicular control using driver modeling. *IEEE Transactions on Intelligent Transportation Systems*, PP(99):1–14, 2014.

[17] T. Simonite. Data Shows Googles Robot Cars Are Smoother, Safer Drivers Than You or I. `http://www.technologyreview.com/news/520746/data-shows-googles-robot-cars-are-smoother-safer-drivers-than-you-or-i/`, 2013. Accessed: 2015-03-19.

[18] C. Urmson, J. Anhalt, D. Bagnell, C. Baker, R. Bittner, M. N. Clark, J. Dolan, D. Duggins, T. Galatali, C. Geyer, M. Gittleman, S. Harbaugh, M. Hebert, T. M. Howard, S. Kolski, A. Kelly, M. Likhachev, M. McNaughton, N. Miller, K. Peterson, B. Pilnick, R. Rajkumar, P. Rybski, B. Salesky, Y.-W. Seo, S. Singh, J. Snider, A. Stentz, W. R. Whittaker, Z. Wolkowicki, J. Ziglar, H. Bae, T. Brown, D. Demitrish, B. Litkouhi, J. Nickolaou, V. Sadekar, W. Zhang, J. Struble, M. Taylor, M. Darms, and D. Ferguson. Autonomous driving in urban environments: Boss and the urban challenge. *Journal of Field Robotics*, 25(8):425–466, 2008.

[19] E. Wahlstrom, O. Masoud, and N. Papanikolopoulos. Vision-Based Methods for Driver Monitoring. In *Intelligent Transportation Systems*, volume 2, pages 903 – 908. IEEE, Oct. 2003.

[20] J. Yen, X. Fan, S. Sun, T. Hanratty, and J. Dumer. Agents with shared mental models for enhancing team decision-makings. In *DECISION SUPPORT SYSTEMS, SPECIAL ISSUE ON INTELLIGENCE AND SECURITY INFORMATICS*, 2005.