

Identifying Modes of Intent from Driver Behaviors in Dynamic Environments

Katherine Driggs-Campbell and Ruzena Bajcsy

Abstract—In light of growing attention of intelligent vehicle systems, we propose developing a driver model that uses a hybrid system formulation to capture the intent of the driver. This model hopes to capture human driving behavior in a way that can be utilized by semi- and fully autonomous systems in heterogeneous environments. We consider a discrete set of high level goals or intent modes, that is designed to encompass the decision making process of the human. A driver model is derived using a dataset of lane changes collected in a realistic driving simulator, in which the driver actively labels data to give us insight into her intent. By building the labeled dataset, we are able to utilize classification tools to build the driver model using features of based on her perception of the environment, and achieve high accuracy in identifying driver intent. Multiple algorithms are presented and compared on the dataset, and a comparison of the varying behaviors between drivers is drawn. Using this modeling methodology, we present a model that can be used to assess driver behaviors and to develop human-inspired safety metrics that can be utilized in intelligent vehicular systems.

I. INTRODUCTION

In recent years, there has been a growing interest in intelligent vehicle systems, including advanced driver assistance, driver monitoring, and autonomous systems [9], [12]. Many of these systems require a model of the driver, which has been shown to improve control algorithms and can be used to avoid safety heuristics [16]. These human centered systems are going to be increasingly important as more and more autonomy is introduced onto the roads [13].

In this paper, we consider a driver model that assesses driver intent to analyze what influences human decision making while driving. By developing this model, we hope to achieve two goals: (1) accurately capture a driver’s intent in dynamic environments (i.e. the road properties and surrounding vehicles) and (2) design a system that is flexible and portable enough to be used in a variety of applications (e.g. driver feedback or autonomous decision making). While humans are prone to distractions while driving [15], we have many desirable qualities like flexibility, adaptability, and efficient high level decision making.

This material is based upon work supported by the National Science Foundation Award ECCS-1239323 and 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).

Ideally, an autonomous system would be able to mimic these positive aspects of drivers while mitigating the drawbacks, and in doing so would improve social understanding and acceptance. This is motivated by work in [17], as the authors demonstrated that sharing a mental model improves human-robot collaborations. In [5], the authors showed that intent is directly tied to human motion, which others can detect and use to anticipate and understand intent; thus implying that understanding intent is “deeply rooted” to social interaction.

We propose a hybrid system inspired model that identifies the intent of the driver, which can be used to assess driver behavior and/or assist in autonomous decision making in a heterogeneous environment (i.e. an environment with a mixture of human driven and autonomous vehicles). In order to ensure portability of the model to systems with mixed controllers, the detection relies on the sensor measurements of the environment surrounding the ego vehicle and not directly on the control inputs or driver state. By using a state of the art simulator, a dataset is collected with a real-time labeling system to gather sensor information from the environment and the humans intent as she drives. Using this dataset, we are able to analyze the influences on decision making, assess a driver behaviors in different modes of intent, and develop an algorithm that can accurately determine the intent of a driver based on the measurements of a dynamic environment to fit a hybrid model formulation.

This work is unique from other modeling methods for the following reasons: (1) a real-time, human labeled dataset of lane changes is collected, that is better able to capture the human decision making process and (2) the switching of modes are determined by the dynamics of the environment to better encompass the decision making process than previous methods, that identify modes by heuristics and time.

The paper is organized as follows. In the subsequent section, we briefly discuss related works. In Section II, we present our modeling methodology. Section III describes the experimental setup and Section IV presents the results and analysis of the data and the corresponding models. Finally, the results are discussed and summarized in Section V.

A. Related Works

There are many works that consider predicting driver behavior by monitoring the driver [9], and have shown promising results in terms of human-in-the-loop and shared control for semi-autonomous frameworks [16]. When it comes to driver modeling intent, many rely heavily on the driver state [7] or on driver input [11]. While effective, these

methods rely on heuristics to determine when a lane change begins and use windowing to train the classifier [7]. This ultimately assumes that these high level decisions are made as a function of time, and not by the dynamic state of the environment, which is difficult to predict due to the high variability over long time horizons [16].

One of the desired outcomes of this work is to identify a model that does not depend on human input and can be used in human driven or autonomous systems. Ideally, the model could be used in human-inspired driving applications. Driving styles in terms of discrete control actions were mimicked in [3], using inverse reinforcement learning.

By requiring that the human driver explicitly label the current mode being driven, we can use supervised classification approaches to generate system identification parameters for these modes of intent. The resulting model based on observations of the surrounding vehicles that can be detected using current sensor technology, with features similar to the cues that humans perceive while making driving decisions. As previously mentioned, this allows the algorithm to be used as a decision system for an autonomous system or an assistance system for a human driven or semi-autonomous vehicle that can effectively function in a mixed environment.

II. METHODS

In this work, we model each mode of intent as a discrete state, which has a different controller (or control objective) associated with that mode. As the driver navigates through the environment, she transitions between the modes of intent, switching controllers in each mode. We assume that the switching of modes is only influenced by the state of the surrounding environment (i.e. the driver will not spontaneously change lanes when there are no other cars nearby).

A. Hybrid System Approach

A hybrid system is a representation of a dynamical system that has continuous dynamics that depend upon a discrete state or mode of operation. Suppose we are given the vehicle dynamics for the ego vehicle, which can be approximated using a bicycle tire model extended with roll dynamics, as found in [14]. We compactly represent these dynamics in the following discretized form:

$$x_{k+1} = f(x_k, u_k), \quad \forall k \in \mathbb{N} \quad (1)$$

where $x_k \in \mathbb{R}^n$ is the state of the vehicle and $u_k \in U$ is the input from the input space at time k . The following states are those of interest:

$$x_k = [p_x \quad p_y \quad v_x \quad v_y \quad \theta]^\top \quad (2)$$

where p_x and p_y are the positions with corresponding velocities v_x and v_y , and θ is the heading angle of the vehicle. We can decompose the dynamics into the following form:

$$x_{k+1} = f(x_k) + g_q(x_k, u_k), \quad \forall k \in \mathbb{N} \quad (3)$$

where all variables are as before and $g_q(x_k, u_k)$ is the control law that defines the input to the vehicle that changes with the associated mode $q \in Q$. If human driven, this assumes that

the behaviors or inputs of the driver will change dynamically based on intent. If semi- or fully autonomous, this algorithm might update the high level objective, constraints, or cost function for a control algorithm like model predictive control or would identify a new target set for control schemes using motion planning.

We suppose that the vehicle will be acting in a dynamic environment that will determine the mode of the ego car. In our scenario, the set of modes Q describe the various “types” of driving that occur during normal driving. These might include lane changing, emergency avoidance maneuvers, lane keeping, etc.

However, as previously stated, we assume that the mode is determined not by the dynamics themselves, but by the observable, relative dynamics of the surrounding vehicles, which can be obtained via sensor measurements that are updated at each time step. Suppose the sensors map the true dynamics of the environment to a noisy estimate through some function h :

$$\hat{y}_k^i = h(x_k^i) \quad \forall i = \{1, \dots, m\} \quad (4)$$

where x_k^i is the true state (as in Eq. 2) and \hat{y}_k^i is the measured representation of vehicle i , supposing that we can detect and measure m vehicles within a predefined radius of the ego vehicle. Here, we suppose \hat{y}_k^i for $i = 1, \dots, m$ can be passed to a detection algorithm, \mathcal{A} , to calculate $\sigma_q \in \Sigma$, which is the set of discrete inputs that detects and signals changes in the environment that instigate a change of modes. Formally,

$$\sigma_q = \mathcal{A}(x_k, \hat{y}_k^1, \dots, \hat{y}_k^m) \quad (5)$$

where all variables are as previously described. Note that \mathcal{A} will compute and select features from the vehicle states, and will exclude control inputs to maintain flexibility.

B. Modes of Intent

In this paper, we focus on identifying the transitions between driver intent modes that a driver encounters while driving. We predefine the modes of behavior, collect the driver data to build a dataset, and learn the hybrid system model inspired by human driving.

For simplicity, we examine the scenario of driving in a two-lane, one way road, in a non-urban setting. We define three modes: *lane keeping*, *preparing to lane change*, and *lane changing*. The model is visualized in Figure 1. While *lane keeping* and *lane changing* are self-explanatory, *preparing to lane change* can be thought of as the mode when the driver begins planning to change lanes and is waiting for the proper moment. This can be thought of as the moment when the driver turns on her turning signal or blinker to indicate that she wants to change lanes soon. By identifying this mode, we inherently have a predictive nature in the system, by determining the lane change prior to the maneuver actually occurring. But instead of putting a distinct time stamp on the prediction horizon, we suppose that this relies on new instances on the environment. In Section IV, we show the empirical results for the prediction time horizon.

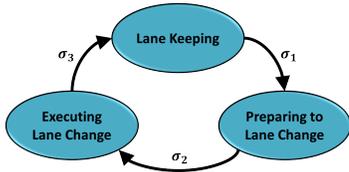


Fig. 1: Illustration of discrete modes in our hybrid model of driver intent, where we model the transitions as discrete inputs, σ_q .

We introduce this in-between mode as a precursor to executing the lane change for two reasons. First, detecting lane changes from a dataset has traditionally been done by determining when a lane change occurs by some heuristic (e.g. when the heading angle passes a particular threshold or when the vehicle exits the lane). These models look at the data leading up to this point in order to predict that a lane change will occur in the next few seconds [7], [11]. This, however, does not capture the decision making process of the human, or capture the idea that these decisions occur as a function of the environment, not just time.

Second, we introduce this in-between state in hopes to capture the variation that humans exhibit as their intent change. As humans drive, we are constantly assessing the environment, and building a mental model of what is happening around us by determining the intent of other vehicles. While we often rely on turning indicators or blinkers to convey our intent to surrounding vehicles, humans can often determine whether or not a vehicle wants to change lanes without the signal. If we assume that vehicles are to be integrated into society, we can assume that they will need to be predictable and convey their intent to human drivers beyond just using the indicator.

This is helpful in building our system, because (1) it mimics the decision making process more accurately by introducing an intermediary step between the two obvious actions and (2) it allows the model to vary the amount of prediction for when the driver would change lanes, by relying on the state of the environment dynamics, not a predetermined time horizon.

To learn these transitions or discrete inputs (Eq. 5), we build a dataset that consists of the environment, as represented in Equation 4, with corresponding mode labels. Details about the experiments and data collection are presented in Section III. Since we have collected a set of labeled data, we may use supervised classification techniques to learn and analyze the transitions between modes.

C. Features from the Dataset

As described, we were interested in analyzing the effects of the dynamic environment on driver’s intent and decisions. Therefore, we utilized the observable and measurable dynamics of the surrounding vehicles given current sensing technology as the starting point of the features. This means that from a dataset of x and \hat{y}^i we generate features with associated label $q \in Q$, which match the predefined modes. The distances to the surrounding vehicles are the primary raw features, as relying on heading angle and inputs would

eliminate the portability of the algorithm.

To make the features robust to changes in position, velocity, and control input, we considered relative positions and velocities, heading relative to the road, and as well as two time metrics: time-to-collision (TTC) and time headway (THW), which are commonly used as a metric for threat and driver perception [4], as well as the relative TTC and THW, denoted rTTC and rTHW. These are defined as follows:

$$\text{TTC}_i = \frac{d_i}{v_e}, \quad \text{THW}_i = \frac{v_e}{d_i} \quad (6)$$

$$\text{rTTC}_i = \frac{d_i}{v_i}, \quad \text{rTHW}_i = \frac{v_i}{d_i} \quad (7)$$

where v_e is the speed of the ego vehicle and d_i and v_i are the relative position and velocity, respectively, to vehicle i .

We define a detection region of $50m$ for which features of detected vehicles are included. For consistency, the feature vector is ordered such that we consider a small grid around the ego vehicle. The first position corresponds to the vehicle in front of the ego vehicle and in the same lane. The second position is the in front of the mid-way point of the ego vehicle and in the opposite lane. The third and final position is in the opposite lane as the ego vehicle, behind the mid-way point. This grid is visualized in Fig. 2a, and is mirrored when the ego vehicle is in the left lane. If no vehicle is in range in one of the three defined position, padding is inserted to maintain ordering¹. Using the normalized feature vector and corresponding labels, we use classification techniques to identify the driver model.

D. Identifying Transitions

We translate this problem into a classification problem, assuming we have a complete, labeled data set of driver behaviors. Formally, we wish to define our algorithm \mathcal{A} as a function that will take the measurements from the environment, calculate our features, and identify the current driver mode using a classification algorithm. We present results using the following techniques: Support Vector Machines (SVM), Random Forests (RF), and Logistic Regression (LR).

SVM and RF were tested due to their efficiency and robustness to overfitting. These methods are also logical choices when finding regions associated with classes or when identifying separating hyperplanes. LR was tested as it is the most popular of soft classification algorithms, meaning it assigns a score or a probability that a sample belongs to a class given the features. This should be able to capture the “gray” area that is present in human decision making by identifying the likelihood or mixture of the modes in a given sample. For further information, we guide the reader to [10].

¹Let it be noted that a variety of feature sets were tested in developing this model, including sets that were applied to different classifiers based on the number of surrounding vehicles. We’d also like to note that the modes were varied, using the three modes presented here but labeled for the left and right lane, creating a total of six modes. While this showed promising results, to keep the model simple, we limited the number of modes. Additionally, preliminary studies show promising results that this can easily be extended to roads with more lanes and of varying curvature. However, as the scenarios expand more sufficient data must be collect as well as more modes, so was not explored here given time constraints.

III. EXPERIMENTAL SETUP

In this section, we describe our experimental setup for collecting the dataset and the features used to discriminate modes. When studying human-in-the-loop systems, one of the challenges is collecting realistic data in a safe manner. To address this, we have developed an experimental setup for studying human-in-the-loop systems in vehicles, particularly in driving applications. The resulting testbed was designed to be a flexible, realistic platform that allows us to both observe the driver with monitoring devices, but also control and measure the environment [8].

Driver data was collected using a Force Dynamics CR401, a 4-axis motion platform simulator, which recreates the forces experienced while driving [1]. This system has been integrated with PreScan software, which provides vehicle dynamics and customizable driving environments allowing us to recreate various driving environments needed for data collection [2]. Using this human-in-the-loop test bed, we are able to reliably and realistically obtain driver data that can illustrate the utility of our models and provide useful motion feedback to the drivers. The motion simulator and visualization seen by the driver is shown in Fig. 2b.

To set up the experiment, multiple scenarios were created in which the driver traverses a straight two lane road attempting to maintain a speed between 15 and 20 m/s. Scenarios were generated by creating combinations of the simulation parameters to collected a complete dataset. The following parameters were varied: (1) the initial speed and lane location of ego vehicle; (2) the number and location of surrounding vehicles, varied from one to three; and (3) the initial and final speed of each surrounding vehicle.

For example, in some scenarios, the lead vehicle would slow down, forcing the driver to change lanes only if there was room in the next lane. Thus, the key here is finding the configurations of the environment states that cross the boundary or safety margin of the human and allows us to identify their likely action between staying in the lane (i.e. braking) or changing lanes to maintain her desired speed.

While driving, the human driver actively labels the data into classes or modes for the training and testing sets during the experiment. The driver presses a button on the steering wheel to signal that she was preparing to change lanes and then presses a paddle on the steering wheel to indicate that she was executing the lane change. This allowed us to separate the data into the modes as previously mentioned.

Five drivers were recruited for this experiment, as a pilot study and proof of concept. The drivers were asked to drive for an hour on two separate occasions to collect the data. An optional practice session was offered, to get the driver accustomed to the simulator and the data collection method. The dataset ultimately resulted in about 200 lane changes, per driver, which was divided into two datasets for training and testing. We note that some scenarios did not require a lane change (e.g. the relative speed of the lead vehicle was initialized such that the driver never felt the need to overtake them), while other scenarios which heavy traffic

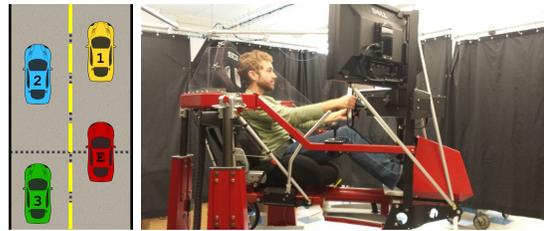


Fig. 2: (a) Illustration of the ordering of feature vector depending on location of vehicle relative to the ego vehicle (denoted E). (b) Picture of the driver in the motion platform simulator.

caused multiple lane changes, but varied depending on the driver’s behaviors in the simulation. The vehicle dynamics, driver mode, and environment is sampled at the synchronized rate of 60 Hz.

IV. RESULTS

In this section, we examine the influence of the features on driver behavior and describe the models developed from the aforementioned algorithms.

A. Analyzing Driver Behaviors

From this dataset that describes driver intent, we are able to visualize the variability in human driving and gain insight to the connection between a driver’s perception of her surrounds to her discrete state of intent. This data can also verify the utility of introducing the *preparing to lane change* mode; the advantage of the human active labeling method over the traditional heuristic labels and windowing; and compare the differences in behaviors between drivers.

To justify separating the lane keeping modes into two modes, we consider the distribution of lateral deviations from the center of a lane, which is assumed to be the driver’s internal goal in this mode. We observe that drivers tend to distribute themselves closer to the center lane and are more likely to edge away from traffic when simply lane keeping, and move toward next lane when preparing to lane change. The empirical distributions of these behaviors are shown in Figure 3 ².

To analyze the behavior in modes and across drivers, we also consider the time spent in the *prepare* mode, T_P , as well as the time difference, ΔT , between when the instance that human identified a mode transition and the instance that the lane change occurred according to a baseline heuristic, which was determined as the time the bounding box of the vehicle exited the lane. This evaluation is presented in Figure 4 for a single driver.

By visualizing these time metrics, we observe that the time between modes varies a great deal, implying that the decision making process does not solely depend on time, but on the dynamics of the environment. This also empirically shows the predictive capability of this model, showing that

²The subsets of the data were compared using the two-sample Kolmogorov-Smirnov test to see if they came from the same distribution. The hypothesis was rejected with $p \ll 0.01$.

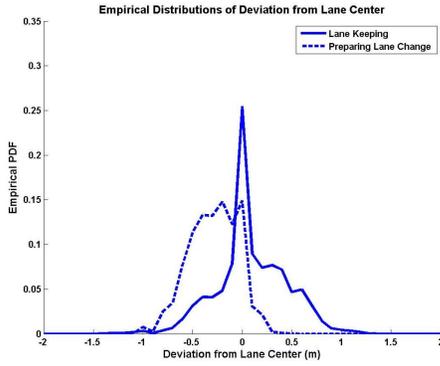


Fig. 3: Empirical probability density functions that describe the different driving behaviors exhibited in the similar, yet distinct, lane keeping and preparing to lane change modes, for a particular driver. This data consists of data from the right lane, but we noted similar, but reflected behaviors, in the other lane.

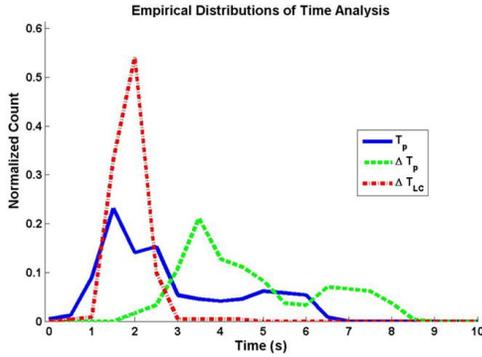


Fig. 4: Empirical probability distributions showing our time analysis metrics, where T_P is the time spend preparing to lane change and ΔT is the change in time from when the driver entered the specified mode compared to the baseline. The mean and standard deviations for all drivers are shown in Table I.

TABLE I: Results from the timing analysis and the TTC and THW metrics are presented, showing the mean and standard deviation of the means for all subjects. Subscript P denotes result for *prepare* mode and LC denotes *lane change* mode.

Metric	Mean (s)	St. Dev. (s)
T_P	3.05	0.29
ΔT_P	2.66	1.17
ΔT_{LC}	1.52	0.28
TTC_P	1.34	0.17
TTC_{LC}	1.20	0.14
THW_P	0.80	0.11
THW_{LC}	0.96	0.12

we are able to predict the intent of the vehicle, prior to the maneuver actually occurring. Similarly, the TTC and THW metrics can also be analyzed to identify the typical thresholds for the mode transitions as labeled by a driver. The averages for the timing metrics and these features for all subjects are presented in Table I.

B. Model Performance

It is worth noting that we cannot expect a perfect model, as there is uncertainty near the boundaries of the human's decision making process and noise in the actual labeling of data. Thus, errors will likely occur at the boundaries,

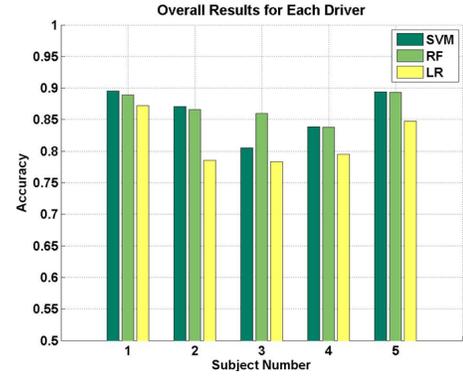


Fig. 5: Plot showing the accuracy for each of the five subjects.

TABLE II: This table presents the overall model accuracy (in %) and the accuracy in different scenarios, denoted Scen. i , where i indicates the number of vehicles in a 50m radius of the ego vehicle, for a particular driver only on the test dataset.

Method	Overall	Scen. 1	Scen. 2	Scen. 3
SVM	89.5	88.5	85.9	92.1
RF	88.9	86.6	85.9	91.7
LR	87.2	86.1	85.4	88.7

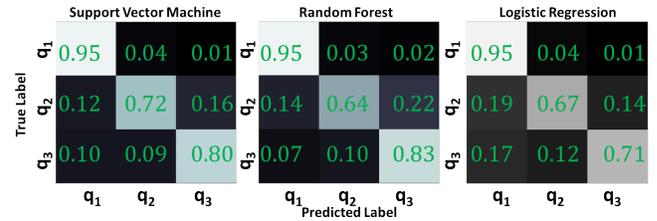


Fig. 6: Confusion matrices showing normalized accuracy results, on the test dataset for a given driver. Entry i, j shows the ratio of samples that are from mode q_i , but predicted by the algorithm to be q_j . Accuracy is labeled and denoted by the color, where white indicates 100% and black indicates 0%.

but we expect the model to identify the optimal separating hyperplane to balance the noisy data.

By using the obtained parameters from cross validation³ and applying the generated classifiers to testing data to validate the model, we obtain the performance results presented Figure 5. The results for a sample driver are shown in Table II. The accuracy is presented as the overall accuracy, and by scenarios that are defined by the number of vehicles in the immediate vicinity of the ego vehicle.

We can see that the generated classifiers determines the driver mode with relatively high accuracy, despite the uncertainty that is inherent in human actions. To visualize the classification accuracy without biases in the number of instances, the confusion matrices showing the normalized results are shown below in Figure 6.

As shown, this methodology shows promising results for identifying a hybrid model of an individual driver's decision

³This was completed using the leave- k -out method. SVM was implemented using LibSVM [6], using the radial basis function kernel (the parameters for which were found in cross validation), and one-against-one method for multi-class classification. In RF, the number of trees was swept in cross validation. LR relied on the logit link function to identify the nominal model. Details for these methods can be found in [10].

making process. The following observations were made:

1) *SVM*: We note that it is generally conservative, transitioning from the lane keeping modes later than the other methods, but is still able to predict lane changes prior before the maneuver occurs. This method also provides smoother transitions than the other methods. This is to be expected given the method SVM balances the separating hyperplanes with the noisy labels.

2) *RF*: This classifier can accurately distinguish the behavior modes, but is not always smooth about the borders, which is intuitive due to fine parsing the algorithm exhibits when identifying the separation boundaries. While in theory this method is robust to overfitting, we noted that this method performed best in instances most similar to the training set. This implies that this method is more reliable under distinct operating conditions and does not generalize as well as SVM.

3) *LR*: The accuracy of this method was determined by using the most likely mode as the classification. While it can be seen that this method produced relatively low accuracy, it is interesting as it shows the modal mixtures near the transitions between modes as shown in Figure 7.

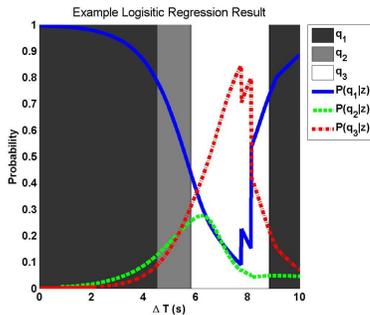


Fig. 7: Example illustration of how the logistic regression changes over the execution of a lane change, where the shaded regions show the actual mode, and the lines show the probability of a mode conditioned on the input features at that timestep.

4) *Overall Performance*: We note that the accuracy of the classification varies a great deal between subjects, which is likely due to the inaccuracies of human labeling and the nature of human decisions. When the data is combined in an attempt to generate a global model for all drivers, extremely poor performance was found. This is likely due to the fact that the internal decision making process is highly specific to the individual—especially when considering decisions like preparing to take an action, which if removed might produce better results. However, bounds on behaviors can still be identified for general use. If this method were to be used in practice, an expert driver would be selected (i.e. one that generates a reliable model). This expert could be used as a standard to compare other drivers’ behaviors to or to generate a high level control algorithm for use in semi- or fully autonomous vehicles.

Additionally, no windowing was used, although it is common practice in similar studies (see Section I-A). While this might improve accuracy, it was decided that using minimal features was desirable, to maintain flexibility and eliminate reliance on the trajectories generated on the simulator.

V. DISCUSSION

We have compiled a dataset of human labeled lane changes and developed a hybrid mental model to mimic a driver’s decision making process, by learning a detection module to identify the separating hyperplane between driver modes. This model is unique as it provides a robust classification of human decisions assuming that the transitions are not determined by time, but by the dynamic environment. The models have been shown to be highly robust to variations in the environment and exhibited high accuracy. Future work consists of expanding the dataset to consider more scenarios and more complex and uncertain environments. Also, there is a great deal of consideration that must go into demonstrating the feasibility of the model, as it was constructed in a simulated environment. Thus, verifying the methodology on real-world data must be explored. We would also like verify the generalizability and portability to other vehicles and work towards implementing this work as a high-level controller for autonomous systems in heterogeneous environments.

REFERENCES

- [1] Force Dynamics. <http://www.force-dynamics.com>, 2014. Accessed: 2014-09-30.
- [2] PreScan Simulation Software. <https://www.tassinternational.com/prescan>, 2014. Accessed: 2014-09-30.
- [3] P. Abbeel and A. Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of ICML*, 2004.
- [4] D. Abbink, E. Boer, and M. Mulder. Motivation for continuous haptic gas pedal feedback to support car following. In *Intelligent Vehicles Symposium, 2008 IEEE*, pages 283–290, June 2008.
- [5] C. Becchio, V. Manera, L. Sartori, A. Cavallo, and U. Castiello. Grasping intentions: from thought experiments to empirical evidence. *Frontiers in Neuroscience*, 6(117), 2012.
- [6] C.-C. Chang and C.-J. Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [7] 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.
- [8] K. Driggs-Campbell and R. Bajcsy. Experimental design for human-in-the-loop driving simulations. Master’s thesis, EECS Department, University of California, Berkeley, May 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. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning (2nd edition)*. Springer-Verlag, 2009.
- [11] N. Kuge, T. Yamamura, and O. Shimoyama. A driver behavior recognition method based on a driver model framework. *SAE Technical Paper 2000-01-0349*, 2000.
- [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] R. Rajamani. *Vehicle Dynamics and Control*. Springer-Verlag, 1st edition, 2005.
- [15] M. Regan, J. D. Lee, and K. L. Young. *Driver Distraction: Theory, Effects, and Mitigation*. CRC, 2008.
- [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] 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.