Interactive Planning for Autonomous Urban Driving in Adversarial Scenarios

Yuanfu Luo*[1,5], Malika Meghjani*[2,3], Qi Heng Ho*[3], David Hsu[1], and Daniela Rus[4]

***Abstract*— Autonomous urban driving among human-driven cars requires a holistic understanding of road rules, driver intents and driving styles. This is challenging as a short-term, single instance, driver intent of lane change may not correspond to their driving styles for a longer duration. This paper presents an interactive behavior planner which accounts for road context, short-term driver intent, and long-term driving style to infer beliefs over the latent states of surrounding vehicles. We use a specialized Partially Observable Markov Decision Process to provide risk-averse decisions. Specifically, we consider adversarial driving scenarios caused by irrational drivers to validate the robustness of our proposed interactive behavior planner in simulation as well as on a full-size self-driving car. Our experimental results show that our algorithm enables safer and more travel time-efficient autonomous driving compared to baselines even in adversarial scenarios.**

I. INTRODUCTION

Autonomous urban driving has been gaining popularity in recent years. A majority of research in this domain has an underlying assumption of *exo-vehicles* (referred to as nearby vehicles) being driven by reasonably rational drivers. In this paper, we specifically focus on adversarial scenarios for autonomous urban driving which are caused by irrational drivers. However, our proposed approach is applicable to complex driving scenarios with a combination of rational and irrational drivers. We achieve this with a holistic analysis of short-term driver intents and their long-term driving styles with respect to the contextual road rules for time-efficient and safe driving. For example, consider the scenario that the ego-vehicle is driving on a street with two lanes and an exo-vehicle is driving in front. If the exo-vehicle intents to keep lane, the ego-vehicle can overtake by changing to the adjacent lane for efficiency. However, considering the intent alone sometimes leads the ego-vehicle to perform naive actions resulting in unforeseen outcomes. In addition to the driver intent, the ego-vehicle also needs to consider the driving style of the exo-vehicle. If the exo-vehicle is constantly aggressively steering, the ego-vehicle needs to keep a safe distance and change lane only when there is safe space in the lateral direction. Specifically, the intents of the exo-vehicles help the ego-vehicle anticipate their future actions so that it can make time efficient decisions. Whereas, the driving styles help the ego-vehicle to stay at a safe distance from dangerous drivers.

The driver intents and driving styles of exo-vehicles are unfortunately, unknown to the ego-vehicle and need to be inferred. We infer the probability distributions over them for surrounding exo-vehicles using recurrent neural networks, based on their states and the road contextual information. The road context such as lane directions and distance to lane center can potentially help narrow down the space of possible driver intents and driving styles and in addition, speed up the planning process by filtering out invalid actions [1].

We propose, CID, *Context, Intent, and Driving style* aware planner which uses a Partially Observable Markov Decision Process (POMDP) framework to encode intents and driving styles as the hidden states, and simulate the possible future trajectories of neighboring exo-vehicles using the road context. We solve the POMDP using a risk-averse planner, IS-DESPOT [2], to get the optimal action for the ego-vehicle. Our contributions in this paper are three-folds: (a) an inference module for driving style analysis, (b) CID, an interactive behavior planner which integrates intent and driving style inferences with respect to road contextual information, (c) validation of our proposed CID planner using adversarial driving scenarios in simulation and on real-world data for qualitative analysis. We also validated our planner by comparing it to important baselines for quantitative analysis. Experimental results show that our planner can significantly outperform the baselines in terms of safety, especially for the adversarial scenarios, while achieving the same level of travel time efficiency.

II. RELATED WORK

A. Driver Intent and Driving Style Inference

Both inferences have been studied intensively but independently. In [3], speed and lateral position changes are used for driving style detection. The work in [4] evaluates the Euclidean norm on accelerometer data and uses the deviations from the average of the norm to classify the driving styles using fuzzy logic. DrivingSense [5] classifies dangerous driving style based on smartphone auto-calibration. The work in [6] uses linear regression to map a selected set of features to driving styles. These approaches abstractly classify driving styles into aggressive driving, normal driving, etc.. Our approach distinctively defines driving style into longitudinal and lateral erratic for risk-averse planning. Driver intent prediction approaches use HMM and SVM [7], [8], [9]

*authors contributed equally.

Yuanfu Luo and David Hsu are with the National University of Singapore[1], Singapore. dyhsu@comp.nus.edu.sg

Yuanfu Luo is also with SZ DJI Technology Co., Ltd., China[5]. logan.luo@dji.com

Malika Meghjani is with Singapore University of Technology and Design[2], Singapore. malika_meghjani@sutd.edu.sg

Malika Meghjani and Qi Heng Ho are with the Singapore-MIT Alliance for Research and Technology[3], Singapore. qi.ho@colorado.edu.

Daniela Rus is with the Massachusetts Institute of Technology[4], Cambridge, MA, USA. rus@csail.mit.edu

which work well for small data sets with few outliers. Other approaches for intent prediction are based on neural networks [10], [11], [12]. However, these require a huge amount of data for learning a range of intents. In contrast to the existing work, our recurrent neural network approach leverages the road context for driver intent and driving style inferences and combine their outcome to obtain informative observations (e.g. right lane change with aggressive driving style) for interactive planning. This allows us to generalize our approach even with a small noisy data set.

B. Planning for Autonomous Driving

The approaches for planning under uncertainty often model either driver intents or driving styles. For example, [13], [14], [15], [16], [17], [18] have modelled intents. MPDM [13] evaluates a fixed set of policies on the trajectories rolled out from the sampled intents. However, it only plans for one-time interaction with other agents. Intention-POMDP [14] modelled the autonomous driving problem as a POMDP to handle uncertain intentions. Luo et al. [15] improved on [14] by using an interactive motion model PORCA for the agents. However, both of them are designed for driving in pedestrian environments. The work in [18] has used a rich representation of road contexts to help predict vehicles' intent, but it also induces high computational complexity compared to our approach. Gao et al. [16] extract vehicle's intent from a planned path and feed the intent to a neural network to generate the vehicle's control. However, their approach only considers the intent of the ego-vehicle, ignoring those of exo-vehicles. Driving styles have also been studied previously for planning under uncertainty in autonomous driving applications. The work in [19] predicts vehicle's driving style: whether aggressive or patient. The predicted belief on driving style is modelled as hidden states in POMDP for autonomous driving. The work in [20] considers courteous driving style by placing a courtesy term in the cost function. However, the courtesy is only considered for the ego-vehicle. In contrast, our cost function considers driver intents and driving styles of exo-vehicles. Our previous work [1] used only driver intent as the hidden states to formulate the autonomous driving problem as a POMDP. Similar to other existing work, it assumes drivers are rational with no adversarial behaviors. This paper focuses on adversarial driving scenarios caused by irrational drivers. We consider multiple exo-vehicles and use road contextual information for both driver intent and driving styles inference to achieve time efficient and safe decision making.

III. Driver Intent and Driving Style Inference

We use road contextual information to infer driver intents and driving styles of the exo-vehicles. The driver intent refers to actions such as, lane-keep and lane-change. Whereas, the driving style represents driving patterns that classify a driver as either rational or irrational driver based on the frequency of lane changes and brakes in their driving patterns. Our goal is to analyze and infer the driver intents and driving styles of other road users to help in interactive, safe and time-efficient path planning for autonomous urban driving.

A. Driver Intent Inference

We implement a Long Short Term Memory (LSTM) recurrent neural network for inferring the short-term driver intent beliefs of neighboring exo-vehicles [21]. We formalized the intent as the lane-keep, left-lane change or right-lane change. The input to the LSTM network are based on the geometric features derived from the vehicle trajectory and road contextual information. The features are (a) δx: Change in lateral pose, (b) δy: Change in longitudinal pose, (c) ll: If left lane exists (Boolean $0/1$), (d) rl: If right lane exists (Boolean $0/1$) and (e) d_{center}: Distance to the center of current lane. d_{center} values range from $-1.5m$ to $1.5m$ as we consider average kerbside lane width to be $3m$ [22]. We train the LSTM network on sequences with 4 time steps of 0.25 seconds each. The training data for intent inference is obtained from NGSIM vehicle trajectory dataset [23] complemented with data collected by our ego-vehicle. The intent labels for our ego-vehicle data are marked by a safety driver using the indicator for referencing the start and the end of the intent.

B. Driving Style Inference

The driving style inference is a critical aspect of autonomous safe planning which helps us in distinguishing a rational exo-vehicle driver from an irrational one for risk-aware decision making. We define driving style according to the following categories: normal, longitudinally erratic, laterally erratic, and both (longitudinally and laterally) erratic. The decoupling of laterally and longitudinally erratic driving styles allows us to differentiate between the exo-vehicles that are aggressively steering and the ones that are recklessly accelerating and braking, respectively.

Consider a trajectory for an exo-vehicle which is defined as a temporal sequence of poses where a pose at time t is represented as $\mathbf{p}_t = (x_t, y_t, \theta_t)$. We divide the driving trajectory into time intervals of 1 second comprising of 4 time-steps with equal time. We train another LSTM with the following features as inputs to the network: (a) δx: Change in lateral position, (b) δy: Change in longitudinal position and (c) d_{center}: Distance to the center of current lane. The aforementioned features were specifically selected as they efficiently represent the changes in the driving styles with respect to the road contextual information. The output of the LSTM network is a discrete probability vector ϕ_t over four driving style classes ζ: i.e., normal, longitudinally erratic, laterally erratic, and both erratic which is provided every 0.25 seconds. We trained the LSTM network for driving style inference on the NGSIM dataset. We extracted trajectories in local lane coordinates from the dataset and clustered them into discrete groups by comparing their deviation to candidate reference trajectories derived from our transition model as described in IV-B. These clustered trajectories are then filtered and labelled by a human expert, before being divided into sub-sequences of 1 second. We validated our

proposed LSTM network by comparing it with a rule-based method and time series forest (TSF) [24], [25]. The rule-based approach considers the variance of δx and δy of the exo-vehicles and applies fixed thresholds to classify their driving styles. We use the same input features as LSTM for TSF, with default parameters except for the number of trees (500) and maximum tree depth (10). We report the weighted average precision, recall and F1-score for our test data in TABLE I. Our proposed method outperforms TSF and rule-based method, as neural network based methods are robust to noisy and unbalanced data.

Algorithm	Accuracy	Precision	Recall	F1 Score
Rule-based	0.64	0.62	0.64	0.63
TSF	0.75	0.74	0.75	0.72
LSTM network	**0.83**	**0.83**	**0.81**	**0.82**

TABLE I: Classification report comparison

In scenarios where we can observe a vehicle for a longer time period (longer than 1 second), we aggregate the past inferences. The advantage of aggregating the inferences as compared to classifying a longer series of driving data is that it allows for a shorter lead time for the belief inference. From the classifier output ϕ_t at each time-step t, we use a multinomial naive Bayes classifier in the following recursive form to maintain a posterior class belief distribution of an exo-vehicle's driving style β:

$$p(\zeta = \beta|\phi_{t,\beta}) \propto p(\phi_{t,\beta}|\zeta = \beta)p(\zeta = \beta|\phi_{t-1,\beta}) \quad (1)$$

where $\phi_{t,\beta}$ is the output probability of the LSTM network corresponding to the class β at time t. We also add a Laplacian smoothing factor in order to ensure $p(\zeta = \beta) > 0$ for each class β.

IV. Context, Intent and Driving Style Aware POMDP Planning

We formulate the behavior planning for autonomous driving as a POMDP [26] since it provides a principled way to model the uncertainty on intents and driving style for autonomous driving. We refer to our method as Context Intent and Driving style (CID) aware POMDP planner.

A. POMDP Preliminary

A POMDP models an agent acting in a partially observable stochastic environment. It is formally defined as a tuple (S, A, Z, T, O, R, b_0), where S, A, and Z are the state, action, and observation spaces respectively. $T = p(s'|a, s)$ is the probabilistic state transition from $s \in S$ to $s' \in S$ when the agent takes the action $a \in A$. It models the imperfect control of the agent and dynamically changing environment. $O = p(z|s, a)$ is the observation function defining the probability of observing $z \in Z$ when the agent takes action a and reaches state s. This models the sensor noise. $R(s, a)$ is the reward function defining the reward the agent can get by executing action a at state s. Since the agent does not know the exact state it is currently in due to partial observability, it maintains a belief, i.e., probability distribution, over its current state, and b_0 is its initial belief. POMDP planning aims to find a *policy* π, a mapping from a belief b to an action a, that maximizes the expected total discounted rewards:

$$V_\pi(b) = \mathbb{E}\left(\sum_{t=0}^{\infty} \gamma^t R(s_t, \pi(b_t)) \;\middle|\; b_0 = b\right), \quad (2)$$

where t is the time step, and $\gamma \in (0, 1]$ is a discount factor which places preference for immediate rewards over future ones. The expectation is taken over the sequence of uncertain states and observations in the future.

B. POMDP Model

1) State Modelling: The state in our problem consists of the road contexts Ω, the pose $\mathbf{p}$, the acceleration acc and the speed v of ego and exo-vehicles along with the intent ξ and the driving style β for exo-vehicles, at current time.

2) Action Modelling: We define our action space as $\{\text{LANE-KEEP}, \text{LEFT-LANE-CHANGE}, \text{RIGHT-LANE-CHANGE}\}$. The steering and speed of the ego-vehicle is computed according to the planned action using our transition model.

3) Observation Modelling: The observation comprises the Ω, $\mathbf{p}$, v and acc of each vehicle. To focus on modelling the hidden intent ξ and driving style β, we assume the observation over Ω, $\mathbf{p}$ and v has no noise.

4) Transition Modelling: The transition function models the movements of each vehicle. We used a driving style enhanced Time-To-Collision (TTC) trajectory predictor to model the movements. Our predictor improves on our previous trajectory predictor [1]. The TTC trajectory predictor predicts the trajectory of a vehicle using 5th order and 4th order polynomial curve fitting for lateral displacement and longitudinal displacement, respectively. It uses a time-to-collision model to predict the end state:

$$\dot{s}_1 = \min\left(v_{\max}, \sqrt{\max(0, v_{\text{f}}^2 + 2a_{\max}(d_{\text{f}} - d_{\text{safe}})}\right), \quad (3)$$

where $v_{\max}$ and $a_{\max}$ are the allowed maximum speed and acceleration, respectively, and v_{f} is the current speed of the front vehicle. d_{f} is the distance to the front vehicle and d_{safe} is a predefined and fixed-value safe distance. The end acceleration is computed accordingly by

$$\ddot{s}_1 = (\dot{s}_1 - \dot{s}_0)/t_1. \quad (4)$$

We improve on this TTC-based predictor by incorporating the driving style to it. Drivers with different driving style tend to keep different safe distances [27]. Instead of using predefined and fixed-value d_{safe}, we adjust d_{safe} based on the driving style inference of the surrounding vehicles. For example, we use a larger d_{safe} for conservative drivers, to model its behavior of safe driving. In contrast, if a driver is aggressive, we use a smaller d_{safe}, to model its behavior of aggressive driving. With our driving style enhanced trajectory predictor, we build our transition function by adding a Gaussian noise on the trajectory. Specifically, for each time step, the transition of the pose is defined by:

$$p(\mathbf{p}_{t+1} \mid \mathbf{p}_t, \xi, \Omega, \beta, acc_t, v_t) = f\left(\left\|\mathbf{p}_{t+1} - \mathbf{p}_{t+1}^{\text{pred}}\right\| \mid 0, \sigma^2\right), \quad (5)$$

where $\mathbf{p}_{t+1}^{\text{pred}}$ is the pose extracted from our predicted trajectory, and f is the probability density function of the Gaussian

distribution with mean 0 and variance σ^2. acc_{t+1} and v_{t+1} are computed using $\mathbf{p}_{t+1}$, $\mathbf{p}_t$ and v_t accordingly.

5) Reward Modelling: We designed the reward function considering both the time efficiency and the safety following the idea in [1]. We achieve time efficiency by penalizing the ego-vehicle when it drives on the lane farther to its destination by a penalty $R_d = -W_d \times (d/d_{\max})$ where d is the distance from current lane to the destination lane and $d_{\max}$ is the maximum inter-lane distance, and we also penalize for speed of the ego-vehicle by $R_v = W_v \times \frac{\|v - v_{\max}\|}{v_{\max}}$ to encourage it to select the lane on which it can drive faster. In addition, we penalize collisions with exo-vehicles with a penalty of $R_c = -W_c \times \max[(4-d)^2, 1]$, where $d < 4$ meters is the expected distance threshold between the two vehicles W_d, W_v and W_c are weights controlling the importance of each reward component and are empirically set to 100, 20, and 1000, respectively. The final reward is sum of R_d, R_v and R_c.

6) Initial Beliefs: The initial beliefs over the intents and driving style at each time step is inferred by the intent and the driving style predictors as proposed in Section III.

C. POMDP Solver

We use IS-DESPOT [2], a state-of-the-art online POMDP solver to solve our POMDP. To handle the high computational complexity of solving a POMDP, most solvers leverage Monte Carlo sampling in planning. However, this naive sampling strategy often overlooks rare but critical events, resulting in risky policy. A key feature of IS-DESPOT is the use of importance sampling [28] to handle rare but critical events. In autonomous driving, collisions are the rare but critical events. We use IS-DESPOT to handle collisions, generating risk-averse driving policies. To use IS-DESPOT, it requires us to design an importance sampling distribution q, which, in this case, can be interpreted as a deformed transition function with increased probability for critical states and decreased probability for others. We design q in a way similar to [2]. We increase the probability of the collision state for vehicles with erratic driving style.

V. EXPERIMENTAL RESULTS

We validate our proposed behavior planner, CID, using both simulated and real-world data. We compare the performance of CID with five baselines. All planners run at 4 Hz.

A. Baseline algorithms

We compare CID with the following baselines:

1) Reactive-Controller: Reactive-Controller [14] performs reactive actions based on ego-vehicle's distance to neighboring exo-vehicles.

2) SimmobilityST: SimmobilityST [29] uses a hierarchical decision making process for controlling the driving behavior of each vehicle centrally.

3) Pessimistic-Planner: This is a variant of our CID planner with an assumption that all exo-vehicles are adversarial.

4) Optimistic-Planner: This is also a variant of our planner with an assumption that all exo-vehicles are rational. The outcomes of this planner are similar to our previous work [1] where the driving styles of other road users is not accounted.

5) Max-Likelihood-Planner: This planner plans vehicle actions assuming the world state to be the most likely inferred state by our predictor. It plans by doing forward simulations starting from the current state and selecting the action branch with the highest reward.

B. Experiments in Simulation

We used SimMobility simulator for exhaustively testing a wide-range of scenarios with different intensities of adversarial behaviors for the exo-vehicles. The exo-vehicles are centrally controlled by SimMobilityST for normal driving styles and are modified for the erratic ones. The ego-vehicle is controlled by the corresponding baseline algorithm. We used real road network and the road contextual information for the simulations. For the qualitative analysis of our planner CID, we specifically designed five exemplar scenarios. Each of these scenarios present adversarial situations for the ego-vehicle's planner by executing either laterally or longitudinally erratic behaviors. These scenarios and their results from our planner are presented in Fig. 1.

Scenario 1 (Fig. 1, Row 1) presents one slow moving, laterally erratic, exo-vehicle which constantly oscillates within the lane and frequently changes lanes. Without inferring the laterally erratic driving style of the exo-vehicle, the ego-vehicle can easily mistake it for a lane change intent and can potentially result in a collision. For example, Optimistic-Planner and SimmobilityST baseline algorithms expect a rational driving style from the exo-vehicle and attempt to overtake it, resulting in a collision. In contrast, CID accounts for the laterally erratic driving style of the exo-vehicle and slows down to maintain a safety distance.

Scenario 2 (Fig. 1, Row 2) generates one slow moving, longitudinally erratic vehicle that remains in lane but constantly brakes and randomly changes speed. This driving style of the exo-vehicle can potentially cause rear-end collision with ego-vehicle. -Controller is saved from the potential collision as it changes lane to increase its headway distance, however, it ends up in the lane which takes a longer path to the destination. SimmobilityST and Pessimistic-Planner do not collide as they maintain a safe distance to the exo-vehicle though at the cost of taking a longer travel time to reach the destination. Optimistic-Planner attempts to overtake the exo-vehicle with the assumption that it would be rational and occasionally ends up in a collision when the exo-vehicle suddenly accelerates aggressively. Our planner, CID, infers the longitudinally erratic behavior of the slow moving exo-vehicle and changes lane. It then waits for a safety target gap and successfully overtakes the longitudinally erratic vehicle to align with the shortest path to the destination, thus accounting for both safety and time efficiency.

Scenario 3 (Fig. 1, Row 3) increases the adversity for the ego-vehicle with two longitudinally erratic vehicles in two lanes, providing only one option to avoid collision and reach the destination. Reactive-Controller surprisingly overtakes the first exo-vehicle however, with a very small safety margin. SimmobilityST and Pessimistic-Planner constantly follows the first exo-vehicle at a safe distance, taking a really

long time to reach the destination. Similar to Scenario 1, Optimistic-Planner attempts to overtake the first exo-vehicle with little gap between the two exo-vehicles and occasionally ends up in a collision due to a premature lane change. Our CID planner changes lane after inferring longitudinally erratic driving style of the first exo-vehicle. After changing lane, it finds the second vehicle and infers its behavior as longitudinally erratic as well, it then slows down and maintains safety distance to the second exo-vehicle until a gap is available for an overtake. It finally overtakes the first exo-vehicle and safely reaches the destination.

Scenario 4 (Fig. 1, Row 4) is a combination of two vehicles where the first one is longitudinally erratic and the second one is laterally erratic. Similar to Scenario 3, CID changes lane after inferring that the first exo-vehicle is longitudinally erratic. It then encounters the second exo-vehicle and infers that it is laterally erratic. Our planner then slows down and waits for a safe gap until the laterally erratic vehicle clears from the scene. It then increases its speed for continuing to the destination. Unlike our planner, all the other baselines except Pessimistic-Planner, occasionally results in a side-way collision with the laterally erratic exo-vehicle.

Scenario 5 (Fig. 1, Row 5) presents two laterally erratic exo-vehicles which eventually become normal. This scenario showcases the ability of our planner to dynamically update its actions while accounting for changes in behavior. Our CID planner first infers both exo-vehicles as erratic and slows down to keep a safe distance. Once the behavior of the first exo-vehicle is updated to normal, our planner passes it. Next, our planner guides the ego-vehicle to stay in lane and keep a safe distance to the second exo-vehicle. Lastly, when the behaviour of the second exo-vehicle is also updated to normal, the ego-vehicle increases its speed and safely continues its path to the destination. All the other baselines except Pessimistic-Planner, occasionally result in a side-way collision with one of the laterally erratic exo-vehicles.

We randomly sample from the aforementioned adversarial scenarios to generate a range of adversarial driving sequences for quantitative analysis of our planner with respect to the baselines. We executed 100 iterations per algorithm. The performance criteria for the quantitative results is based on two fundamental metrics for autonomous driving: safety and efficiency. The safety is measured by collision rate and the efficiency by the travel time to the destination. A summary of our results for CID planner compared to the baselines in adversarial scenarios is presented in TABLE II. It can be observed that Pessimistic-Planner and our CID planner, both have zero collision rate. However, the Pessimistic-Planner takes longer travel time, showcasing poor efficiency. In contrast, CID has reasonable travel time compared to average travel time for all baselines. Both Optimistic-Planner and Max-Likelihood-Planner achieve better travel time than our planner. However, they have higher collision rates because Optimistic-Planner believes all other drivers drive rationally while Max-Likelihood-Planner is too confident on the state it infers which could be wrong.

In addition, we also performed a comparative study when all exo-vehicles exhibit rational driving styles. The scenarios for the rational behavior are adapted from our previous work [1]. These results are in TABLE III. Since all the exo-vehicles behave rationally, the collision rate is observed to be zero for all planners. For travel time, our planner and Max-Likelihood-Planner perform the best with the least time.

Algorithm	Collision Rate	Travel Time (s)
Reactive-Controller	0.21	61
SimMobilityST	0.21	74.5
Pessimistic-Planner	0	121.5
Optimistic-Planner	0.12	60.8
Max-Likelihood-Planner	0.17	62
CID	0	72

TABLE II: Adversarial Scenarios: Average performance comparison over 100 randomly generated scenarios.

Algorithm	Collision Rate	Travel Time (s)
Reactive-Controller	0.0	78.7
SimMobilityST	0.0	71.7
Pessimistic-Planner	0.0	162
Optimistic-Planner	0.0	63
Max-Likelihood-Planner	0.0	61.3
CID	0.0	63.2

TABLE III: Normal scenarios: Average performance comparison over 100 randomly generated scenarios.

C. Experiments in Real World

We validated our planner using a self-driving car, SCOT [30] in a real urban environment of Singapore. We specifically tested Scenarios 1 and 2 with single exo-vehicle which is laterally and longitudinally erratic. The two scenarios are presented in Fig. 2 and 3. The high-level action of our CID planner is compared with the actions of an unbiased safety driver for the ego-vehicle who was not provided any information about the test scenarios. The speed of the ego-vehicle inferred by our transition model is also qualitatively compared with the human driver's decisions.

In Scenario 1, the vehicle initially infers the exo-vehicle as laterally erratic from its movement. CID outputs a lane keep action with slow speed to keep a safety distance. This action is in-line with the safety driver's decision. In frame 2 and 3, the exo-vehicle swerves between the two lanes, and CID outputs a right lane change action (indicated by red strip), while still keeping a slow speed, as the best option. Similarly, the safety driver decides to lane change. Finally, in frames 4 and 5, the exo-vehicle again attempts to block the ego-vehicle as an adversary. Both CID and the safety driver identifies the exo-vehicle to be lane-changing and maintains a safe speed in the current lane (blue strip).

For Scenario 2, frame 1, the ego-vehicle observes the exo-vehicle to be laterally erratic and maintains a safety distance while keeping in lane. In frame 2, CID outputs right lane change action. The safety driver also decides the same as seen in the following frame (frame 3). During frame 3, CID momentarily infers the exo-vehicle to be laterally erratic and slows down. It was observed that the safety driver had also slowed down. A potential reason is due to the fact that the exo-vehicle swayed towards right as there is an incoming merging lane on the left. In frames

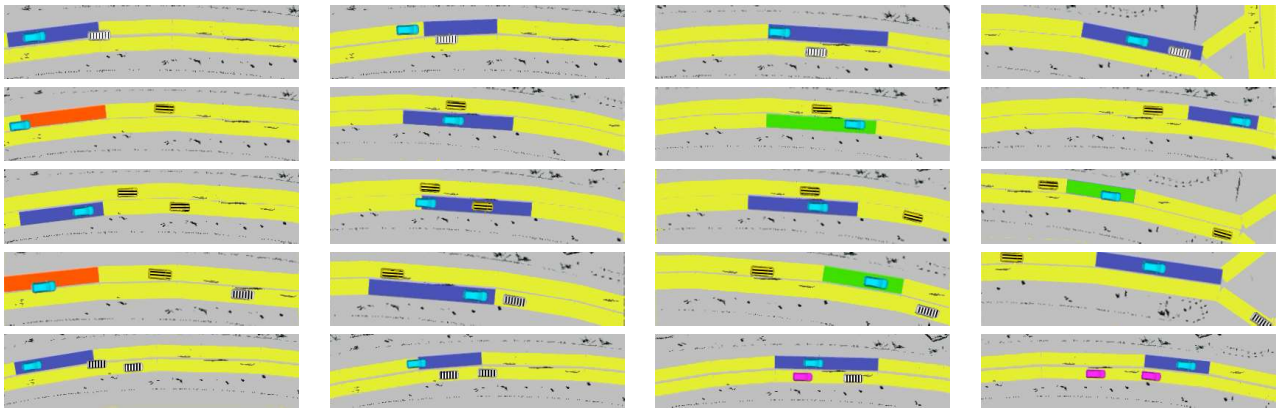

frame 1 *frame 2* *frame 3* *frame 4*

Fig. 1: Scenario 1 to 5 (row 1 to row 5), each with 4 frames. Ego-vehicle is in cyan. Exo-vehicles with normal, longitudinally erratic, and laterally erratic driving style are in violet, yellow and black stripes, and white and black stripes respectively. Actions are represented by blue, green and red lane strips for lane keep, left lane change and right lane change respectively.

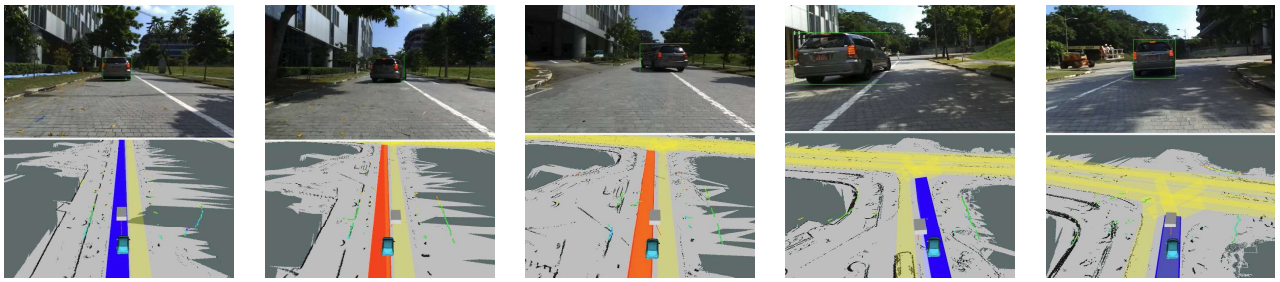

frame 1 *frame 2* *frame 3* *frame 4* *frame 5*

Fig. 2: Scenario 1 in real world. The ego-vehicle (cyan) encountered a slowly-moving laterally erratic vehicle. The actions suggested by the planner matched with those the driver actually executed.

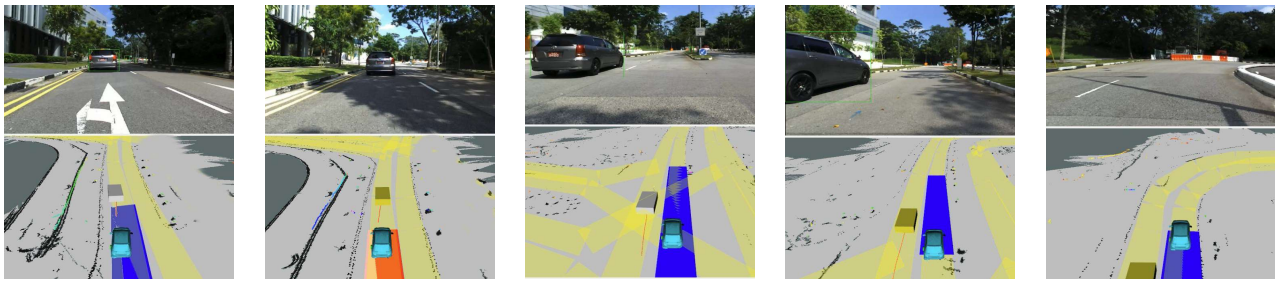

frame 1 *frame 2* *frame 3* *frame 4* *frame 5*

Fig. 3: Scenario 2 in real world. The ego-vehicle (cyan) encountered a slowly-moving longitudinally erratic vehicle. The actions suggested by the planner matched with those the driver actually executed.

4 and 5, our planner correctly infers the exo-vehicle to be longitudinally erratic (indicated as yellow) and accelerates in its lane. The safety driver is also observed to increase speed while passing by the exo-vehicle. The online video (`http://bit.ly/cid_planner_2021`) shows sample runs of CID both on the real vehicle and in simulation.

VI. DISCUSSION AND CONCLUSION

We presented a novel and interactive, context, intent and driving style aware POMDP planner, CID, for autonomous urban driving in adversarial scenarios. The road contextual information and driver intent inference used for our planner helped in time efficient planning, whereas the driving style inference accounted for safety. We validated our planner in challenging dynamic urban environments with adversarial scenarios. Our experimental results in simulation and on real-world data showcased significantly better performance of our planner compared to five different baseline algorithms.

ACKNOWLEDGMENT

This research was supported by the National Research Foundation, Prime Minister's Office, Singapore, under its AI Singapore Program (AISG Award No: AISG2-RP-2020-016), National University of Singapore (R-252-000-A87-114) and CREATE program, Singapore-MIT Alliance for Research and Technology (SMART), Future Urban Mobility, IRG.

REFERENCES

[1] Malika Meghjani, Yuanfu Luo, Qi Heng Ho, Panpan Cai, Shashwat Verma, Daniela Rus, and David Hsu. Context and intention aware planning for urban driving. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, 2019.

[2] Yuanfu Luo, Haoyu Bai, David Hsu, and Wee Sun Lee. Importance sampling for online planning under uncertainty. *Int. J. Robotics Research*, 38(2-3):162–181, 2019.

[3] Ishak Mohamad, Mohd Alauddin Mohd Ali, and Mahamod Ismail. Abnormal driving detection using real time global positioning system data. In *Proc. of IEEE Int. Conf. on Space Science and Communication*, 2011.

[4] Ahmad Aljaafreh, Nabeel Alshabatat, and Munaf S Najim Al-Din. Driving style recognition using fuzzy logic. In *IEEE Int. Conf. on Vehicular Electronics and Safety*, 2012.

[5] Chunmei Ma, Xili Dai, Jinqi Zhu, Nianbo Liu, Huazhi Sun, and Ming Liu. DrivingSense: Dangerous driving behavior identification based on smartphone autocalibration. *Mobile Information Systems*, 2017, 2017.

[6] Ernest Cheung, Aniket Bera, and Dinesh Manocha. Efficient and safe vehicle navigation based on driver behavior classification. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition Workshops*, pages 1024–1031, 2018.

[7] Nachiket Deo, Akshay Rangesh, and Mohan M Trivedi. How would surround vehicles move? a unified framework for maneuver classification and motion prediction. *IEEE Trans. on Intelligent Vehicles*, 3(2):129–140, 2018.

[8] Adam Houenou, Philippe Bonnifait, Véronique Cherfaoui, and Wen Yao. Vehicle trajectory prediction based on motion model and maneuver recognition. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, pages 4363–4369, 2013.

[9] Matthias Schreier, Volker Willert, and Jürgen Adamy. Bayesian, maneuver-based, long-term trajectory prediction and criticality assessment for driver assistance systems. In *17th Int. IEEE Conf. on Intelligent Transportation Systems*, pages 334–341. IEEE, 2014.

[10] Derek J Phillips, Tim A Wheeler, and Mykel J Kochenderfer. Generalizable intention prediction of human drivers at intersections. In *IEEE Intelligent Vehicles Symp.*, pages 1665–1670. IEEE, 2017.

[11] Sajan Patel, Brent Griffin, Kristofer Kusano, and Jason J Corso. Predicting future lane changes of other highway vehicles using rnn-based deep models. *arXiv preprint arXiv:1801.04340*, 2018.

[12] Florent Altché and Arnaud De La Fortelle. An LSTM network for highway trajectory prediction. In *IEEE 20th Int. Conf. on Intelligent Transportation Systems*, pages 353–359. IEEE, 2017.

[13] Enric Galceran, Alexander G Cunningham, Ryan M Eustice, and Edwin Olson. Multipolicy decision-making for autonomous driving via changepoint-based behavior prediction: Theory and experiment. *Autonomous Robots*, 41:1367–1382, 2017.

[14] Haoyu Bai, Shaojun Cai, Nan Ye, David Hsu, and Wee Sun Lee. Intention-aware online POMDP planning for autonomous driving in a crowd. In *Proc. IEEE Int. Conf. on Robotics & Automation*, pages 454–460. IEEE, 2015.

[15] Yuanfu Luo, Panpan Cai, Aniket Bera, David Hsu, Wee Sun Lee, and Dinesh Manocha. PORCA: Modeling and planning for autonomous driving among many pedestrians. *IEEE Robotics and Automation Letters*, 3:3418–3425, 2018.

[16] W. Gao, D. Hsu, W.S. Lee, S. Shen, and K. Subramanian. Intention-net: Integrating planning and deep learning for goal-directed autonomous navigation. In S. Levine and V. Vanhoucke andK. Goldberg, editors, *Conference on Robot Learning*, volume 78 of *Proc. Machine Learning Research*, pages 185–194. 2017.

[17] Constantin Hubmann, Jens Schulz, Marvin Becker, Daniel Althoff, and Christoph Stiller. Automated driving in uncertain environments: Planning with interaction and uncertain maneuver prediction. In *IEEE Trans. on Intelligent Vehicles*, pages 5–17. IEEE, 2018.

[18] Wei Liu, Seong-Woo Kim, Scott Pendleton, and Marcelo H. Jr. Ang. Situation-aware decision making for autonomous driving on urban road using online POMDP. In *IEEE Intelligent Vehicles Symposium*. IEEE, 2015.

[19] Tirthankar Bandyopadhyay, Kok Sung Won, Emilio Frazzoli, David Hsu, Wee Sun Lee, and Daniela Rus. Intention-aware motion planning. In *Algorithmic Foundations of Robotics X*, pages 475–491. Springer, 2013.

[20] Liting Sun, Wei Zhan, Masayoshi Tomizuka, and Anca D Dragan. Courteous autonomous cars. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, 2018.

[21] Malika Meghjani, Shashwat Verma, You Hong Eng, Qi Heng Ho, Daniela Rus, and Marcelo H. Ang Jr. Context-aware intention and trajectory prediction for urban driving environment. In *Int. Symp. on Experimental Robotics*. Springer, 2018.

[22] Land Transport Authority of Singapore. REVISED ROAD CROSS-SECTION DETAILS OF SAFEGUARDED ROADS . https://www.corenet.gov.sg/media/2186957/a-revised-road-cross-section-details-of-safeguarded-roads-b-including-proposals-involving-2-landed-houses-in-the-development-control-lodgement-to-road-transport.pdf, 2018.

[23] U.S. Federal Highway Administration. Next generation simulation program (ngsim), 2006. https://ops.fhwa.dot.gov/trafficanalysistools/ngsim.htm.

[24] Houtao Deng, George Runger, Eugene Tuv, and Martyanov Vladimir. A time series forest for classification and feature extraction. *Information Sciences*, 239:142–153, 2013.

[25] Markus Löning, Anthony Bagnall, Sajaysurya Ganesh, Viktor Kazakov, Jason Lines, and Franz J Király. sktime: A Unified Interface for Machine Learning with Time Series. In *Workshop on Systems for ML at NeurIPS 2019*.

[26] Richard D Smallwood and Edward J Sondik. The optimal control of partially observable Markov processes over a finite horizon. *Operations Research*, 21:1071–1088, 1973.

[27] Tie-Qiao Tang, Jia He, Shi-Chun Yang, and Hua-Yan Shang. A car-following model accounting for the drivers attribution. *Physica A: Statistical Mechanics and its Applications*, 413:583–591, 2014.

[28] M.H. Kalos and P.A. Whitlock. *Monte Carlo Methods*, volume 1. John Wiley & Sons, New York, 1986.

[29] Muhammad Adnan, Francisco C Pereira, Carlos Miguel Lima Azevedo, Kakali Basak, Milan Lovric, Sebastián Raveau, Yi Zhu, Joseph Ferreira, Christopher Zegras, and M Ben-Akiva. Simmobility: A multi-scale integrated agent-based simulation platform. In *95th Annual Meeting of the Transportation Research Board Forthcoming in Transportation Research Record*, 2016.

[30] Scott Drew Pendleton, Hans Andersen, Xiaotong Shen, You Hong Eng, Chen Zhang, Hai Xun Kong, Wei Kang Leong, Marcelo H Ang, and Daniela Rus. Multi-class autonomous vehicles for mobility-on-demand service. In *IEEE/SICE Int. Sym. on System Integration*, 2016.