APPEARED IN
Proc. IEEE/RSJ International Conference on Intelligent Robots & Systems, 2019

Context and Intention Aware Planning for Urban Driving

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

***Abstract*— We present a novel autonomous driving system which uses the road contextual information and intentions of other road users for urban driving. Unlike highways, urban environments require the drivers to follow traffic signs and signals while using their best judgment for anomalous situations. In such scenarios, a self-driving car needs to understand and take into account the uncertainties in the environment to plan and decide its action accordingly. Our planner models the intentions of the surrounding vehicles leveraging a neural network, and integrates the road contextual information to reduce its environment uncertainties and also speed up the decision making process. We validate our planner in simulation and in a real urban environment. Our experimental results show that integrating intention inference and road contextual information for prediction, planning and decision making helps improve safety and efficiency of our autonomous driving system.**

I. INTRODUCTION

We propose a novel autonomous driving system which uses the road contextual information and intention of other road users to provide safe and efficient high-level driving actions for urban driving. Urban environments pose a unique set of challenges for self-driving cars. Consider a scenario for overtaking an illegally parked car from the opposite side of a two-way street with a single lane on each side. In order to decide whether to overtake the parked car, the ego-vehicle has to understand the road contextual information such as the number of lanes, direction of the lane, lane width and distance to the nearest intersection, etc.. Otherwise, it can cause severe safety hazards due to misjudgments. Another challenge in urban driving is the need for long-term planning while interacting with multiple exo-vehicles. A driving system has to perform long-horizon planning in a large state space composed of all neighboring vehicles, so that the ego-vehicle avoids collisions with them while efficiently navigating to its goal. The key aspect, in this case, is to predict the long-term behaviours of exo-vehicles and plan for the ego-vehicle, accordingly.

We developed a hierarchical prediction model for long-term planning. At the high level, we model the *intentions* of exo-vehicles, such as keeping the current lane, changing to the left lane, or changing to the right lane. At the low level, we use a polynomial curve fitting method to predict the vehicles' actual motion conditioned on their intentions. This hierarchical model allows us to learn the driving intentions independently of driving trajectories, thus requiring limited amount of training data to learn the correlations between the road contextual information and driver's intentions. Since the trajectories can vary significantly across different drivers, the online polynomial curve fitting model helps us capture this variance, aptly.

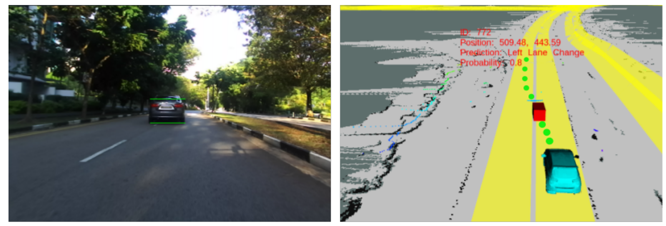

Fig. 1: Intention and trajectory prediction of the exo-vehicles given their past trajectories and the road contextual information. The perception system detects the vehicle with a bounding box (left). The intention and trajectory prediction results are presented on the right in text and as green markers respectively.

The planning process can be formalized as a Partially Observable Markov Decision Process (POMDP) [1] which provides a principled way to handle uncertainties such as partial observability, action noise and sensing noise. However, POMDP planning suffers from its well-known high computational complexity. In order to achieve real-time planning, we further use road contextual information to assist the search by pruning invalid actions and shaping the rewards.

Our contributions in this paper are: (a) an autonomous driving framework that can track and predict intentions and trajectories of multiple exo-vehicles, (b) a context-aware prediction model which decouples intention and trajectory prediction for exo-vehicles, (c) a context and intention aware planner that determines long-term high-level actions of the ego-vehicle under uncertainties of exo-vehicles' intentions. We evaluate our driving system qualitatively and quantitatively in a range of scenarios using a simulator which integrates real-world road networks and the relevant contextual information. We also analyzed our system performance on real-world data. Our results show that integrating road contextual information and intention inference into long-term planning helps improve the efficiency and safety of the system.

*Both authors contributed equally.

Malika Meghjani, Qi Heng Ho and Shashwat Verma are with the Singapore-MIT Alliance for Research and Technology[1], Singapore. Malika is also affiliated with Singapore University of Technology and Design[2]. `malika_meghjani@sutd.edu.sg`, `qiheng@smart.mit.edu`, `shashwatverma14@gmail.com`

Yuanfu Luo, Panpan Cai and David Hsu are with the National University of Singapore [3], Singapore `{yuanfu,caipp,dyhsu}@comp.nus.edu.sg`

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

II. RELATED WORK

A. Road contextual information

Given the widespread availability of precise and high resolution digital map information [2], [3], the possibility of using the road contextual information as a prior for behavior prediction is becoming increasingly prevalent. Specifically, in [4], the authors propose a one-stage detector and behavior predictor using two different 2D convolution neural networks, each one processing 3D point cloud data and dynamic map information. Similar to our work, their map information comprises static road features such as lanes, intersections, crossing and traffic signs. They categorize behavior prediction into high-level actions and motion estimation using the same neural network framework. In our work, we use road contextual information for not merely behaviour prediction but also for high-level planning.

B. Intention Inference and Trajectory Prediction

Intention and trajectory prediction of dynamic obstacles are two key components for decision making of autonomous vehicles. A survey of state-of-the-art intention and trajectory prediction algorithms is presented in [5]. Neural networks are popularly known to be useful for intention prediction [6]. For example, neural networks were used for both intention and trajectory prediction for vehicles on highways in [7] and [8], and for trajectory prediction only in [9]. These approaches, however, require a huge amount of data for learning a range of driving behaviors. We overcome the challenge of large data requirements by using a neural network to only predict the intention whereas the trajectory of predicted intention is obtained based on polynomial fitting and extrapolating the real-time vehicle state [10]. This allows us to predict the intention and trajectory in real-time.

C. Planning with human intentions

Several previous work use POMDPs to handle the uncertainty in human drivers' intentions. Some of this research focuses on intersection scenarios with a small number of agents. In [11], the authors modeled exo-vehicles' intended behaviours: to drive aggressively or patiently, as hidden variables of the POMDP. They then solved the POMDP offline to control ego-vehicle's speed at the intersection. Due to the high computational cost of offline planning, the approach has only been tested in two-vehicle interactions. A similar approach was taken in [12], but they infer the intention of other vehicles using a reaction-based probabilistic model. Noticeably, the work used a rich representation of road contexts to help predict other vehicles' motion. However, this representation also induces high computational complexity. Thus their results only discussed interactions among 2$\sim$3 vehicles. Our system uses a much more concise notion of road contexts. Recently, POMDP planning is applied to leverage a road network known a-priori to the robot vehicle for driving at intersections [13]. Their method models the intended paths of exo-vehicles on the road network as hidden variables and also control the vehicle speed.

We argue that the lane merging problem is much harder than the intersection case, because exo-vehicles have a lot more freedom: drivers can choose to merge lane at any moment when they feel it promising. The decision making of lane merging scenarios has been studied by several previous work. A simplified approach [14] was proposed to make the lane merging problem tractable: evaluate a fixed set of policies by sampling exo-vehicles' intentions and rolling out future interaction trajectories. A similar multi-policy decision making approach [15] has been applied to navigation among multiple pedestrians. These multi-policy methods only plan for one-time interaction with other agents. However, in real-life driving scenarios, long-term interactions are often required, e.g., executing multiple lane merges to reach a faster lane. Another set of work [16], [17] addresses the lane changing problem from the perspective of active information gathering. Both work apply exploration bonuses on the reward function to encourage probing actions that bring information to better understand human drivers' inner states. Again, these work only focused on the interactions among a small number of vehicles.

Another group of work studied the interaction with multiple agents, typically pedestrians, which incurs exponentially higher complexities than the aforementioned scenarios. It has been proposed in [18] to model pedestrian intentions as a finite set of goals and apply a simple goal-directed motion model to predict pedestrian motions. The method used a state-of-the-art online POMDP planning algorithm DESPOT [19] to handle moderately dense crowds. A recent work [20] proposed a more sophisticated pedestrian motion model, PORCA, and integrated it into parallel POMDP planning [21] to drive an autonomous vehicle among many pedestrians. Different from these work for free-walking pedestrians, our system interacts with multiple vehicles on urban roads, in which case it is important to leverage road contexts in intention inference, motion prediction, and decision making.

III. OVERVIEW

The overview of our autonomous driving system is presented in Fig. 2. It comprises three sub-systems: perception, high-level decision maker, and low-level controller. In the perception system [22], a LiDAR-based point cloud clustering module and a vision-based obstacle detection and classification module are used for identifying the road region and the vehicles. Their outputs are fused in the sensor fusion module for vehicle tracking. The high-level decision maker receives the tracking trajectories, infers a belief over the intentions of each vehicle, then plans lane-keeping/lane-changing action based on the inferred intention belief and the road contextual information. The planned action can then be sent to the low-level controller to plan a path and provide the steering and throttle control that tracks the path. The low-level controller is however not addressed in this paper.

This work focuses on developing the decision maker to plan for high-level action, specifically, the action of keeping lane, changing to left lane, or changing to right lane, for the ego-vehicle, to achieve context and intention aware planning

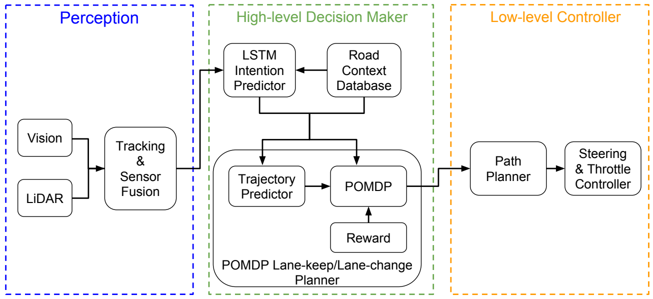

Fig. 2: System overview comprising of perception, high-level decision making and low-level control modules.

for autonomous driving. We first use an LSTM intention predictor to infer a belief over the current intention, i.e., the high-level action, of each nearby exo-vehicle, based on their trajectories and the road contextual information. The intention belief and the road contextual information are then input to a POMDP planner to plan for the lane-keeping/lane-changing action. At each time step, the planner uses belief tree search [19] to reason about the uncertain intentions of exo-vehicles. The belief tree search takes as input the current belief and the state of all vehicles and performs Monte Carlo simulations for the uncertain future to provide the largest rewarding high-level action.

IV. CONTEXT AND INTENTION AWARE PLANNING

We present our high-level decision maker in this section. The decision maker receives the tracking trajectories of exo-vehicles, and outputs the lane-keeping/lane-changing action for ego-vehicle. It mainly consists of three modules: a road context database, an LSTM intention predictor, and a POMDP high-level planner. Inside the POMDP planner, a trajectory predictor is used as transition function to simulate the future trajectories of all vehicles. We condition both the LSTM intention predictor and the POMDP planner on the road contextual information.

A. Road Contextual Information

The road contextual information is obtained from the Open Street Map [2]. The contextual information comprises the number of lanes, the direction of the lanes, the lane width, and the positions of the start and end points of each lane. We further pre-compute the path corresponding to the lane center which are feasible to navigate by a non-holonomic vehicle, and the shortest distance between the end points of each lane pair. The lane center and the lane directions are overlaid on the map and illustrated in Fig. 3.

B. Intention Inference

We use a recurrent neural network architecture to infer the belief over the current intentions of the exo-vehicles. We formalize the intention as the high-level action, i.e., the action of lane-keeping, left-lane-changing, or right-lane-changing, of the exo-vehicle. Concretely, we implement a Long Short Term Memory (LSTM) network to infer the intention belief. The LSTM is trained using generic features

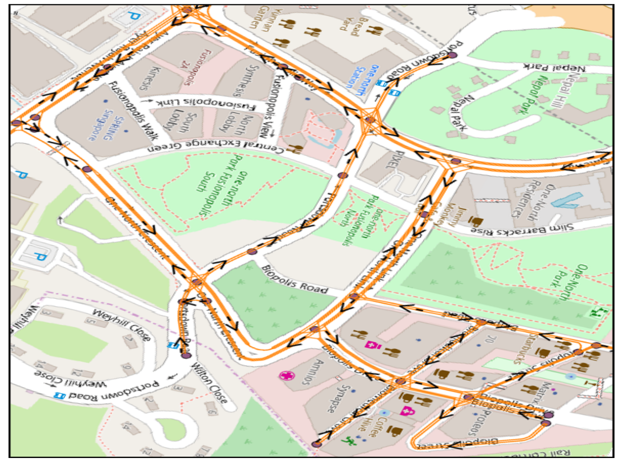

Fig. 3: Road contextual information provided by Open Street Map. The lane centers is marked in yellow along with the direction in black.

which are representative of the lane-keeping/lane-changing intention and also provide the road contextual information [10]. This allows the predictor to collectively learn the correlation between the intentions and the road contextual information. These features include,

- δx: Change in lateral pose
- δy: Change in longitudinal pose
- ll: If left lane exists (Boolean $0/1$)
- rl: If right lane exists (Boolean $0/1$)
- d_{center}: Distance from the center of the lane.

where, d_{center} value ranges from -1.5m to 1.5m for an average lane width of 3m. The LSTM network is trained with the above features for trajectories of 4 time steps where each time step is of 0.25 second. The hidden state vector of the LSTM is updated at each time step based on the hidden state at the previous time step and the input vector at the current time step. The output of the final state of the LSTM is fed into a fully-connected (FC) layer with a softmax activation function. The output of the FC layer is a belief, i.e., probability distribution, over three classes: lane-keeping, right lane-changing, and left lane changing. We use 256 units for the LSTM layer and fully connected layer. The model is trained to minimize the sum of the cross-entropy losses of the predicted and ground truth classes, and trained using Adam (adaptive moment estimation) optimization algorithm [23] with a learning rate of 0.0001. The training data for LSTM is obtained using our ego-vehicle so that we can precisely label the start and end points of desired intention sequence.

C. Trajectory Prediction

Given the road contextual information and the past poses of the target vehicle, we predict the vehicle's trajectory for each given intention. The past poses of the vehicle are modeled by 5^{th} and 4^{th} order polynomial representing the lateral and longitudinal components of the trajectory as suggested in [24]. The lateral displacement of the vehicle $d(t)$ changing with time t is represented by:

$$d(t) = c_5 t^5 + c_4 t^4 + c_3 t^3 + c_2 t^2 + c_1 t + c_0, \qquad (1)$$

where $c_{i,i=\{0,1,2,3,4,5\}}$ are the unknown coefficients. Given the initial and final states of the vehicle at a start time $t_0 = 0$ and a predefined end time t_1 respectively, we can solve for the unknown coefficients of the polynomial using:

$$\begin{bmatrix} t_0^5 & t_0^4 & t_0^3 & t_0^2 & t_0 & 1 \\ t_1^5 & t_1^4 & t_1^3 & t_1^2 & t_1 & 1 \\ 5t_0^4 & 4t_0^3 & 3t_0^2 & 2t_0 & 1 & 0 \\ 5t_1^4 & 4t_1^3 & 3t_1^2 & 2t_1 & 1 & 0 \\ 20t_0^3 & 12t_0^2 & 6t_0 & 2 & 0 & 0 \\ 20t_1^3 & 12t_1^2 & 6t_1 & 2 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} c_5 \\ c_4 \\ c_3 \\ c_2 \\ c_1 \\ c_0 \end{bmatrix} = \begin{bmatrix} d_0 \\ d_1 \\ \dot{d}_0 \\ \dot{d}_1 \\ \ddot{d}_0 \\ \ddot{d}_1 \end{bmatrix} \tag{2}$$

where d_i, $\dot{d}_i$ and $\ddot{d}_i$, $i = \{0, 1\}$ are the displacement, speed and acceleration in the lateral direction at the start time ($i = 0$) and the end time $i = 1$, respectively. We can compute d_0, $\dot{d}_0$ and $\ddot{d}_0$ easily from the vehicle's past trajectory. The final displacement, d_1 is selected based on the predicted intention whereas, the final lateral speed $\dot{d}_1$ and acceleration $\ddot{d}_1$ are set to be zero because the vehicle usually will not move in the lateral direction after it reaches the intended lane.

The longitudinal displacement $s(t)$ of the target vehicle is represented as:

$$s(t) = c_4t^4 + c_3t^3 + c_2t^2 + c_1t + c_0. \tag{3}$$

Similar to the lateral component, the unknown coefficients $c_{i,i=\{0,1,2,3,4\}}$ for the longitudinal component of the trajectory are solved using Eq. (4), given the initial state at $t_0 = 0$ and the final state at the predefined end time t_1.

$$\begin{bmatrix} t_0^4 & t_0^3 & t_0^2 & t_0 & 1 \\ 4t_0^3 & 3t_0^2 & 2t_0 & 1 & 0 \\ 4t_1^3 & 3t_1^2 & 2t_1 & 1 & 0 \\ 12t_0^2 & 6t_0 & 2 & 0 & 0 \\ 12t_1^2 & 6t_1 & 2 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} c_4 \\ c_3 \\ c_2 \\ c_1 \\ c_0 \end{bmatrix} = \begin{bmatrix} s_0 \\ \dot{s}_0 \\ \dot{s}_1 \\ \ddot{s}_0 \\ \ddot{s}_1 \end{bmatrix} \tag{4}$$

We can easily compute the displacement s_0, the speed $\dot{s}_0$ and the acceleration $\ddot{s}_0$ in the longitudinal direction for the initial state using the vehicle's past trajectory. However, the final longitudinal speed $\dot{s}_1$ and acceleration $\ddot{s}_1$ are unknown and need to be inferred. Constant velocity (Const-Vel) and constant acceleration (Const-Acc) models are two most simple and popular models to infer them. Const-Vel sets $\dot{s}_1 = \dot{s}_0$ and $\ddot{s}_1 = 0$. Const-Acc sets $\dot{s}_1 = \dot{s}_0 + \ddot{s}_0 t_1$ and $\ddot{s}_1 = \ddot{s}_0$.

However, Const-Vel and Const-Acc do not predict well in the presence of other vehicles, since they do not consider the interactions between vehicles. To address this issue, we propose to use a time-to-collision model (TTC) to predict the longitudinal speed. TTC computes the speed $\dot{s}_1$ for the target vehicle based on the distance d_f to the front vehicle in its target lane, along with the speed and maximum deceleration, assuming the worst case scenario, i.e., front vehicle suddenly decelerates with its maximum deceleration. Since it is impossible to know the maximum deceleration of the exo-vehicles, we assume it is the same as the maximum deceleration a_max of the ego-vehicle. The speed of the target vehicle is computed by:

$$\dot{s}_1 = \min\left(v_\text{max}, \sqrt{\max(0, v_\text{f}^2 + 2a_\text{max}(d_\text{f} - d_\text{safe})}\right), \tag{5}$$

where v_max is the maximum speed allowed, v_f is the current speed of the front vehicle and d_safe is a predefined safe distance. The acceleration is computed accordingly by

$$\ddot{s}_1 = (\dot{s}_1 - \dot{s}_0)/t_1. \tag{6}$$

Now with $d(t)$ and $s(t)$, we can predict the lateral and longitudinal displacements of the vehicle. Using the road contextual information and predicted intention, we can project those displacements into the target lane and get the predicted trajectory of the vehicle.

D. Context and Intention-aware POMDP

We formulate the high-level decision making process in urban environments as a POMDP model. POMDP provides a principled general framework for planning under uncertainty.

1) POMDP Preliminaries: A POMDP model is formally defined by a tuple (S, A, Z, T, O, R, b_0), where S, A and Z represent the state space, the action space, and the observation space, respectively. The transition function $T(s, a, s') = p(s'|s, a)$ represents the probability of reaching a state s' after the robot executes an action a in state s; it characterizes the imperfect robot control and environment dynamics. The observation function $O(s, a, z) = p(z|a, s)$ represents the probability of receiving an observation z after the robot executes a and reaches s; it models the sensor noises. The reward function $R(s, a)$ defines a real-value reward for executing a in s. The robot does not know the exact state it is in because of imperfect sensing. Therefore, it maintains a belief, i.e., probability distribution, over S. Initially, the robot's belief is b_0, and it gets updated via Bayes' rule at each time step. The aim of POMDP planning is to find a *policy* π, a mapping from a belief b to an action a, that maximizes the expected total discounted rewards:

$$V_\pi(b) = \mathbb{E}\left(\sum_{t=0}^{\infty} \gamma^t R(s_t, \pi(b_t)) \;\middle|\; b_0 = b\right), \tag{7}$$

where s_t is the state at time t, $\pi(b_t)$ is the action that the policy π chooses at time t, and $\gamma \in (0, 1)$ is a discount factor that places preferences for immediate rewards over future ones.

2) Context and Intention Aware POMDP: We construct the context and intention aware POMDP model for lane-keep/lane-change decision making in urban environments.

State Modeling. In our problem formulation, the state is defined as a combination of road contextual information c, pose (x, y, θ) of each vehicle, and the intention g of each exo-vehicle. The intentions of exo-vehicles are not observable to the ego-vehicle; we formalize them as hidden variables in the state.

Action Modeling. We plan the high-level action for the ego-vehicle at each time step, including LANE-KEEP, LEFT-LANE-CHANGE and RIGHT-LANE-CHANGE. We further prune the forbidden actions in different lanes with the help of road contextual information.

Observation Modeling. The observation consists of the road contextual information, poses, speeds and intentions of all vehicles. We do not consider sensor noises on the observations, to focus on modelling the uncertainty in intentions of the exo-vehicles.

Transition Modeling. The transition function models the movements of each vehicle under different intentions. For each vehicle, we use the trajectory predictor from Section IV-C to predict its next-step pose, given its intention and the road contextual information. By adding a Gaussian noise on the pose, for each vehicle, we obtain a transition function:

$$p(x_{t+1}, y_{t+1}, \theta_{t+1} | h(t), g, c), \tag{8}$$

where $h(t) = \{x_t, y_t, \theta_t, ..., x_{t-3}, y_{t-3}, \theta_{t-3}\}$ is a 4-time-step history of the past poses. For the ego-vehicle, its intention is represented by its high-level action.

Reward Modeling. The reward function is designed for the sake of safety, efficiency and smoothness of driving. To achieve safety, we penalize the collision with exo-vehicles with a huge penalty $R = -1000 \times \max[(4-d)^2, 1]$, where $d < 4$ meters is the distance between the two vehicles in collision. We divide the reward for efficiency into global reward and local reward. For global reward, we assign a reward $R = 0$ when the vehicle reaches its destination and a penalty $R = -100 \times (d/d_{\max})$ where d is the distance from current lane to the destination lane and $d_{\max}$ is the maximum inter-lane distance, to penalize the ego-vehicle for choosing a lane that is farther to the destination. For local reward, we assign a penalty $R = 20 \times \frac{v - v_{\max}}{v_{\max}}$ to encourage the ego-vehicle to choose a lane on which it can drive faster. For smoothness, we assign a penalty $R = -1$ for doing lane changes to avoid excessive lane changes. The final reward is the weighted sum of the aforementioned individual rewards.

Initial Belief. We use the LSTM intention predictor from Section IV-B to infer the belief over the intentions of each exo-vehicle, and use the inferred belief as the initial belief of the POMDP model.

3) Solving POMDP: We use Determinized Sparse Partially Observable Trees (DESPOT) [19] for solving the lane-keeping/lane-changing POMDP. DESPOT is one of the fastest online POMDP solvers. The key idea of DESPOT is to search a belief tree under K sampled scenarios only, which greatly reduces computational complexity, making it an efficient solver for our POMDP model.

V. EXPERIMENTAL RESULTS

We validate our POMDP high-level planner with both simulated and real-world data, both qualitatively and quantitatively. Specifically, for qualitative analysis, we designed four scenarios to validate the behavior of our planner, demonstrating the benefit of using the contextual information, intention inference, and long-term planning. For quantitative analysis, we randomly generated scenarios in simulation and tested the performance of our planner, and compared the average results with those of baselines.

In the following, we will first introduce our baseline algorithms and the criteria for our performance comparison. Then we will introduce the four scenarios we designed and analyze our results with both simulated and real-world data.

A. Baseline algorithms

We compared our algorithm, Context-Intention-POMDP, with three baselines: Reactive-Controller, Greedy-Controller, and SimMobilityST.

1) Reactive-Controller: This controller reacts based on the distances from the ego-vehicle to the front exo-vehicles in different lanes. It first gets the headway distances from the perception system, i.e., the distances to the nearest front vehicle, in the current lane, the headway distance in the left lane (if left lane exists), and the headway distance in the right lane (if right lane exists). It then compares those headway distances with a distance threshold D. It chooses to keep lane if the headway distance of the current lane is larger than D. Otherwise, it chooses to change to the lane with the largest headway distance. We set $D = 20$ meters in our experiment.

2) Greedy-Controller: The controller greedily chooses to drive in the lane that is shortest to the destination lane at each time step regardless of exo-vehicles.

3) SimMobilityST: SimMobilityST is a rule-based algorithm used in SimMobility short-term simulator [25] to centrally control the driving behavior of the simulated vehicles. SimMobilityST uses a four-level decision maker to model the individual driving behavior of each vehicle. These decisions include, (a) target lane selection, (b) gap acceptance of the leading and lagging vehicles, (c) the target gap and (d) desired acceleration. The first decision of target lane selection, is based on a global path following model. The vehicles select the target lane such that they follow the global path assigned to them. Once the target lane is decided, the second decision of accepting the gap distance to the leading and lagging vehicles is made by comparing the available gap with the critical gap. If the gap to the neighbouring vehicles is acceptable, the target gap of the ego-vehicle after reaching the target lane is decided based on the expected speeds of the neighbouring vehicles. Finally, the desired acceleration is decided based on the driver's state: car-following (in original lane), lane-changing, and recovering after lane-changing (in target lane).

B. Criteria for Performance Comparison

We compared our planner, Context-Intention-POMDP, with the baseline algorithms in terms of safety, efficiency, and smoothness. We measure the safety by the collision rate, the efficiency by the success rate and travel time, and the smoothness by the number of lane changes per 100 meters. One trial is considered unsuccessful if the ego-vehicle did not reach the goal within 3 minutes or collision occurs. We computed the average results for travel time and number of lane changes only for successful trials.

C. Experiments in Simulation

We experimented in a simulator, SimMobility, in order to test a wide-range of scenarios for our high-level motion planning algorithm, without any influence of the perception errors. SimMobility incorporates the real-world road network and contextual information. The exo-vehicles in the simulator are controlled by SimMobilityST algorithm.

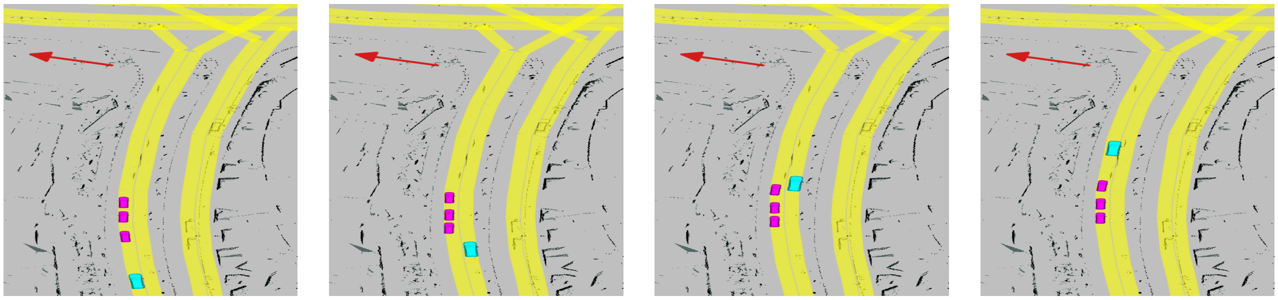

(*a*) (*b*) (*c*) (*d*)

Fig. 4: Scenario 1. The ego-vehicle (cyan) overtakes the slowly-moving exo-vehicles (violet) by changing to the right lane first and then changing back. The red arrow indicates the destination of the ego-vehicle. The yellow lines represent the lanes.

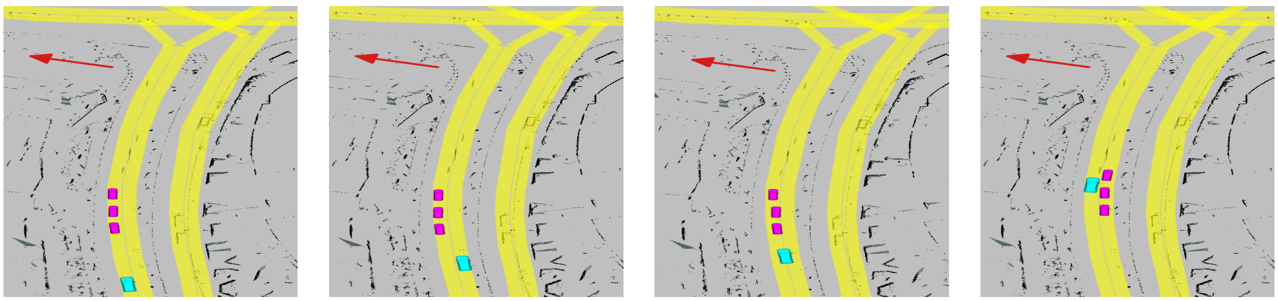

(*a*) (*b*) (*c*) (*d*)

Fig. 5: Scenario 2. The ego-vehicle (cyan) first changes to the right lane. After inferring the exo-vehicles' intentions of driving to the right lane, it immediately changes to the left lane which saves travelling time and avoids near collisions.

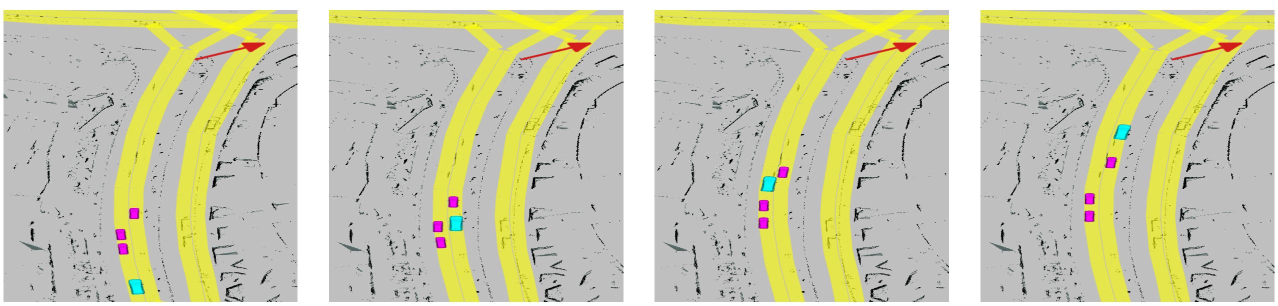

(*a*) (*b*) (*c*) (*d*)

Fig. 6: Scenario 3. The ego-vehicle (cyan) executes a sequence of complicated maneuvers to arrive at its destination faster: changing to the right lane, then waiting for the gap to increase, and at last overtaking the front vehicle.

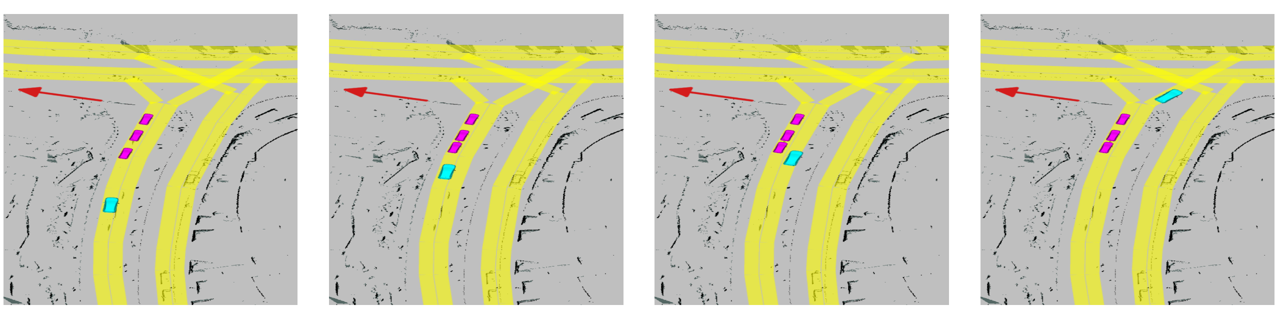

(*a*) (*b*) (*c*) (*d*)

Fig. 7: Scenario 4. The ego-vehicle (cyan) encounters a traffic jam. After waiting for a while, it decides to change to the right lane, which leads to a detour but is free to go.

We first analyzed our planner qualitatively using four specially designed scenarios. Fig. 4 to Fig. 7 shows the driving behaviour of our planner on those four scenarios. The direction towards the destination location is indicated by the arrow sign in each scenario.

Scenario 1 (Fig. 4) demonstrates the need of integrating road contextual information into planning. Without knowing the contextual information, the ego-vehicle will not be able to overtake or change back to the left lane after it overtakes the exo-vehicles. The latter is required as the right lane leads to a bigger detour to the destination. For example, in this scenario, a vehicle using Greedy-Controller fails to overtake the front vehicles while, Reactive-Controller fails to change back to the left lane after passing the slow vehicles.

Scenario 2 (Fig. 5) shows the benefit of integrating intention inference into planning. Using our Context-Intention-POMDP planner, the ego-vehicle starts to change to the left lane (Fig. 5(c)), which is closer to its destination, immediately after it infers that the exo-vehicles intend to change to the right lane. Both Reactive-Controller and SimMobilityST start to change to the left lane only after all the exo-vehicles complete their lane changes, which is inefficient and may cause collisions.

Scenario 3 (Fig. 6) and Scenario 4 (Fig. 7) illustrates the benefit of long-term planning. In Scenario 3, the ego-vehicle using Context-Intention-POMDP executes a sequence of complicated maneuvers: changing to the right lane, waiting for a safe gap to overtake the left vehicle and lastly returning to the right lane to go towards its destination. This requires long-term planning. In Scenario 4, we simulate the traffic jam by letting all exo-vehicles stop at the intersection. Using Context-Intention-POMDP, the ego-vehicle first waits for a while, then changes to the right lane, which takes it farther away from its destination but at a faster speed.

We also quantitatively evaluated our Context-Intention-POMDP planner and compared them to the baseline methods, using 200 randomly generated scenarios.

Algorithm	Collision Rate	Travel Time (s)	Num. Lane Changes per 100m	Success Rate
Reactive-Controller	0.0	78.7	0.6	0.725
Greedy-Controller	0.0	84.9	**0**	1.0
SimMobilityST	0.0	71.7	0.4	1.0
Context-Intention-POMDP	0.0	**65.7**	0.7	1.0

TABLE I: Average performance comparison on 200 randomly generated scenarios. The average results are computed using only the successful trials.

In each scenario, we fixed the start and destination of all the vehicles, but randomly generated speed, and time of lane changes for 5 exo-vehicles along a path of length 324.4 meters. We set the maximum speed of the ego-vehicle to be 6.5m/s. The average performance of our algorithms and the baselines are shown in TABLE I. Overall, our planner can reach the destination much faster with a slightly larger number of lane changes. One of the potential reason for increased number of lane changes for our planner is related to our choice of smoothness reward which has the least weight compared to safety and efficiency reward values.

D. Experiments in Real World

We validated our planner on a self-driving car (SCOT) in a real urban environment, in Singapore, with several parked cars on the roadside, slowly moving vehicles and pedestrians at crossings. Specifically, we created three scenarios similar to Scenario 1, 2 and 3 from Section V-C with a leading vehicle. The ego-vehicle was controlled by a safety driver and his spontaneous actions are compared with the high-level action outputs of our planner. The safety driver was not provided with any information about the scenarios to avoid any biasness in the results.

The high-level actions of our planner for Scenario 1 are presented in Fig. 8. The LANE-KEEP, LEFT-LANE-CHANGE and RIGHT-LANE-CHANGE actions are represented by blue, green and red colours, respectively. Initially in frame 1, the ego-vehicle detects the leading vehicle at a reasonable distance and keeps lane. It then infers that the exo-vehicle is moving slowly and suggests a RIGHT-LANE-CHANGE in frame 2. The safety driver eventually changes lane from opposite side of the road while the planner is still suggesting RIGHT-LANE-CHANGE in frame 3. Lastly in frame 4, the planner suggests to merge back to its original lane as soon as a safe gap is observed. This once again matched the safety driver's action. Our perception system does not differentiate trajectories between pedestrians and vehicles. But interestingly, our planner still manages to handle the scenario with pedestrians: after returning to the original lane in frame 5, a pedestrian is observed and the planner briefly suggests RIGHT-LANE-CHANGE based on its limited observations, and after predicting the pedestrian's trajectory of moving towards the right lane, it consistently suggests to keep lane in frame 6, similar to the safety driver's action.

We successfully demonstrated that our context and intention aware planner suggested actions in advance which matched the safety driver's actions for all three scenarios. These real urban environment scenarios and the simulated results are presented in the following video link: https://youtu.be/psm6juPltJs.

VI. DISCUSSION AND CONCLUSION

We presented a complete framework for autonomous driving in urban environments based on our novel context and intention aware planner. The use of road contextual information allowed us to reduce the uncertainties in intention and trajectory predictions of exo-vehicles, as well as, reduce the computation cost of our planner for the ego-vehicle. The presented simulation and real world results suggest that we can plan safe and time efficient decisions over a range of urban road scenarios when compared to the baseline algorithms.

Our work, however, does have some limitations. Specifically, our intention prediction module considers the exo-vehicles independently, ignoring the influence of interactions

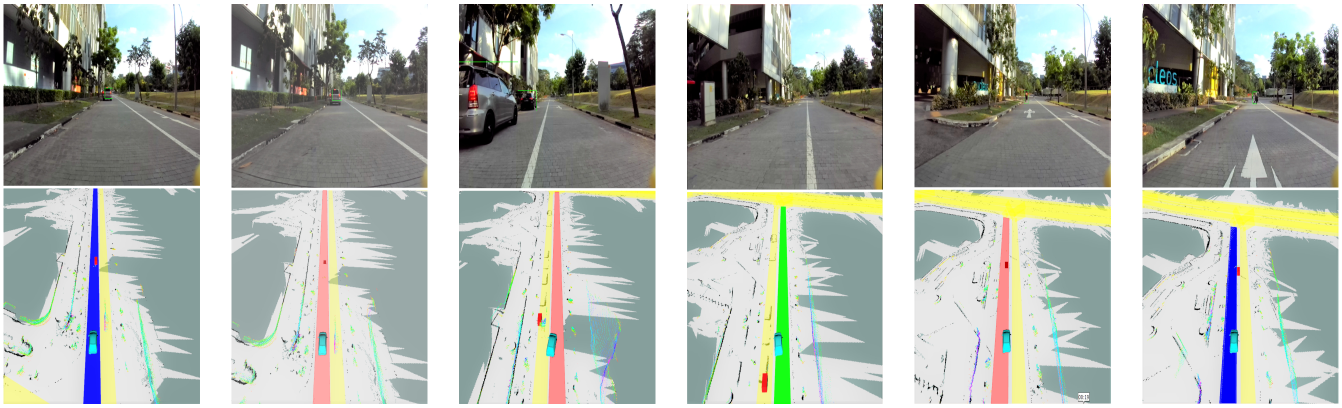

frame 1 *frame 2* *frame 3* *frame 4* *frame 5* *frame 6*

Fig. 8: Scenario 1 in real world. The ego-vehicle (cyan) encountered a slowly-moving vehicle and a pedestrian (red boxes). The actions suggested by the planner matched with those the driver actually executed.

between them. This can be learned more efficiently such that the planner can make informed decision with increased certainty. Also, the high-level actions provided by our planner are decoupled from the low-level control. It is possible that low-level controller cannot execute the high-level action given by the planner. We hope to address these concerns in our future work for autonomous driving in complex urban environments.

ACKNOWLEDGMENT

This research was supported by the National Research Foundation, Prime Minister's Office, Singapore, under its CREATE programme, Singapore-MIT Alliance for Research and Technology (SMART) Future Urban Mobility (FM) IRG.

REFERENCES

[1] 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.

[2] OpenStreetMap contributors. Planet dump retrieved from https://planet.osm.org . https://www.openstreetmap.org, 2017.

[3] Gabriel Svennerberg. *Beginning Google Maps API 3*. Apress, 2010.

[4] Sergio Casas, Wenjie Luo, and Raquel Urtasun. Intentnet: Learning to predict intention from raw sensor data. In *Conf. on Robot Learning*, pages 947–956, 2018.

[5] Brian Paden, Michal Čáp, Sze Zheng Yong, Dmitry Yershov, and Emilio Frazzoli. A survey of motion planning and control techniques for self-driving urban vehicles. *IEEE Trans. on Intelligent Vehicles*, 1:33–55, 2016.

[6] 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.

[7] 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.

[8] 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.

[9] Chenxi Ding, Wuhong Wang, Xiao Wang, and Martin Baumann. A neural network model for drivers lane-changing trajectory prediction in urban traffic flow. *Mathematical Problems in Engineering*, 2013, 2013.

[10] 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.

[11] 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.

[12] 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.

[13] 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.

[14] 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.

[15] Dhanvin Mehta, Gonzalo Ferrer, and Edwin Olson. Autonomous navigation in dynamic social environments using multi-policy decision making. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, pages 1190–1197. IEEE, 2016.

[16] Dorsa Sadigh, S Shankar Sastry, Sanjit A Seshia, and Anca Dragan. Information gathering actions over human internal state. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, pages 66–73. IEEE, 2016.

[17] Min Chen, Emilio Frazzoli, David Hsu, and Wee Sun Lee. POMDP-lite for robust robot planning under uncertainty. In *Proc. IEEE Int. Conf. on Robotics & Automation*, pages 5427–5433. IEEE, 2016.

[18] 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.

[19] Nan Ye, Adhiraj Somani, David Hsu, and Wee Sun Lee. DESPOT: Online POMDP planning with regularization. *J. of Artificial Intelligence Research*, 58:231–266, 2017.

[20] 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.

[21] Panpan Cai, Yuanfu Luo, David Hsu, and Wee Sun Lee. HyP-DESPOT: A hybrid parallel algorithm for online planning under uncertainty. *Proc. Robotics: Science & Systems*, 2018.

[22] Shashwat Verma, You Hong Eng, Hai Xun Kong, Hans Anderson, Malika Meghjani, Wei Kang Leong, Xiaotong Shen, Charles Zhang, Marcelo H. Ang Jr., and Daniela Rus. Vehicle detection, tracking and behavior analysis in urban driving environments using road context. In *Proc. IEEE Int. Conf. on Robotics & Automation*, 2018.

[23] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[24] 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.

[25] Carlos Lima Azevedo, Neeraj Milind Deshmukh, Balakumar Marimuthu, Simon Oh, Katarzyna Marczuk, Harold Soh, Kakali Basak, Tomer Toledo, Li-Shiuan Peh, and Moshe E Ben-Akiva. SimMobility short-term: An integrated microscopic mobility simulator. *Transportation Research Record*, 2622:13–23, 2017.