

SUMMIT: A Simulator for Urban Driving in Massive Mixed Traffic

Panpan Cai*, Yiyuan Lee*, Yuanfu Luo, David Hsu



(a) Singapore-Highway




(b) Magic-Roundabout



(c) Meskel-Intersection



(d) Singapore-Highway in SUMMIT



(e) Magic-Roundabout in SUMMIT



(f) Meskel-Intersection in SUMMIT

Fig. 1: Benchmark scenes in the real world and corresponding scenes in SUMMIT.

Abstract—Autonomous driving in an unregulated urban crowd is an outstanding challenge, especially, in the presence of many aggressive, high-speed traffic participants. This paper presents SUMMIT, a high-fidelity simulator that facilitates the development and testing of crowd-driving algorithms. By leveraging the open-source OpenStreetMap map database and a heterogeneous multi-agent motion prediction model developed in our earlier work, SUMMIT simulates dense, unregulated urban traffic for heterogeneous agents at any worldwide locations that OpenStreetMap supports. SUMMIT is built as an extension of CARLA and inherits from it the physics and visual realism for autonomous driving simulation. SUMMIT supports a wide range of applications, including perception, vehicle control and planning, and end-to-end learning. We provide a context-aware planner together with benchmark scenarios and show that SUMMIT generates complex, realistic traffic behaviors in challenging crowd-driving settings.

I. INTRODUCTION

The vision of using autonomous driving to improve the safety and convenience of our daily life is coming closer. However, driving in *unregulated*, crowded urban environments, like in uncontrolled roads or unsignalised intersections in less-developed countries (Fig. 1), remains an open problem. Human participants can be fairly aggressive in these scenarios. One may disregard or be unaware of traffic rules, leading to behaviors like close following, inappropriate overtaking, illegal turning and crossing, etc. The road condition can become highly chaotic when involving many

participants. Technical challenges for driving in unregulated urban crowds come from the complexity of both crowd behaviors and map environments. Traffic agents can be significantly different from each other. Cars, buses, bicycles, and motorcycles have different geometry, kinematics, and dynamics. Human participants also have different behavioral types - being conservative or aggressive, attentive or distracted, etc. In terms of the map environment, urban roads can have complex and versatile layouts: multi-lane roads, intersections, roundabouts, etc. Road structures significantly influence the motion of traffic agents and thus generate very different crowd behaviors in different locations. Such environments raise enormous difficulties for perception, control, planning, and decision-making of robot vehicles.

High-quality data for developing, training, and testing crowd-driving algorithms are, however, difficult and expensive to acquire due to the cost of devices, regulations and safety constraints. Although there are publicly available data sets like KITTI [1], BDD100K [2], Oxford RobotCar [3], etc., that provide real-world driving data with rich sensor inputs, these data are not *interactive*, i.e., one cannot model the reactions of exo-agents to the robot's decisions. Such data are, however, extremely important for robust planning and

*The authors contributed equally.

The authors are with School of Computing, National University of Singapore, 117417 Singapore. {caipp, leeyiyuan, yuanfu, dyhsu}@comp.nus.edu.sg.

TABLE I: Comparison between SUMMIT and existing driving simulators.

Simulator	Real-world Maps	Unregulated behaviors	Dense Traffic	Realistic Visuals & Sensors
SimMobilityST [5]	✓	×	✓	×
SUMO [6]	✓	×	✓	×
TORCS [7]	×	✓	×	✓
Apollo [8]	×	×	×	✓
Sim4CV [9]	×	×	×	✓
GTAV [10]	×	×	×	✓
CARLA [4]	×	×	×	✓
AutonoViSim [11]	×	✓	✓	✓
Force-based [12]	×	×	✓	✓
SUMMIT (ours)	✓	✓	✓	✓

learning. A promising source of interactive data are driving simulators that can offer a virtually unlimited amount of controllable scenarios. However, existing driving simulators do not capture the full complexity of unregulated urban crowds such as complex road structures and traffic behaviors, and are thus insufficient for testing or training robust driving algorithms. We aim to fill this gap.

We develop a new simulator, SUMMIT¹, that generates high-fidelity, interactive data for unregulated, dense urban traffic on complex real-world maps. SUMMIT uses real-world maps fetched from online sources to provide a virtually unlimited source of complex environments. Given arbitrary locations, the simulator automatically generates crowds of heterogeneous traffic agents with sophisticated, unregulated behaviors. The simulator leverages road contexts of real-world maps to guide the behaviours of traffic agents topologically and geometrically in order to construct realistic traffic conditions. We implemented SUMMIT based on CARLA [4] to leverage the high-fidelity physics, rendering, and sensors. Through a python-based API, SUMMIT reveals rich sensor data, semantic information, and road contexts to external algorithms, enabling the application in a wide range of fields such as perception, vehicle control and planning, end-to-end learning, etc. We provide both qualitative and quantitative results to show that SUMMIT can generate complex, realistic mixed traffic in real-world urban environments.

We further provide a context-aware planner, Context-POMDP², as a reference for future crowd-driving algorithms. The planner explicitly reasons about interactions among traffic agents under the uncertainty of human intentions and driving types. By further conditioning planning on available road contexts, Context-POMDP achieves safe and efficient driving in very challenging scenarios.

II. RELATED WORK

A. Existing Driving Simulators

Driving simulators have been extensively applied to boost the development of autonomous driving systems. Recent simulators (Table I) have brought realistic visuals and sensors,

but do not capture the complexities of urban environments and unregulated traffic behaviors.

Multi-car simulators like TORCS [7], [13], [14] focus on interactions between multiple robot-vehicles. These simulators suit the study of complex interactions between agents, but can hardly scale up to crowded urban scenes. CARLA [4], Sim4CV [9], and GTA [10] explicitly feature detailed physics modeling and realistic rendering for end-to-end learning. CARLA also provides a rich set of sensors such as cameras, Lidar, depth cameras, semantic segmentation, etc. However, these simulators rely on predefined maps, limiting the variety of environments. The simulated traffic also have relatively low density and simple rule-based behaviors. Another class of simulators [6], [5], [11], [12] feature traffic simulation and control in urban environments. Among them, SUMO [6] and SimMobilityST [5] support real-world maps but use simple rule-based behaviors, while another class [11], [12] apply more sophisticated motion models but are restricted to predefined maps. We aim to model the complexities in both urban maps and traffic behaviors in an automatic and unified framework.

B. Crowd Simulation Algorithms

Existing crowd simulation algorithms, e.g., social force and velocity obstacles, can in principle be applied to generate crowd behaviors in urban environments. Social force [15], [16], [17], [18] assume that traffic-agents are driven by attractive forces exerted by the destination and repulsive forces exerted by obstacles. Social force can simulate large crowds, but the quality of interactions are constrained by model simplicity. Velocity Obstacle (VO) [19] and Reciprocal Velocity Obstacle (RVO) [20], [21], [22] compute collision free motion by optimizing in the feasible velocity space. Variants such as GVO [23], NH-ORCA [24], B-ORCA [25], PORCA [26] explicitly handle non-holonomic traffic agents. Some variants model behavioral types of crowd agents such as patience [26] and attention [27]. A recent model, GAMMA [28], can simulate heterogeneous traffic agents with different geometry, kinematics, and behavioral types in a unified velocity-space framework. The behavior model in SUMMIT extends the framework of GAMMA to encode topological road contexts such as lanes and pedestrian sidewalks to closely represent real-world scenarios.

III. SUMMIT SIMULATOR

SUMMIT focuses on simulating complex unregulated behaviours of dense urban traffic in complex real-world maps. It is designed for generating high-fidelity interactive data to facilitate the development, training, and testing of crowd-driving algorithms. SUMMIT automatically generates massive mixed traffic using topological road contexts and optimization-based unregulated crowd behaviors. SUMMIT fetches real-world maps from the OpenStreetMap [29], and constructs two topological graphs: a lane network for vehicles, and a sidewalk network for pedestrians. These networks

¹Code available at <https://github.com/AdaCompNUS/summit>

²Code available at <https://github.com/AdaCompNUS/Context-POMDP>

³We only check-mark simulators explicitly featuring crowd behaviours.

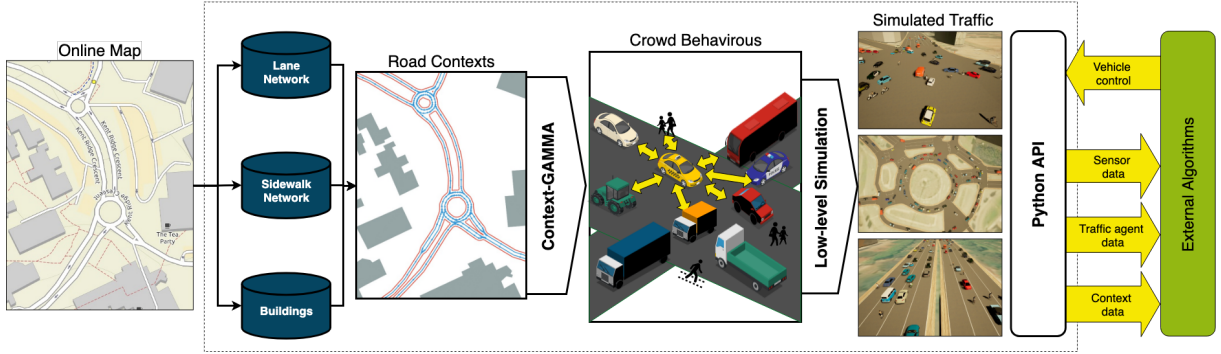


Fig. 2: An overview of SUMMIT that simulates massive mixed traffic at any location in the world.

form a representation of the road contexts. Then, our behavior model, Context-GAMMA, takes road contexts as input to guide the traffic behaviors geometrically and topologically. At the microscopic level, Context-GAMMA uses velocity-space optimization to generate sophisticated, unregulated crowd behaviors. The low-level structure of SUMMIT is based on CARLA, retaining its desirable features such as high-fidelity physics, realistic rendering, weather control, and rich sensors. Fig. 2 provides an overview of SUMMIT.

A. Representing Real-world Maps

1) *Lane Network*: A lane network in SUMMIT defines the connectivity of the road structure at the fidelity of individual lanes. The network consists of directed lane segments and connections between them. SUMMIT relies on SUMO [6] to automatically convert OSM maps to lane networks. The extensive suite of network editing tools provided by SUMO can also be leveraged to improve and customize maps. The lane network interface allows users to locate traffic agents on the lane network and retrieve connected lane segments. The interface closely follows CARLA’s waypoint interface, so that CARLA users can easily adapt to it.

2) *Sidewalk Network*: A sidewalk network in SUMMIT defines the behaviors of pedestrians, which usually walk along road edges and occasionally cross roads. The network contains sidewalks near road edges defined as poly-lines and connections between sidewalks defined as cross-able roads. The sidewalk poly-lines are extracted from the geometry of roads. Similar to the lane network, the sidewalk network interface allows users to locate pedestrians on the network and retrieve the opposite sidewalk for road-crossing.

3) *Occupancy Map Interface*: We additionally provide an occupancy map interface to expose drive-able regions for the ego-vehicle. An occupancy map is the top-down projection of the road geometry, aligned with the ego-vehicle’s location and heading direction. It can be used either for collision checking in control and planning algorithms or as bird-view input to neural networks.

4) *Landmarks*: SUMMIT also makes use of landmark data in OSM maps such as buildings and forests to provide structurally rich and realistic visuals. We additionally support randomization of the landmark textures to generate more

versatile visual inputs and enable techniques such as domain randomization [30].

B. Crowd behavior Modelling

SUMMIT uses Context-GAMMA, a context-aware crowd behavior model, to generate sophisticated interactive behaviors of traffic agents. Context-GAMMA extends GAMMA [28] to incorporate road contexts and models them as constraints in velocity space. The realism and accuracy of GAMMA has been validated in various real-world datasets. For completeness, we briefly introduce GAMMA, and present the extensions in Context-GAMMA.

GAMMA formulates the motion of traffic agents as constrained geometric optimization in velocity space. It assumes that each traffic agent optimizes its velocity based on the navigation goal, while being constrained by kinematic constraints (e.g. non-holonomic motion of car) and geometric constraints (collision avoidance with nearby agents). For a given agent A , let K_A represent the set of velocities that satisfy kinematic constraints and G_A^τ represent the set of velocities that satisfy geometric constraints for at least τ time. Then GAMMA selects for A a new velocity from their intersection:

$$v_A^{\text{new}} = \arg \min_{v \in G_A^\tau \cap K_A} \|v - v_A^{\text{pref}}\|, \quad (1)$$

where v_A^{pref} is A ’s preferred velocity computed from its goal. When computing K_A and G_A^τ , GAMMA also takes into account responsibility and attention of the agent to generate more human-like motions. We refer readers to [28] for more details of the construction of K_A and G_A^τ . Geometrically, K_A is a convex velocity set and G_A^τ is the intersection of velocity-space half planes. See Fig. 3a for an example of K_A and G_A^τ .

GAMMA handles heterogeneous traffic agents with different kinematics and geometry in a unified velocity-space framework. It has been proven experimentally to accurately predict behaviors of real-world traffic agents [28].

However, as GAMMA does not make explicit use of road contexts, it often fails to generate realistic simulation for complex urban roads. GAMMA agents can aggressively head towards their goals and be trapped by the complex road structure. In the real-world, road contexts can effectively

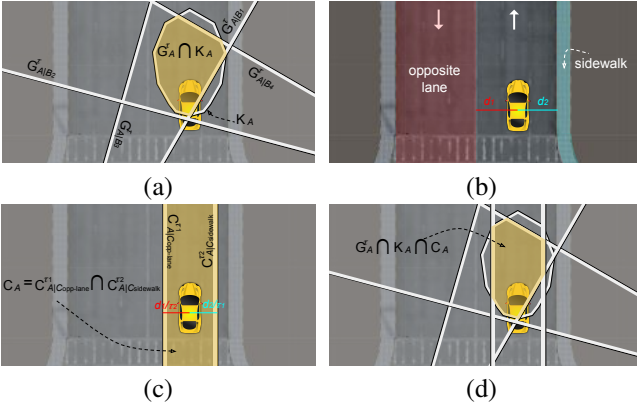


Fig. 3: Crowd behaviour modelling with Context-GAMMA: (a) the feasible velocity space (yellow) of an agent A as the intersection of kinematic constraints K_A and geometry constraints G_A^T . (b) A car constrained by the opposite lane and the sidewalk. (c) The corresponding contextual constraints $C_A^{T1|C_{opp-lane}}$ and $C_A^{T2|C_{sidewalk}}$. (d) The augmented feasible velocity space.

guide and constrain traffic agents' behaviors: vehicles tend to follow particular lanes when the road is clear and avoid to drive along the wrong direction. Moreover, as traffic agents are heterogeneous, they are affected by different sets of static obstacles: pedestrians consider sidewalks as open spaces, but vehicles consider them as obstacles.

Our new model, Context-GAMMA, provides a general way to embed road contexts as objectives and constraints in velocity space. Context-GAMMA extracts the preferred velocity v^{pref} of traffic agents from the lane and sidewalk networks. To diversify agent behaviors, it randomly selects a lane from all feasible lanes ahead of the agent, and points v^{pref} to a look-ahead waypoint along the selected lane.

Context-GAMMA can also model traffic rules by casting contextual constraints, e.g., no wrong-direction driving, into half planes in agents' velocity space, forcing them to select velocities complying with the road rule. Denote the intersection of all the contextual half planes for an agent A as C_A . Context-GAMMA optimizes the agent velocity in the augmented feasible velocity space $K_A \cap G_A^T \cap C_A$. Note that $K_A \cap G_A^T \cap C_A$ is also convex by construction and the objective function (1) is quadratic. Therefore, the optimization problem can be efficiently solved in linear time.

Fig. 3b and Fig. 3c show an example of contextual constraints. To prevent car A in Fig. 3b from driving to the opposite lane within a time window τ_1 , the lateral speed of the car should be constrained under d_1/τ_1 , where d_1 is the distance from the car to the opposite lane. This constraint forms a half-plane in the velocity space, $C_A^{T1|C_{opp-lane}}$, defined by a separation line parallel to the opposite lane with an offset of d_1/τ_1 from the origin. Any velocity in $C_A^{T1|C_{opp-lane}}$ would be feasible. Similarly, collision avoidance with the sidewalk can result in another half-plane $C_A^{T2|C_{sidewalk}}$ in the velocity space with an offset of d_2 from the origin. The intersection of the two half-planes forms A 's contextual constraint, C_A (Fig. 3c), which is further imposed on $K_A \cap G_A$ to form the feasible space of A (Fig. 3d).

C. Interfaces

The Python API of SUMMIT extends that of CARLA, exposing to external algorithms not only sensor data and agent states, but also road contexts like lane networks, sidewalk networks, and map occupancy grids. Algorithms can also send vehicle control back to the simulation including steering, acceleration, braking, reversing, etc.. SUMMIT thus enables a wide range of applications such as perception, sensor-based control, model-based reasoning, and end-to-end learning.

IV. CONTEXT-AWARE POMDP PLANNING

SUMMIT also offers an expert planner for autonomous driving. Planning for driving in an unregulated dense traffic is extremely challenging. The robot has to be smart enough to make efficient progress, instead of being "frozen" and stuck in the crowd. In the meantime, a highly dynamic and interactive crowd makes the task safety critical. Mistakes in planning can lead to severe or even fatal accidents.

The key to success is to explicitly model interactions among agents as well as the uncertainty on human inner states. Such sophisticated planning requires long-term reasoning in the belief-space, which brings combinatorial complexities. To make the problem tractable, we propose to condition planning on road contexts. We formulate driving as a context-aware POMDP, and solves it efficiently using online belief tree search [31]. We refer to this planner as Context-POMDP.

Context-POMDP conditions Monte Carlo simulations for the future on both human hidden states and road contexts. This is achieved by forward simulating exo-agents using Context-GAMMA. Context-POMDP consists of two components: a belief tracker that infers a joint belief over exo-agents' hidden states, and an online POMDP solver that computes an optimal driving action for the current belief.

A. Belief Tracking

The belief tracker maintains a joint belief over two dimensions of hidden states:

- The *intention* of the traffic agent: Let $U_i, i \in I_{exo}$ be the set of path candidates for the i th traffic agent extracted from the road contexts such as the lane network and the sidewalk network. This agent may take any of the path candidates in U_i as its actual intention.
- The *type* of a traffic agent: An agent can be either *distracted*, thus not interacting with the ego-vehicle, or be *attentive*, thus cooperatively avoid collision with the ego-vehicle.

The belief tracker is implemented as a factored histogram filter [32]. Each exo-agent is associated with a probability distribution over the set of possible hidden state values. At each time step, we use Context-GAMMA to generate mean motions for an agent conditioned on all possible hidden state values. By comparing the mean motions with the actual observed motion, we compute the likelihood of the observation and update the posterior belief using the Bayes rule.

B. Context-POMDP

The core of Context-POMDP is a context-aware POMDP model solved using a state-of-the-art belief tree search algorithm (HyP-DESPOT) [31]. We present the details of the model as follows.

1) *State and Observation Modelling*: A state in Context-POMDP includes both discrete-domain variables and continuous-domain variables:

- State of the ego-vehicle, $s_c = (x, y, \phi, \mu)$, including the position (x, y) , heading direction ϕ , and the intended driving path μ .
- Observable states of exo-agents, $\{s_i = (x, y, \vec{v})\}_{i \in I_{exo}}$, including the position (x, y) and the current velocity \vec{v} . I_{exo} defines the set of indices of exo-agents.
- Hidden states of exo-agents, $\{\theta_i = (t_i, \mu_i)\}_{i \in I_{exo}}$, including the type and the (sampled) intended path of the i th traffic agent.

We assume that the ego-vehicle can observe its own state and discretized values of the observable states of exo-agents. The hidden states of exo-agents can only be inferred and modelled with beliefs.

2) *Action Modelling*: The action space of the ego-vehicle consists of its steering angle and acceleration. Given the well-known exponential complexity of POMDP planning [33], Context-POMDP decouples the action space of the ego-vehicle to keep the branching factor of the planning problem within a tractable range. Concretely, we restrict the POMDP to compute the acceleration along the intended path, while the steering angle is generated using a pure-pursuit algorithm [34]. The action space contains three possible accelerations for each time step: $\{ACC, MAINTAIN, DEC\}$. The acceleration value for *ACC* and *DEC* is 3 m/s^2 and -3 m/s^2 , respectively. The maximum speed of the ego-vehicle is 6 m/s .

3) *Transition Modelling*: Context-POMDP predicts traffic agents' motion using the following set of models. Distracted traffic agents are assumed to track their intended path with the current speed. Attentive traffic agents also tend to follow the sampled path, but use Context-GAMMA to generate the actual local motion. The motion of all agents, including the ego-vehicle, are constrained by their kinematics, e.g., pedestrians are simulated using holonomic motion and car-like vehicles are simulated using bicycle models. To model stochastic transitions of exo-agents, their motion are perturbed by Gaussian noises on the displacement.

4) *Reward Modelling*: The reward function in Context-POMDP takes in to account safety, efficiency, and smoothness of driving. It assigns large penalties when the ego-vehicle collides with any exo-agent, uses a motion cost to penalize driving at low speed, and finally, penalizes frequent deceleration. Details of the reward function can be found in [26].

V. RESULTS

We want to answer the following questions in the experiments:

- Can SUMMIT simulate realistic dense traffic on complex maps?
- What are the benefits brought by SUMMIT over rule-based models commonly used in simulators?
- Can Context-POMDP drive a vehicle safely and efficiently in dense unregulated urban traffic?

We provide both qualitative and quantitative results to answer these questions.

A. Real-world Benchmark Scenarios

We designed three real-world benchmark scenarios to evaluate the performance of SUMMIT and the Context-POMDP planner.

- Singapore-Highway (Fig. 1a) A highway in Singapore with multiple lanes. Traffic agents try to drive as fast as possible and thus conduct overtaking frequently.
- Magic-Roundabout (Fig. 1b) A roundabout at Swindon, England with very complex layout. Traffic agents meet at the main roundabout and the accompanying intersections, having to coordinate with each other.
- Meskel-Intersection (Fig. 1c) A complex intersection at the Meskel square, Addis Abeba. Traffic agents come from different directions and encounter at the intersection, all of them driving aggressively.

Maps of the scenarios are fetched from online and imported into SUMMIT. It then simulates unregulated traffic on these maps using Context-GAMMA. All scenarios contain 120 heterogeneous traffic agents driving or walking in the region of interest, each conducting aggressive and unregulated behaviors. Once an agent moves out of the region, we replace it with new agents inside the region to maintain the high density of the traffic.

B. Simulation on Benchmark Scenarios

The realism and accuracy of our base model GAMMA have been validated in [28]. In this section, we perform a qualitative study on the simulation performance of Context-GAMMA. Fig. 1(d-f) shows qualitative simulation results on the benchmark scenarios. Comparison with the real-world scenarios shows that the simulated traffic closely represent the reality. More simulation results can be found in the accompanying video or via <https://youtu.be/dNiR0z2dROg>.

C. Comparison with Rule-based Simulation

We compare Context-GAMMA with rule-based behaviours commonly used in simulators to demonstrate the sophistication of behaviours. Particularly, we compare Context-GAMMA with a reactive model that moves agents along lane center-curves and uses time-to-collision (TTC) [35] to calculate the agents' speeds. Performances of the two models are measured using the average speed of traffic agents and a congestion factor defined as the percentage of agents being jammed in the crowd, which are removed after remaining stationary for a substantial amount of time. These measures indicate how smart a behaviour model is




Fig. 4: Performance profile of Context-GAMMA and TTC on the Meskel-Intersection: (a) average speed of traffic agents; (b) congestion factor of the traffic.

TABLE II: Time performance and scalability of Context-GAMMA for simulating different number of agents on a laptop with an i7-9750H CPU and an RTX 2060 GPU.

Number of Agents	150	200	250	300	350	400
Frequency (Hz)	28.4	20.2	16.6	12.7	9.4	7.1
Update Time (ms)	35.3	49.6	60.2	78.9	106.9	141.8

at collision avoidance and driving efficiency, which human drivers are adept at.

Fig. 4 shows a detailed profile of the agent speeds and the congestion factors for different types of agents against the simulation time. Context-GAMMA generates faster and smoother traffic than the TTC in all benchmark scenarios throughout 20 minutes of simulation. The congestion factor of the TTC-controlled traffic grows quickly with the simulation time, indicating that agents fail to coordinate with each other. In contrast, Context-GAMMA consistently produces higher agent speeds and low congestion factors for all agent types. This is because Context-GAMMA explicitly models cooperation between agents and provides an optimal collision avoidance motion using both steering and acceleration.

D. Efficiency and Scalability of Simulation

Efficiency tests in Table II show that Context-GAMMA scales well with the density of the crowd. The simulation runs at high rates even when modelling up to 400 agents, and the growth of computation time is almost linear until the map saturates with agents.

E. Driving Performance of the Context-POMDP Planner

We now validate the performance of the Context-POMDP planner by comparing its driving performance with local-

TABLE III: Comparison on the driving performance of driving algorithms. Roll-out and Context-POMDP run at 3HZ. Context-GAMMA runs at 20 HZ. A “step” is counted as 1/3s.

	Collision / step	Avg. Speed (m/s)	Dec. / step
Roll-out	0.00095	2.1	0.15
Context-GAMMA	0.002	5.53	0.12
Context-POMDP	0.00069	4.53	0.16

collision avoidance and simple planning baselines. For local collision avoidance, we directly use GAMMA to control the ego-vehicle; For simple planning, we use a roll-out algorithm that plans for optimal action by casting multiple roll-outs using a default policy. The default policy applies the following rules: accelerate the ego-vehicle when exo-agents in front are far way (> 4 m away), maintain half-speed when they are in caution range ($2 \sim 4$ m away), and decelerate when they are close-by (< 2 m away).

Table III provides measurements of the collision rate per step, the average vehicle speed, and the frequency of deceleration when driving the ego-vehicle using Context-POMDP, Context-GAMMA, and Roll-out. In summary, the simple planner drives over-conservatively, while local collision avoidance drives over-aggressively. Sophisticated planning using Context-POMDP offers a successful trade-off between aggressiveness and conservatives.

Compared to Roll-out that can barely move in the crowd, Context-POMDP can drive the vehicle through the crowd at significantly higher speed while remaining safe and smooth. Compared to Context-GAMMA which drives aggressively but leads to high collision rate, Context-POMDP achieves similar driving speed with much safer behaviours; We thus conclude that sophisticated long-term planning is important for crowd-driving, and Context-POMDP ensures safe, efficient, and smooth driving.

VI. CONCLUSION

We presented SUMMIT, a simulator for generating high-fidelity interactive data for developing, training, and testing crowd-driving algorithms. The simulator uses online maps to automatically construct unregulated dense traffic at any location of the world. By integrating topological road contexts with an optimization-based crowd behavior model, SUMMIT can generate complex and realistic crowds that closely represent unregulated traffic in the real-world. We also provided Context-POMDP as a reference planning algorithm for future development. We envision that SUMMIT will support a wide range of applications such as perception, control, planning, and learning for driving in unregulated dense urban traffic.

VII. ACKNOWLEDGEMENT

This research is supported in part by the Singapore MoE AcRF Tier 2 grant MOE2016-T2-2-068 and the Singapore A*STAR Undergraduate Scholarship SE/AUS/15/016.

REFERENCES

- [1] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [2] F. Yu, W. Xian, Y. Chen, F. Liu, M. Liao, V. Madhavan, and T. Darrell, "Bdd100k: A diverse driving video database with scalable annotation tooling," *arXiv preprint arXiv:1805.04687*, 2018.
- [3] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 year, 1000 km: The oxford robotcar dataset," *The International Journal of Robotics Research*, vol. 36, no. 1, pp. 3–15, 2017.
- [4] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," *arXiv preprint arXiv:1711.03938*, 2017.
- [5] C. L. Azevedo, N. M. Deshmukh, B. Marimuthu, S. Oh, K. Marczuk, H. Soh, K. Basak, T. Toledo, L.-S. Peh, and M. E. Ben-Akiva, "Sim-mobility short-term: An integrated microscopic mobility simulator," *Transportation Research Record*, vol. 2622, no. 1, pp. 13–23, 2017.
- [6] P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wießner, "Microscopic traffic simulation using sumo," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pp. 2575–2582, IEEE, 2018.
- [7] B. Wymann, E. Espié, C. Guionneau, C. Dimitrakakis, R. Coulom, and A. Sumner, "Torcs, the open racing car simulator," *Software available at <http://torcs.sourceforge.net>*, vol. 4, no. 6, 2000.
- [8] "Apollo simulation," *<http://apollo.auto/platform/simulation.html>*, 2018.
- [9] M. Müller, V. Casser, J. Lahoud, N. Smith, and B. Ghanem, "Sim4cv: A photo-realistic simulator for computer vision applications," *International Journal of Computer Vision*, vol. 126, no. 9, pp. 902–919, 2018.
- [10] S. R. Richter, Z. Hayder, and V. Koltun, "Playing for benchmarks," in *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2213–2222, 2017.
- [11] A. Best, S. Narang, L. Pasqualin, D. Barber, and D. Manocha, "Autonovi-sim: Autonomous vehicle simulation platform with weather, sensing, and traffic control," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 1048–1056, 2018.
- [12] Q. Chao, X. Jin, H.-W. Huang, S. Foong, L.-F. Yu, and S.-K. Yeung, "Force-based heterogeneous traffic simulation for autonomous vehicle testing," in *2019 International Conference on Robotics and Automation (ICRA)*, pp. 8298–8304, IEEE, 2019.
- [13] H. Zhao, A. Cui, S. A. Cullen, B. Paden, M. Laskey, and K. Goldberg, "Fluids: A first-order local urban intersection driving simulator," in *CASE.*, IEEE, 2018.
- [14] M. Naumann, F. Poggenhans, M. Lauer, and C. Stiller, "Coincar-sim: An open-source simulation framework for cooperatively interacting automobiles," in *2018 IEEE Intelligent Vehicles Symposium (IV)*, pp. 1–6, IEEE, 2018.
- [15] D. Helbing and P. Molnar, "Social force model for pedestrian dynamics," *Physical review E*, vol. 51, p. 4282, 1995.
- [16] R. Löhner, "On the modeling of pedestrian motion," *Applied Mathematical Modelling*, vol. 34, pp. 366–382, 2010.
- [17] G. Ferrer, A. Garrell, and A. Sanfeliu, "Robot companion: A social-force based approach with human awareness-navigation in crowded environments," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, 2013.
- [18] S. Pellegrini, A. Ess, K. Schindler, and L. Van Gool, "You'll never walk alone: Modeling social behavior for multi-target tracking," in *Computer Vision, 2009 IEEE 12th International Conference on*, pp. 261–268, IEEE, 2009.
- [19] P. Fiorini and Z. Shiller, "Motion planning in dynamic environments using velocity obstacles," *Int. J. Robotics Research*, vol. 17, pp. 760–772, 1998.
- [20] J. Van den Berg, M. Lin, and D. Manocha, "Reciprocal velocity obstacles for real-time multi-agent navigation," in *Proc. IEEE Int. Conf. on Robotics & Automation*, 2008.
- [21] J. Van Den Berg, S. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Proc. Int. Symp. on Robotics Research*, 2009.
- [22] J. Snape, J. Van Den Berg, S. J. Guy, and D. Manocha, "The hybrid reciprocal velocity obstacle," *IEEE Transactions on Robotics*, vol. 27, pp. 696–706, 2011.
- [23] D. Wilkie, J. Van Den Berg, and D. Manocha, "Generalized velocity obstacles," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, pp. 5573–5578, IEEE, 2009.
- [24] J. Alonso-Mora, A. Breitenmoser, M. Ruffli, P. Beardsley, and R. Siegwart, "Optimal reciprocal collision avoidance for multiple non-holonomic robots," in *Distributed Autonomous Robotic Systems*, pp. 203–216, Springer, 2013.
- [25] J. Alonso-Mora, A. Breitenmoser, P. Beardsley, and R. Siegwart, "Reciprocal collision avoidance for multiple car-like robots," in *Proc. IEEE Int. Conf. on Robotics & Automation*, 2012.
- [26] Y. Luo, P. Cai, A. Bera, D. Hsu, W. S. Lee, and D. Manocha, "Porca: Modeling and planning for autonomous driving among many pedestrians," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3418–3425, 2018.
- [27] E. Cheung, A. Bera, and D. Manocha, "Efficient and safe vehicle navigation based on driver behavior classification," in *Proc. IEEE Conf. on Computer Vision & Pattern Recognition*, pp. 1024–1031, 2018.
- [28] Y. Luo, P. Cai, D. Hsu, and W. S. Lee, "GAMMA: A general agent motion prediction model for autonomous driving," *arXiv preprint arXiv:1906.01566*, 2019.
- [29] OpenStreetMap contributors, "Planet dump retrieved from <https://planet.osm.org>," <https://www.openstreetmap.org>, 2017.
- [30] M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, et al., "Learning dexterous in-hand manipulation," *arXiv preprint arXiv:1808.00177*, 2018.
- [31] P. Cai, Y. Luo, D. Hsu, and W. S. Lee, "HyP-DESPOT: A hybrid parallel algorithm for online planning under uncertainty," in *Proc. Robotics: Science & Systems*, 2018.
- [32] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT press, 2005.
- [33] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, "Planning and acting in partially observable stochastic domains," *Artificial Intelligence*, vol. 101, pp. 99 – 134, 1998.
- [34] R. C. Coulter, "Implementation of the pure pursuit path tracking algorithm," tech. rep., Carnegie-Mellon UNIV Pittsburgh PA Robotics INST, 1992.
- [35] M. Meghiani, Y. Luo, Q. H. Ho, P. Cai, S. Verma, D. Rus, and D. Hsu, "Context and intention aware planning for urban driving," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, 2019.