

Traffic-Aware Autonomous Driving with Differentiable Traffic Simulation

Laura Zheng, Sanghyun Son, and Ming C. Lin
gamma.umd.edu/trafficdriving

Abstract—While there have been advancements in autonomous driving control and traffic simulation, there have been little to no works exploring their unification with deep learning. Works in both areas seem to focus on entirely different exclusive problems, yet traffic and driving are inherently related in the real world. In this paper, we present Traffic-Aware Autonomous Driving (TrAAD), a generalizable distillation-style method for traffic-informed imitation learning that directly optimizes for faster traffic flow and lower energy consumption. TrAAD focuses on the supervision of speed control in imitation learning systems, as most driving research focuses on perception and steering. Moreover, our method addresses the lack of co-simulation between traffic and driving simulators and provides a basis for directly involving traffic simulation with autonomous driving in future work. Our results show that, with information from traffic simulation involved in the supervision of imitation learning methods, an autonomous vehicle can learn how to accelerate in a fashion that is beneficial for traffic flow and overall energy consumption for all nearby vehicles.

I. INTRODUCTION

The ideal autonomous vehicle (AV) should be able to minimize the travel time of a route, maximize the energy efficiency of the vehicle, and provide a smooth and safe experience for the riders. These objectives are not only important to the passenger’s experience, but also for greater global benefits. Annually in the US, traffic congestion accounts for 29 billion dollars in costs [1], transportation accounts for 27% of carbon emissions [2], and motor vehicle accidents are the leading cause of unnatural death as of 2022 [3]. From a learning perspective, the extent an AV can improve on these objectives is affected by the environment around it, *e.g.* the surrounding vehicle traffic and the road networks. Similarly to other multi-agent problems, improving traffic flow and energy efficiency for all vehicles in the system can also benefit each vehicle’s individual objectives. Inversely, the vehicle’s motion also impacts the traffic around it. An AV influences the flow of human-driven traffic by acting as a pace car to those behind it. A “pace car” is commonly used in racing to control the speeds of competing vehicles; this same notion can be applied to traffic control in the physical world. In a driving policy, we can emulate this effect by defining an individual vehicle’s objective in terms of global metrics.

As long as a model of the traffic and its dynamics on a road network are accessible, an autonomous driving policy can directly optimize for *traffic flow, energy efficiency, and smooth acceleration*. Decades of traffic engineering research present

The authors are with Department of Computer Science, University of Maryland at College Park, MD, U.S.A. E-mail: {lyzheng,shh1295,lin}@umd.edu

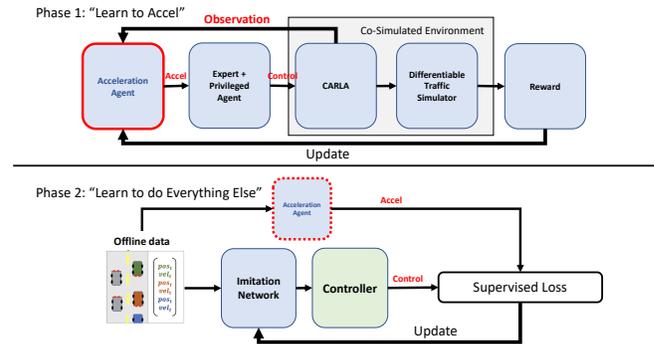


Fig. 1: Training for Traffic-aware Autonomous Driving.

Our method can be adapted to most existing imitation learning frameworks for driving by just adding an additional phase of training, where we isolate “Learn to Accelerate”. An agent whose action space only spans possible acceleration actions navigates through a co-simulated environment and is rewarded when it improves overall traffic flow and fuel consumption. In Phase 2, where we “Learn to do Everything Else”, the acceleration agent is frozen and supervises speed control of an imitation learning agent via distillation.

sophisticated mathematical equations modeling traffic, including car-following models [4]–[7], traffic flow models [8]–[11] and theory [12]. These mathematical models, though often too simplified to account for the uncertainty of driver behaviors, are computationally efficient, differentiable, and do not require data. One possibility is to model the traffic environment with a neural network, as in some of the latest work [13]–[16]. Although these methods are more accurate than ODEs, they also require large amounts of data to generalize and are subject to problems associated with deep learning, such as bias and distributional shifts. In this paper, we couple a learning-based traffic control algorithm with *differentiable* ODE traffic models. We use the gradients of forward traffic dynamics to guide the learning of a driving policy, so long as the policy has access to traffic information. Minimal traffic information involves simulation states of position and velocity over time, which is further explained in Section III-A.

We present a generalizable method for *traffic-aware autonomous driving* (TrAAD) which takes advantage of *differentiable traffic simulation*. By coupling a driving environment with traffic simulation, the driving policy can retrieve traffic information during training and learn behaviors that are both beneficial to individual and global goals.

In TrAAD, we add a phase of training in addition to traditional imitation learning for driving, where the vehicle “learns to accelerate”. This phase involves maximizing the overall traffic flow of a vehicle’s local lane, minimizing the fuel consumption of all vehicles, and discouraging the acceleration actions from being too jerky. Because our method supervises acceleration via distillation, it is generalizable to nearly any standard imitation learning framework, regardless of architecture or design. Our results show that our method, when implemented on top of existing state-of-the-art driving frameworks, improves traffic flow, minimizes energy consumption for the AV, and enhances the passenger’s ride experience.

In summary, we present the following key contributions:

- 1) A simulated traffic-annotated driving dataset for imitation learning for self-driving cars;
- 2) Use of gradients from differentiable traffic simulation to improve sample efficiency for autonomous vehicles;
- 3) A generalizable method for traffic-aware autonomous driving, which learns to control the vehicle via rewards based on societal traffic-based objectives.

Additional results, materials, code, datasets, and information can be found on our project website.

II. RELATED WORKS

A. Autonomous Driving with Traffic Information

Zhu et al. recently proposed a method for safe, efficient, and comfortable velocity control using RL [17]. Similarly to one of our objectives, they aim to learn acceleration behavior that exceeds the safety and comfort of human expert drivers. One major difference is that our work complements existing end-to-end autonomous driving systems with multi-modal sensor data, and learned acceleration behavior cooperates with learned control behavior from imitation learning rather than learning acceleration in a pure traffic simulation setting. In addition, our objective is to directly optimize on an entire traffic state, not just the objectives for the autonomous vehicle itself. The reward objectives of [17] are also inferred from a partially-observed point of view. Other works have considered learning driving behavior with passenger comfort and safety in mind, but many do not directly involve traffic state information beyond partially-observed settings [18]–[20]. Wegener et al. present a method for energy-efficient urban driving via RL [21] in a partially-observed setting purely in traffic simulation, however, does not address integration with current works for more complex vehicle control. In short, our method addresses a broader method for learning a policy beneficial to both individual and societal traffic objectives, while can be easily integrated into existing state-of-the-art end-to-end driving control methods.

B. Differentiable Microscopic Traffic Simulation

While differentiable physics simulation has been gaining popularity in recent years, differentiable traffic simulation is under-explored, especially in applications for autonomous driving. In 2021, Andelfinger first introduced the potential of

differentiable agent-based traffic simulation, as well as techniques to address discontinuities of control flow [22]. In his work, Andelfinger highlights continuous solutions for discrete or discontinuous operations such as conditional branching, iteration, time-dependent behavior, or stochasticity in forward simulation, ultimately enabling the use of automatic differentiation (autodiff) libraries for applications such as traffic light control. One key difference between our work and [22] is that our implementation of differentiable simulation accounts for learning agents acting independently from agents following a car-following model, and is compatible with existing learning frameworks. In addition, we optimize traffic-related learning by defining analytical gradients rather than relying solely on auto-differentiation. Most recently, Son et al. proposed a novel differentiable hybrid traffic simulator that computes gradients for both macroscopic, or fluid-like, representations and agent-based microscopic representations, as well as the transitions between them [23]. In our work, we focus solely on microscopic agent-based simulation to maintain relevance to autonomous driving frameworks.

C. Deep Learning with Traffic Simulation

Deep reinforcement learning has been used to address futuristic and complex problems for control of autonomous vehicles in traffic. One survey on Deep RL for motion planning for autonomous vehicles by Aradi [24] delineates challenges facing the application of DRL to traffic problems, one of which is the long and potentially unsuccessful learning process. This has been addressed in several ways through curriculum learning [25]–[27], adversarial learning [28], [29], or model-based action choice. In our work, we address this issue via sample enhancement for on-policy deep reinforcement learning. With differentiable traffic simulation and access to gradients of reward with respect to policy action, we can artificially generate “helpful” samples during learning with respect to reward. “FLOW” by Wu et al. [30] presents a deep reinforcement learning (DRL) benchmarking framework, built on the popular microscopic traffic simulator SUMO [31]. Wu et al. provide motivation for integrating traffic dynamics into autonomous driving objectives with DRL, defining the problem/task as “mixed autonomy”. Novel objectives for driving include reducing congestion, carbon emissions, and other societal costs; these are all in futuristic anticipation of mixed autonomy traffic. Based on FLOW, Vinitsky et al. published a series of benchmarks highlighting 4 main scenarios regarding traffic light control, bottleneck throughput, optimizing intersection capacity, and controlling merge on-ramp shock waves [32]. We extend the environments from FLOW’s DRL framework to be *differentiable* and show benchmark results for enhanced DRL algorithms utilizing traffic flow gradients for optimization.

III. BACKGROUND

A. Simulation-related Notation and Definitions

To integrate traffic simulation into learning and optimization frameworks for autonomous driving, we need differentiable forward simulation. Agent-based traffic simulation is

governed by ordinary differential equations (ODEs) known as car-following models. These ODEs describe the position and velocity of each individual agent over time. In the context of traffic simulation, the agent directly in front is considered in calculating the position and velocity. Since the position and velocity of each individual agent can be used to describe the current state of the entire simulation, we define the *agent state* q_n^t and the *simulation state* s^t at time step t for vehicle n out of N total vehicles to be

$$q_n^t = \begin{bmatrix} x_n^t \\ v_n^t \end{bmatrix}^\top, s^t = [q_1^t \dots q_N^t]^\top \quad (1)$$

B. The Intelligent Driver Model (IDM)

We model traffic dynamics of human drivers in our experiments with the Intelligent Driver Model (IDM) [4]. IDM describes how the position and velocity of each individual vehicle change over time. Intuitively, vehicles following the IDM will increase acceleration when the headway to the vehicle in front is large, and decelerate comfortably based on a maximum deceleration parameter. Unlike previous works in traffic simulation, we analytically derive the gradient of forward simulation to optimize the computation of backpropagation, rather than relying on auto-differentiation. We leave the formal definition, speedup analysis, variable definitions, and the full derivation of the analytical IDM Jacobian in the appendix, which can be found on the project website.

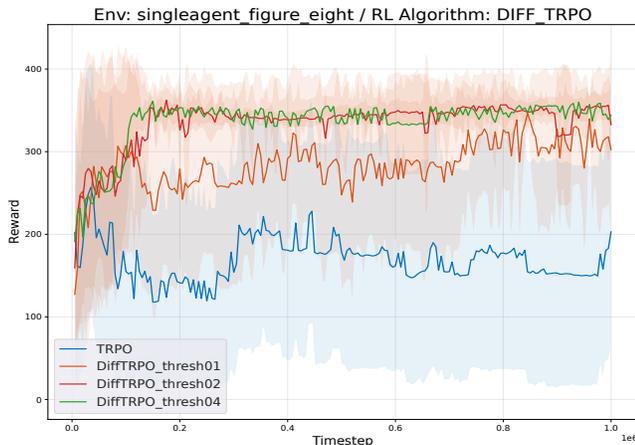


Fig. 2: **Effect of Gradients for On-Policy Learning.** We show the training curve for our sample efficiency enhancement method DiffTRPO versus the on-policy algorithm TRPO [33], as well as various experiments varying the perturbation threshold parameter δ discussed in section IV-C. Curves are averaged over 10 runs each to account for stochasticity. A perturbation threshold δ of 0.2 and 0.4, shown in red and green respectively, yields *faster learning* and *higher reward* than the baseline TRPO in blue. Additional results for PPO are shown in supplemental materials on the project website.

IV. TRAFFIC-AWARE AUTONOMOUS DRIVING

Most, if not all self-driving models focus on replicating expert human behavior as ground truth. Additionally, FLOW [30] introduces the vision of mixed autonomy, where AVs can drive in a way that benefits all members of the

simulation. This behavior is not necessarily observed in the real world or known to humans, and thus ground truth for optimal driving for societal good is not known. Currently, there is no existing research on integrating such behavior into single-vehicle driving control. We present TrAAD as a practical application of the ideas presented by FLOW to current imitative autonomous driving methods.

Differentiable traffic simulation can be directly applied to supervised learning due to the availability of simulation gradients for backpropagation. In our method, we use traffic simulation gradients to improve sample efficiency in on-policy RL.

Dataset Collection. We collect our own driving dataset in the same format as “Learning by Cheating” (LBC) [34]. Driving data is collected by an expert driver, CARLA autopilot, in CARLA Simulator [35] at a rate of 2 Hz, or every 10 frames at 20 frames per second. The expert driver is *not* a learned agent and uses route waypoints and simulation information to drive optimally. Training data comprises 50 routes over towns 1, 3, 5, and 6, while test data comprises 26 routes. These towns describe a small suburban town, a large complex city, a square-grid city, and a highway environment respectively. Sensors include a top-down birds-eye segmentation map and three RGB cameras (left, right, and center dashcam views). Annotations include vehicle position, vehicle control commands, and traffic state of the current vehicle lane in the same format as Equation 1. The dataset and the code to generate it can be found on our project website.

Hardware specs and Software. All experiments are trained on a single NVIDIA RTX A5000 GPU, Intel(R) Xeon(R) W-2255 CPU (20 cores), and 16 GB RAM, with the exception of experiments on TransFuser, which was trained on eight A5000 GPUs in parallel.

A. Phase 1: Learning Acceleration with Online RL

The goal of this phase is to learn acceleration behavior with on-policy reinforcement learning (RL) algorithms, given that other controls such as steering are optimal via the expert driver. For our experiments, we use PPO as well as the gradient-enhanced PPO later described in Section IV-C, which improves sample efficiency by taking advantage of simulation gradients.

The action space of this phase is continuous and defined by a maximum acceleration α_{max} and minimum acceleration α_{min} , which represent units of meters per second squared (m/s^2). The observation space is multi-modal and depends on the inputs of the imitation learning framework. For LBC, we collect a top-down segmentation map M plus the traffic state s given by the simulator. To account for traffic laws and discourage blind acceleration, the predicted acceleration is only used when the expert determines a non-braking state. Reward functions from FLOW [30] can now be used with auto-differentiation since traffic simulation is differentiable.

In our experiments, we use a weighted sum of two global objectives, R_{vel} and R_{mpg} , plus an individual jerkiness constraint for the AV. The first term relates to the average

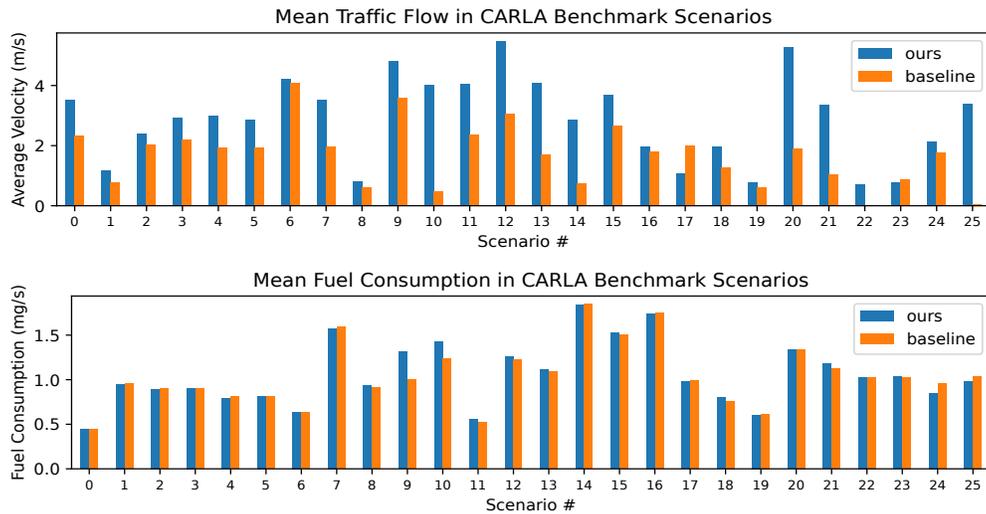


Fig. 3: **Comparison of overall traffic flow and fuel consumption between our method and baseline, for each CARLA test scenario.** Our model consistently improves overall traffic flow over measured baseline traffic flow in nearly all scenarios, showing that our single-vehicle control can influence the traffic flow around it. Our method is also able to maintain similar, if not better, fuel consumption—despite a direct relationship between increased flow and increased fuel usage.

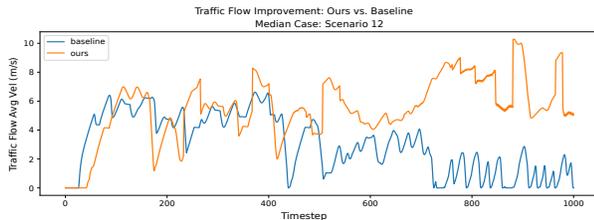


Fig. 4: **Comparison of Traffic Flow Over Time in the median case.** We visualize the average velocity of traffic flow over time in the learned agent’s lane for TrAAD (in orange) versus the baseline (in blue). TrAAD achieves +64.27% on average.

velocity of the traffic state, the second term is a miles-per-gallon metric that relates to energy efficiency, and the third term is an L2 constraint that discourages the next predicted actions to be too far from the last predicted action for smoother acceleration. Intuitively, a passenger will feel safer if a driver slowly presses down the gas pedal to accelerate. The third term can also be viewed as a form of time-dependent regularization on sequences of actions. The reward function R_{comb} is described below, where g_t^i is gallons of fuel consumed per second and v_t^i represents the velocity in meters per second by vehicle i at timestep t after action a_t :

$$R_{vel}(s_t) = \frac{1}{N} \sum_{i=1}^N v_t^i$$

$$R_{mpg}(s_t) = \frac{1}{1609N} \sum_{i=1}^N \frac{v_t^i}{g_t^i}$$

$$R_{comb}(s_t) = \alpha R_{vel}(s_t) + \beta R_{mpg}(s_t) - \lambda \|a_t - a_{t-1}\|_2$$

B. Phase 2: Learn to do Everything Else

This phase involves the integration of the learned Phase 1 acceleration model into an autonomous driving framework for single-vehicle control. In our experiments, we use Learning

by Cheating (LBC) [34] as the backbone for Phase 2. We first train the privileged agent that has access to a top-down ground truth map of the environment $M \in \{0, 1\}^{W \times H \times 7}$. The privileged agent learns to drive an AV in a fully-supervised manner, similar to LBC. Our Phase 1 model, illustrated in the top half of Figure 1, directly provides ground truth for the supervision of throttle control in this step. Instead of learning the throttle based on ground-truth labels provided by the expert agent during data collection, the Phase 1 acceleration agent supervises the throttle command. Each sample is also annotated with the local traffic lane state, which is also provided to the Phase 1 model.

Intuitively, the Phase 1 model is teaching the Phase 2 model how to accelerate; all other controls are learned through the original LBC pipeline. This includes the non-cheating “student” model, which learns from limited single-view information. Supervision of both speed and steering of the student model is done by the trained Phase 2 privileged model, which should have learned optimized acceleration behavior from the Phase 1 acceleration model.

This implementation design provides two main advantages: (1) Phases 1 and 2 are modular, and thus can be re-used to train multiple Phase 2 models; (2) Since traffic information transitions between models as a form of knowledge distillation, TrAAD is able to complement prior works rather than competing against them. The involvement of traffic information in our method can be thought of as knowledge distillation.

C. Improving Sample Efficiency with Traffic Gradients

Since agent-based forward simulation of traffic can be represented by ODEs, gradients of the traffic state at timestep t with respect to that at timestep $t - 1$ can be analytically derived. We define these gradients as “traffic gradients”. We can use traffic gradients in enhancing the sample efficiency of

TABLE I: **Driving Performance Benchmark.** Arrows denote the direction of improvement. We benchmark two baselines, LBC [34] and TransFuser [36], [37], and evaluate our method on CARLA Leaderboard Test Scenarios and Longest6 benchmarks from [37] respectively. Our method, TrAAD, is able to complement and enhance the driving performance of existing methods solely through the distillation of acceleration behaviors, with no modifications to design governing the learning of other controls. For TransFuser, ours is able to improve the RouteCompletion %, yet incurring negligible infractions. See Section V for details.

Method	DrivingScore \uparrow	RouteCompletion \uparrow	PedCollisions \downarrow	VehCollisions \downarrow	OtherCollisions \downarrow	Timeouts \downarrow
LBC	2.876	6.292	0.0	4.148	25.715	0.830
LBC+Ours	6.256	19.487	0.0	3.706	12.766	0.137
TransFuser	33.908	77.657	0.0240	3.366	0.168	0.144
TransFuser+Ours	31.291	79.964	0.0223	2.604	0.200	0.044

TABLE II: **Ablation on Impact of Each Factor in Reward Function.** Arrows denote the direction of improvement. As expected, the model omitting the fuel consumption term is able to freely accelerate without worrying about fuel efficiency. Despite fuel efficiency and traffic flow being inversely related, our model is able to achieve a middle ground between ablation models for the best of both worlds. More details on this experiment can be found in Section V.

Method	TrafficFlow \uparrow	FuelCons \downarrow	Jerk \downarrow	DrivingScore \uparrow	RtCompletion \uparrow	Infractions \downarrow	Collisions $_{Veh} \downarrow$	Collisions $_{Other} \downarrow$
Baseline	1.713	1.0399	2.2862e-3	4.021	6.184	0.302	2.879	10.557
NoVelocity	2.637	1.1010	1.7311e-3	4.197	6.871	0.350	2.309	15.701
NoFuel	2.679	1.0866	0.9465e-3	5.507	7.622	0.25	0.819	11.05
NoJerk	2.290	1.0248	0.9466e-3	3.900	5.263	0.229	0.608	9.731
Ours	2.428	1.0594	0.9465e-3	4.094	6.271	0.320	2.801	9.801

the common on-policy RL algorithms, such as PPO [38] or TRPO [39]. Inspired by the sampling enhancement scheme of [40], which applies the scheme to a model-based RL algorithm, here we present a method that is applicable to the other general on-policy RL algorithms and show its efficacy in the Section V

Specifically, PPO and TRPO are based on evaluating the perturbed policy $\tilde{\pi}$ with the experience from the original policy π [41]. This is possible because the expected reward of the policy $\tilde{\pi}$ can be approximated with the information from the original policy up to the first order. Our strategy is based on manipulating the collected set of experience with our gradients so that we can maximize the efficiency of each update.

Let us denote a single experience unit as $(s_n, a_n, r_n, s_{n+1}, \frac{\partial r_n}{\partial a_n}, \frac{\partial s_{n+1}}{\partial a_n})$, where s_n , a_n , and r_n refer to the state, action, and reward at the time step n , and $\frac{\partial r_n}{\partial a_n}$ and $\frac{\partial s_{n+1}}{\partial a_n}$ stand for the gradient of the reward and the next state with respect to the action. These gradient terms come from our differentiable traffic simulator. Then we can perturb (a_n, r_n, s_{n+1}) with the gradient terms as follows:

$$\tilde{a}_n = a_n + \epsilon, \quad \tilde{r}_n = r_n + \epsilon \cdot \frac{\partial r_n}{\partial a_n}, \quad \tilde{s}_{n+1} = s_n + \epsilon \cdot \frac{\partial s_n}{\partial a_n}$$

$$\text{where } |\epsilon \cdot \frac{\partial s_n}{\partial a_n}| \leq \delta \text{ for some } \delta > 0.$$

Note that the amount of perturbation ϵ is bounded by a threshold factor δ , as we do not want to perturb s_{n+1} too much. This is because we still want to take advantage of the commonly used advanced advantage estimation techniques, such as GAE [39], in estimating the advantage, and they usually consider multiple time steps in a trajectory. With this constraint, we ensure that the perturbed trajectory does not deviate from the original one too much so that we can still

use the advantage estimation techniques that use the whole trajectory.

By perturbing the action, reward, and next state in this fashion, we can expect that our experience buffer is filled with more “important” experience than before, relative to the reward function. Intuitively, an action that produces a higher immediate reward than other actions is more likely to bring about more meaningful results. This intuition is not true in every case, but in many cases, we can observe that it holds. By using this more important experience and manipulating our policy based on it, we can expect our policy to learn more meaningful lessons from the same experience than before. In experiments (Figure 2), gradient-enhanced algorithms are prefixed with “Diff”. Thus, gradient-enhanced TRPO becomes DiffTRPO, and gradient-enhanced PPO becomes DiffPPO.

V. IMPLEMENTATION AND RESULTS

We show results for improved traffic flow via CARLA Leaderboard test scenarios, which describe 10 categories of scenarios based on the National Highway Traffic Safety Administration (NHTSA) pre-crash typology. For experiments with TransFuser in Table I, we use the Longest6 benchmark presented in [36].

A. Integration with SOTA and Improvements

We implement TrAAD on recent state-of-the-art methods for imitation learning to demonstrate its generalizability and benefits on higher-performing benchmarks. Evaluation of learned driving policies for each model is based on CARLA driving benchmark metrics, which quantify route completion, infractions, collisions, and timeouts, among others. We observe from the top half of the results in Table I that TrAAD implemented on Learning By Cheating (LBC) [34] is able to improve Route Completion (%) of scenario route driven,

by a factor of *over 2x*. In addition, we achieve lower overall collisions with pedestrians, other vehicles, and other objects. Agent timeouts are also improved by a factor of 6. Overall, we improve the overall Driving Score by *over 2x* as well, where Driving Score is defined as the product between the Route Completion and an infraction penalty.

Mean traffic flow and mean fuel consumption for each benchmark CARLA scenario are visualized in Figure 3. Our model consistently improves overall traffic flow over baseline models in the local lane, while maintaining similar fuel consumption despite the natural relation between higher velocities and higher fuel consumption.

We also evaluate TrAAD on the most recent state-of-the-art model from 2022, TransFuser [37]. While this method was originally published in 2021, we use the improved 2022 version of the method. Since this method achieves significantly better performance than LBC, we can observe the scaling effects of our method on a driving model with higher benchmark metrics. Despite improving Route Completion and collisions against pedestrians and other vehicles, we observe a slightly worse Driving Score. This is most likely due to our model incurring a marginally higher rate of traffic infractions. Since TrAAD increases overall flow, a slightly higher rate of infractions is likely as a vehicle may be traveling at a higher velocity when encountering red lights or stop signs. However, this slight degradation is minimal and does not actually result in more collisions overall.

B. Ablation: Effect of Traffic Gradients on Performance

We show the impact of traffic gradients and perturbation threshold values δ during training for the on-policy algorithm TRPO [33] in Figure 2. In this experiment, we compare the training reward curve of TRPO versus gradient-enhanced DiffTRPO with different perturbation threshold δ values. We do not evaluate this enhancement on CARLA benchmark scenarios, as not every scenario involves a difficult traffic-related task. If the task is not difficult enough, the benefits of using traffic gradients are less pronounced in learning. To demonstrate improvement in difficult traffic tasks, we use the Figure-Eight benchmark from [30], [32]. In this environment, the ego vehicle is tasked with controlling dense traffic flow to maximize overall traffic flow to a target speed, i.e. *congestion and shockwaves are undesirable in the optimal policy*.

While all DiffTRPO iterations visibly perform better than the baseline TRPO, threshold values of $\delta = 0.2$ and $\delta = 0.4$ converge the fastest to the highest reward. Each training curve is averaged over ten runs to account for stochasticity. More results for PPO can be found in supplemental materials, where $\delta = 0.1$ achieves the best results for DiffPPO.

C. Ablation: Effect of Each Reward Term

We study the effect of each of the three terms of the reward function. In this ablation experiment, we analyze three scenarios: (1) Without velocity term, (2) Without fuel consumption term, (3) Without jerkiness constraint. We evaluate them on societal metrics of traffic flow, fuel consumption, and driving jerkiness, as well as evaluation

metrics from CARLA driving benchmark similar to Section V-A. Results for this experiment can be found in Table II.

Overall, we find that our model expectedly achieves a middle ground between ablated models. In our combined reward function, we optimize for both fuel efficiency and traffic flow. However, these terms are inversely related in the real world; as overall traffic flow increases, more fuel is used in order to produce a higher velocity. Thus, we observe from the results that the model omitting the fuel term is able to achieve the highest overall traffic flow, as fuel efficiency does not constrain the optimization of traffic flow. In addition, omitting the jerk constraint results in fewer infractions. For scenarios involving yielding to traffic laws, such as red lights and stop signs, restricting jerkiness may harm infraction scores as vehicles cannot suddenly decelerate to a complete stop. Overall, we observe that TrAAD achieves the “best of both worlds” across *all* other metrics, while the ablated models achieve the best results solely on individual metrics.

VI. CONCLUSION AND DISCUSSION

In this paper, we present a method for traffic-aware autonomous driving (TrAAD) that benefits both the autonomous vehicle and the traffic around it. TrAAD complements existing work by providing additional supervision on acceleration behavior guided by differentiable traffic simulation. In addition, we improve the sample efficiency of on-policy RL algorithms using analytical gradients of car-following models and show that our method produces a driving policy that not only benefits the surrounding traffic system but also improves the driving performance of each autonomous vehicle on multiple benchmarks. Distillation results also imply interesting future work for distillation for real-world models.

Limitations and Future Work

Independent of our work, the current implementation integrated with CARLA simulation for learning imposes some limitations. For example, the shortage of difficult traffic scenarios on available driving benchmarks within CARLA limits more evaluation of our approach with FLOW on vehicle control; similarly with limited traffic tasks on the current driving benchmarks. Furthermore, current driving benchmarks do not consider societal goals, such as the overall traffic flow and fuel consumption, and thus may downplay the potential positive impact of our method. More driving benchmarks that consider the global objectives of the entire road network system would allow for a more comprehensive evaluation of learned policies for AVs.

REFERENCES

- [1] J. Kim, "Estimating the social cost of congestion using the bottleneck model," Mar 2019. [Online]. Available: <https://papers.ssrn.com/abstract=3356167>
- [2] "Sources of greenhouse gas emissions," <https://www.epa.gov/ghgemissions/sources-greenhouse-gas-emissions>, accessed: 2022-09.
- [3] "Road traffic injuries and deaths—a global problem," <https://www.cdc.gov/injury/features/global-road-safety/index.html>, accessed: 2022-09.
- [4] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Physical Review E*, vol. 62, no. 2, p. 1805–1824, Aug 2000.
- [5] G. F. Newell, "A simplified car-following theory: a lower order model," *Transportation Research Part B: Methodological*, vol. 36, no. 3, p. 195–205, Mar 2002.
- [6] R. Wiedemann, *Simulation des Strassenverkehrsflusses*. Karlsruhe: Institut für Verkehrswesen der Universität Karlsruhe, 1974.
- [7] S. Krauss, "Microscopic modeling of traffic flow: investigation of collision free vehicle dynamics," Apr 1998.
- [8] M. J. Lighthill and G. B. Whitham, "On kinematic waves ii. a theory of traffic flow on long crowded roads," *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 229, no. 1178, pp. 317–345, 1955.
- [9] D. C. Gazis, R. Herman, and R. W. Rothery, "Nonlinear follow-the-leader models of traffic flow," *Operations research*, vol. 9, no. 4, pp. 545–567, 1961.
- [10] A. Aw and M. Rascle, "Resurrection of" second order" models of traffic flow," *SIAM journal on applied mathematics*, vol. 60, no. 3, pp. 916–938, 2000.
- [11] H. M. Zhang, "A non-equilibrium traffic model devoid of gas-like behavior," *Transportation Research Part B: Methodological*, vol. 36, no. 3, pp. 275–290, 2002.
- [12] B. S. Kerner, *Introduction to Modern Traffic Flow Theory and Control*. Springer Berlin Heidelberg.
- [13] I. O. Olayode, L. K. Tartibu, and M. O. Okwu, "Prediction and modeling of traffic flow of human-driven vehicles at a signalized road intersection using artificial neural network model: A south african road transportation system scenario," *Periodica Polytechnica: Transportation Engineering*, vol. 6, p. 100095, Dec 2021.
- [14] F. Diehl, T. Brunner, M. T. Le, and A. Knoll, "Graph neural networks for modelling traffic participant interaction," in *2019 IEEE Intelligent Vehicles Symposium (IV)*. ieeexplore.ieee.org, Jun 2019, p. 695–701.
- [15] W. Jiang and J. Luo, "Graph neural network for traffic forecasting: A survey," *Expert systems with applications*, vol. 207, p. 117921, Nov 2022.
- [16] L. N. N. Do, N. Taherifar, and H. L. Vu, "Survey of neural network-based models for short-term traffic state prediction," *Wiley interdisciplinary reviews. Data mining and knowledge discovery*, vol. 9, no. 1, p. e1285, Jan 2019.
- [17] M. Zhu, Y. Wang, Z. Pu, J. Hu, X. Wang, and R. Ke, "Safe, efficient, and comfortable velocity control based on reinforcement learning for autonomous driving," *Transportation Research Part C: Emerging Technologies*, vol. 117, p. 102662, Aug 2020.
- [18] X. Shen, X. Zhang, T. Ouyang, Y. Li, and P. Raksincharoensak, "Co-operative comfortable-driving at signalized intersections for connected and automated vehicles," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, p. 6247–6254, Oct 2020.
- [19] M. Zhu, X. Wang, and J. Hu, "Impact on car following behavior of a forward collision warning system with headway monitoring," *Transportation Research Part C: Emerging Technologies*, vol. 111, p. 226–244, Feb 2020.
- [20] G. Li, Y. Yang, S. Li, X. Qu, N. Lyu, and S. E. Li, "Decision making of autonomous vehicles in lane change scenarios: Deep reinforcement learning approaches with risk awareness," *Transportation Research Part C: Emerging Technologies*, vol. 134, p. 103452, Jan 2022.
- [21] M. Wegener, L. Koch, M. Eisenbarth, and J. Andert, "Automated eco-driving in urban scenarios using deep reinforcement learning," *Transportation Research Part C: Emerging Technologies*, vol. 126, p. 102967, May 2021.
- [22] P. Andelfinger, "Differentiable agent-based simulation for gradient-guided simulation-based optimization," in *Proceedings of the 2021 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*. New York, NY, USA: ACM, May 2021. [Online]. Available: <https://dl.acm.org/doi/10.1145/3437959.3459261>
- [23] S. Son, Y.-L. Qiao, J. Sewall, and M. C. Lin, "Differentiable hybrid traffic simulation," *ACM Trans. Graph.*, vol. 41, no. 6, nov 2022. [Online]. Available: <https://doi.org/10.1145/3550454.3555492>
- [24] S. Aradi, "Survey of deep reinforcement learning for motion planning of autonomous vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 2, pp. 740–759, 2022.
- [25] Z. Qiao, K. Muelling, J. M. Dolan, P. Palanisamy, and P. Mudalige, "Automatically generated curriculum based reinforcement learning for autonomous vehicles in urban environment," in *2018 IEEE Intelligent Vehicles Symposium (IV)*, Jun 2018, p. 1233–1238.
- [26] M. Bouton, A. Nakhaei, K. Fujimura, and M. J. Kochenderfer, "Cooperation-aware reinforcement learning for merging in dense traffic," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, Oct 2019, p. 3441–3447.
- [27] M. Kaushik, V. Prasad, K. M. Krishna, and B. Ravindran, "Overtaking maneuvers in simulated highway driving using deep reinforcement learning," in *2018 IEEE Intelligent Vehicles Symposium (IV)*, Jun 2018, p. 1885–1890.
- [28] A. Ferdowsi, U. Challita, W. Saad, and N. B. Mandayam, "Robust deep reinforcement learning for security and safety in autonomous vehicle systems," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, Nov 2018, p. 307–312.
- [29] X. Ma, K. Driggs-Campbell, and M. J. Kochenderfer, "Improved robustness and safety for autonomous vehicle control with adversarial reinforcement learning," in *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE Press, Jun 2018, p. 1665–1671.
- [30] C. Wu, A. Kreidieh, K. Parvate, E. Vinitzky, and A. M. Bayen, "Flow: Architecture and benchmarking for reinforcement learning in traffic control," *arXiv preprint arXiv:1710.05465*, vol. 10, 2017.
- [31] 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 *The 21st IEEE International Conference on Intelligent Transportation Systems*. IEEE, 2018. [Online]. Available: <https://elib.dlr.de/124092/>
- [32] E. Vinitzky, A. Kreidieh, L. L. Flem, N. Kheterpal, K. Jang, C. Wu, F. Wu, R. Liaw, E. Liang, and A. M. Bayen, "Benchmarks for reinforcement learning in mixed-autonomy traffic," in *Proceedings of The 2nd Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, A. Billard, A. Dragan, J. Peters, and J. Morimoto, Eds., vol. 87. PMLR, 29–31 Oct 2018, pp. 399–409. [Online]. Available: <https://proceedings.mlr.press/v87/vinitzky18a.html>
- [33] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, F. Bach and D. Blei, Eds., vol. 37. Lille, France: PMLR, 07–09 Jul 2015, p. 1889–1897.
- [34] D. Chen, B. Zhou, V. Koltun, and P. Krähenbühl, "Learning by cheating," in *Conference on Robot Learning (CoRL)*, 2019.
- [35] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator," in *Proceedings of the 1st Annual Conference on Robot Learning*, 2017, pp. 1–16.
- [36] K. Chitta, A. Prakash, B. Jaeger, Z. Yu, K. Renz, and A. Geiger, "Transfuser: Imitation with transformer-based sensor fusion for autonomous driving," *Pattern Analysis and Machine Intelligence (PAMI)*, 2022.
- [37] A. Prakash, K. Chitta, and A. Geiger, "Multi-modal fusion transformer for end-to-end autonomous driving," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [38] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [39] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, "High-dimensional continuous control using generalized advantage estimation," *arXiv preprint arXiv:1506.02438*, 2015.
- [40] Y.-L. Qiao, J. Liang, V. Koltun, and M. C. Lin, "Efficient differentiable simulation of articulated bodies," in *International Conference on Machine Learning*. PMLR, 2021, pp. 8661–8671.
- [41] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *International conference on machine learning*. PMLR, 2015, pp. 1889–1897.