Learning Realistic Traffic Agents in Closed-loop

Anonymous Author(s)
Affiliation
Address
email

Abstract: Realistic traffic simulation is crucial for developing self-driving software in a safe and scalable manner prior to real-world deployment. Typically, imitation learning (IL) is used to learn human-like traffic agents directly from real-world observations collected offline, but without explicit specification of traffic rules, agents trained from IL alone frequently display unrealistic infractions like collisions and driving off the road. This problem is exacerbated in out-of-distribution and long-tail scenarios. On the other hand, reinforcement learning (RL) can train traffic agents to avoid infractions, but using RL alone results in unhuman-like driving behaviors. We propose Reinforcing Traffic Rules (RTR), a holistic closed-loop learning objective to match expert demonstrations under a traffic compliance constraint, which naturally gives rise to a joint IL + RL approach, obtaining the best of both worlds. Our method learns in closed-loop simulations of both nominal scenarios from real-world datasets as well as procedurally generated long-tail scenarios. Our experiments show that RTR learns more realistic and generalizable traffic simulation policies, achieving significantly better tradeoffs between human-like driving and traffic compliance in both nominal and long-tail scenarios. Moreover, when used as a data generation tool for training prediction models, our learned traffic policy leads to considerably improved downstream prediction metrics compared to baseline traffic agents.

Keywords: Traffic simulation, Imitation learning, Reinforcement learning

1 Introduction

Simulation is a critical component to safely developing autonomous vehicles, particularly for identifying potential failures in safety-critical situations. Learning traffic agents that are realistic is fundamental to build high-fidelity simulation systems that have a low domain gap to the real world. Learning realistic agents is, however, challenging as we need to capture the idiosyncratic nature of *human-like* driving and avoid unrealistic traffic *infractions* like collisions or driving off-road. Existing approaches used in the self-driving industry lack realism: they either replay logged trajectories in a non-reactive manner [1, 2] or use heuristic policies which yield rigid, unhuman-like behaviors. Using data-driven approaches to learn more realistic policies is a promising alternative.

The dominant data-driven approach has been imitation learning (IL), where nominal human driving data is used as expert supervision to train the agents. However, while expert demonstrations provide supervision for human-like driving, pure IL methods lack explicit knowledge of traffic rules and infractions which can result in unrealistic policies. Furthermore, the reliance on expert demonstrations can be a disadvantage, as long-tail scenarios with rich interactions are very rare, and thus learning is overwhelmingly dominated by more common scenarios with a much weaker learning signal.

Reinforcement learning (RL) approaches encode explicit knowledge of traffic rules through hand-designed rewards that penalize infractions [3, 4, 5, 6, 7, 5]. These approaches do not rely on expert demonstrations and instead learn to maximize traffic-compliance rewards through trial and error. In the context of autonomy, this allows training on synthetic scenarios that do not have expert demon-

Submitted to the 7th Conference on Robot Learning (CoRL 2023). Do not distribute.

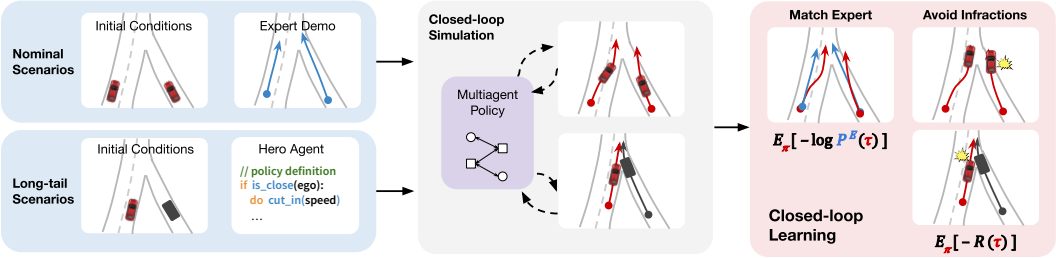

Figure 1: Our multi-agent policy is trained in closed-loop to match expert demonstrations under a traffic compliance constraint using both nominal offline data and simulated long-tail scenarios as a rich learning environment. This gives rise to an IL objective which supervises the policy using real-world expert demonstrations and an RL objective which explicitly penalizes infractions.

strations in order to improve the robustness of learned policies [5]. However, traffic rules alone cannot describe all the nuances of human-like driving, and it is still an open question if one can manually design a reward that can completely capture those intricacies.

Towards learning human-like and traffic-compliant agents, we propose Reinforcing Traffic Rules (RTR), a holistic closed-loop learning method to match expert demonstrations under a traffic-compliance constraint using both nominal offline data and additional simulated long-tail scenarios (Figure 1). We show our formulation naturally gives rise to a unified closed-loop IL + RL objective, which we efficiently optimize by exploiting differentiable dynamics and a per-agent factorization. In contrast to prior works that combine IL and RL [1], our closed-loop approach allows the model to understand the effects of its actions and suffers significantly less from compounding error. Furthermore, exploiting simulated long-tail scenarios improves learning by exposing the policy to more interesting interactions that would be difficult and possibly dangerous to collect from the real world at scale. Our experiments show that, unlike a wide range of baselines, RTR learns realistic policies that better generalize to both nominal and long-tail scenarios *unseen during training*. The benefits carry forward to *downstream tasks* such as simulating scenarios to train autonomy models; prediction models trained on data simulated with RTR have the strongest prediction metrics on real data, serving as further evidence that RTR has learned more realistic traffic simulation. We believe this serves as a crucial step towards more effective and scalable applications of traffic simulation for training and testing autonomous vehicles.

2 Related Work

Traditional traffic simulation: To generate general traffic flow, simulators [8, 9, 10, 11] typically use heuristic models [12, 13, 14] as models of human driving. While these heuristic models are useful in capturing high-level traffic characteristics like flow and density, they are lacking in capturing the lower-level nuances of human driving, thus limiting their applicability in self-driving. For more realistic traffic models, we explore using machine learning as a more promising approach.

Imitation learning: IL methods learn a control policy from expert demonstrations. In the context of autonomous vehicles, [15] pioneered the use of behavior cloning (BC) to learn a driving policy in a simple environment. Since then, BC has been explored for both autonomy [16, 17, 18] and traffic simulation [19]. BC methods primarily suffer from distribution shift due to compounding error [20], and so various techniques like data augmentation [21, 16] and uncertainty-based regularization [22, 23] have been proposed to alleviate the problem. Closed-loop imitation learning approaches [24, 25, 26, 27], which address distribution shift by exposing the policy to a self-induced state distribution during training, have also been explored in traffic simulation [28, 29, 30, 31]. While IL exploits expert demonstrations, there is a lack of explicit knowledge on safety-critical aspects like avoiding infractions. Methods like differentiable common-sense penalties [28, 21] and additional finetuning [32] have been proposed to complement the standard IL loss. In this work, we use reinforcement learning to explicitly encode general non-differentiable traffic rules.

Reinforcement learning: RL methods [33, 34, 35] do not require expert demonstrations and instead learn through interacting with the environment and a reward function. In self-driving, knowledge of infractions can be encoded in the reward [3, 4, 5, 6, 7]. Because RL does not require expert demonstrations, it is possible to train on procedurally generated scenarios for improved infraction avoidance [5]. However, it is difficult to learn realistic driving behavior using reward alone. RL methods can be sample inefficient [3, 4], and specifying human-like driving with a scalar reward is difficult. In this work, we supplement RL with IL to learn more human-like driving while still enjoying the explicit learning signal provided from the reward.

Combined IL + RL: Pretrained IL policies can be used as initialization to guide exploration [36, 37] or regularize learning [38, 39, 40, 41], and offline data can be used to bootstrap learning and help with sparse rewards [42, 43]. Offline RL methods also use IL to help with out-of-distribution generalization and overestimation [44, 45]. In self-driving, IL has been used as a pre-training phase improve sample efficiency [46]. Recent work augments open-loop BC with RL [1] to learn more robust autonomy models. While promising, the open-loop nature of BC leaves the policy susceptible to distribution shift. In this work, we explore a holistic closed-loop IL + RL method for traffic simulation. We also investigate using simulated scenarios where expert demonstration does not exist, rather than relying solely on mining real data which can be unsafe and expensive to collect.

3 Learning Infraction-free Human-like Traffic Agents

To learn realistic infraction-free agents, we propose a unified learning objective to match expert demonstrations under an infraction-based constraint. We show how our formulation naturally gives rise to a joint closed-loop IL + RL approach which allows learning from both offline collected human driving data when possible, and additional simulated long-tail scenarios containing rich interactions that would otherwise be difficult or impossible to collect in the real world.

3.1 Preliminaries

We model multi-agent traffic simulation as a Markov Decision Process $\mathcal{M} = (\mathcal{S}, \mathcal{A}, R, P, \gamma)$ with state space, action space, reward function, transition dynamics, and discount factor respectively. As our focus is *traffic simulation* where we have access to all ground truth states, we opt for a fully observable and centralized multi-agent formulation where a single model jointly controls all agents. This enables efficient inference by sharing computation and easier interaction modeling[1].

State, action and policy: We define the state $\boldsymbol{s} = \{s^{(1)}, \ldots, s^{(N)}, \boldsymbol{m}\} \in \mathcal{S}$ to be the joint states of N agents where N may vary across different scenarios, as well as an HD map $\boldsymbol{m}$ which captures the road and lane topology. We parameterize the state of the i-th agent $s^{(i)}$ with its position, heading, and velocity over the past H history timesteps. The state also captures 2D bounding boxes for each agent. Likewise, $\boldsymbol{a} = \{a^{(1)}, \ldots, a^{(N)}\} \in \mathcal{A}$ is the joint action which contains the actions taken by all the agents. The i-th agent's action $a^{(i)}$ is parameterized by its acceleration and steering angle. Agents are controlled by a single centralized policy $\pi(\boldsymbol{a}|\boldsymbol{s})$ which maps the joint state to joint action.

Trajectories and dynamics: We define a trajectory $\tau_{0:T} = (\boldsymbol{s}_0, \boldsymbol{a}_0, \ldots, \boldsymbol{s}_{T-1}, \boldsymbol{a}_{T-1}, \boldsymbol{s}_T)$ as a sequence of state action transitions of length T for all agents. We use the kinematic bicycle model [47] as a simple but realistic model of transition dynamics $P(\boldsymbol{s}_{t+1}|\boldsymbol{s}_t, \boldsymbol{a}_t)$ for each agent. Trajectories can be sampled by first sampling from some initial state distribution ρ_0 before unrolling a policy π through the transition dynamics i.e. $P^{\pi}(\tau) = \rho_0(\boldsymbol{s}_0) \prod_{t=0}^{T-1} \pi(\boldsymbol{a}_t|\boldsymbol{s}_t) P(\boldsymbol{s}_{t+1}|\boldsymbol{s}_t, \boldsymbol{a}_t)$.

Reward: Let $R^{(i)}(\boldsymbol{s}, a^{(i)})$ be a per-agent reward which is specific for the i-th agent, but dependent on the state of all agents to model interactions such as collision. The joint reward is then $R(\boldsymbol{s}, \boldsymbol{a}) = \sum_i^N R^{(i)}(\boldsymbol{s}, a^{(i)})$, with $R(\tau) = \sum_{t=0}^{T-1} \gamma^t R(\boldsymbol{s}_t, \boldsymbol{a}_t)$ as the γ-discounted return of a trajectory.

Policy learning: Both imitation learning (IL) and reinforcement learning (RL) can be described in this framework. IL can be described as an f-divergence minimization problem: $\pi^* =$

[1] In our experiments, our model easily scales to 50 agents and a map ROI of $1000m \times 400m$ per simulation.

$\arg\min_\pi D_f\left(P^\pi(\tau) \parallel P^E(\tau)\right)$ where P^E is the expert-induced distribution. RL on the other hand aims to find the policy which maximizes the expected reward $\pi^* = \arg\max_\pi \mathbb{E}_{P^\pi}\left[R(\tau)\right]$.

3.2 Learning

To learn a multiagent traffic policy that is as human-like as possible while avoiding infractions, we consider the reverse KL divergence to the expert distribution with an infraction-based constraint

$$\begin{aligned} \arg\min_\pi \quad & D_{\text{KL}}\left(P^\pi(\tau) \parallel P^E(\tau)\right) \\ \text{s.t.} \quad & \mathbb{E}_{P^\pi}\left[R(\tau)\right] \geq 0 \end{aligned} \tag{1}$$

$$R^{(i)}(\boldsymbol{s}, a^{(i)}) = \begin{cases} -1 & \text{if infraction} \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

where $R^{(i)}$ is a per-agent reward function that penalizes any infractions (collision and off-road events). For a rich learning environment, we consider both a dataset D of nominal expert trajectories $\tau^E \sim P^E$ collected by driving in the real world, and additional simulated *long-tail scenarios*. These scenarios contain what we denote as *hero* agents, which induce interesting interactions like sudden cut-ins, etc. (details in Section 3.4). More precisely, let π_θ be our learner policy. Let $\boldsymbol{s}_0^S \sim \rho_0^S$ be the initial state sampled from the long-tail distribution and $\pi^S_{\boldsymbol{s}_0}$ represent the policy of the hero agent. The overall multiagent policy is given as

$$\pi(\boldsymbol{a}_{i,t}|\boldsymbol{s}_t) = \begin{cases} \pi^S_{\boldsymbol{s}_0}(a_t^{(i)}|\boldsymbol{s}_t) & \text{if agent } i \text{ is hero} \\ \pi_\theta(a_t^{(i)}|\boldsymbol{s}_t) & \text{otherwise} \end{cases} \tag{3}$$

The overall initial state distribution is then given as $\rho_0 = (1-\alpha)\rho_0^D + \alpha\rho_0^s$, where ρ_0^D corresponds to the offline nominal distribution, and $\alpha \in (0, 1)$ is a hyperparameter that balances the mixture.

By taking the Lagrangian of Equation 1, the objective is naturally decomposed into an IL and RL component,

$$\mathcal{L} = \mathbb{E}_{P^\pi}\left[\underbrace{-\log P^E(\tau)}_{\text{IL}} - \lambda \underbrace{R(\tau)}_{\text{RL}}\right] - H(\pi) = \mathcal{L}^{\text{IL}} + \lambda\mathcal{L}^{\text{RL}} - H(\pi) \tag{4}$$

where λ is a hyperparameter balancing the two terms, and $H(\pi)$ is an additional entropy regularization term [24] [2]. Notably, we optimize IL and RL *jointly in a closed-loop manner*, as the expectation is taken with respect to the on-policy distribution $P^\pi(\tau)$. Compared to open-loop behavior cloning, the closed-loop IL component allows the model to experience states induced by its own policy rather than only the expert distribution, increasing its robustness to distribution shift. Furthermore, while the additional reward constraint may not change the optimal solution of the unconstrained problem (the expert distribution may be infraction-free), it can provide additional learning signal through RL.

The RL component $\mathbb{E}_{P^\pi}\left[R(\tau)\right]$ can be optimized using standard RL techniques and exploits both offline-collected nominal scenarios and simulated long-tail scenarios containing rich interactions. However, the imitation component $\mathcal{L}^{\text{IL}}$ is only well-defined when expert demonstrations are available and thus only applied to nominal data. We start from an initial state $\boldsymbol{s}_0^E \sim \rho_0^D$ and have the policy π_θ control all agents in closed-loop simulation. The loss is the distance between the ground truth and policy-induced trajectory [3].

$$\mathcal{L}^{\text{IL}} = \mathbb{E}_{\tau^E \sim D}\left[\mathbb{E}_{\tau \sim P^\pi(\cdot|\boldsymbol{s}_0^E)}\left[D(\tau^E, \tau)\right]\right]. \tag{5}$$

It is difficult to obtain accurate action labels for human driving data in practice, so we only consider states in our loss, i.e. $D(\tau^E, \tau) = \sum_{t=1}^T d(\boldsymbol{s}_t^E, \boldsymbol{s}_t)$ where d is a distance function (e.g. Huber).

Optimization: To optimize Equation 4, we first note that the $\mathcal{L}^{IL}$ component is differentiable by using the reparameterization trick [48] when sampling from the policy[4] and differentiating through

[2] The causal entropy term is included as an entropy regularizer in some learning algorithms such as PPO [33]. In our setting, we empirically found that it was not necessary to include.

[3] As we do not have access to P^E directly to query log-likelihood, using a distance is essentially making the assumption that $P^E(\tau) \propto \exp\left[-D(\tau^E, \tau)\right]$.

[4] We found that directly using the mean action provides good results without the need for sampling.

the transition dynamics (kinematic bicycle model). We refer the reader to the supplementary for more details. To optimize the $\mathcal{L}^{RL}$ component, we design a centralized and fully observable variant of PPO [33]. While it is possible to directly optimize the policy with the overall scene reward $R(\boldsymbol{s}, \boldsymbol{a}) = \sum_{i=1}^{N} R^{(i)}(\boldsymbol{s}, a^{(i)})$, we instead optimize each agent individually with their respective individual reward $R_i(\boldsymbol{s}, a_i)$. While this factorized approach may ignore second-order interaction effects, it considerably simplifies the credit assignment problem leading to more efficient learning. More precisely, we compute factorized value targets $V^{(i)} = \sum_{t=0}^{T} \gamma^t R_t^{(i)}(\boldsymbol{s}_t, a_t^{(i)})$, and the factorized PPO policy loss is given as $\mathcal{L}^{\text{policy}} = \sum_{i=1}^{N} \min(r^{(i)} A^{(i)}, \text{clip}(r^{(i)}, 1-\epsilon, 1+\epsilon) A^{(i)})$ where the probability ratio is factorized, i.e. $r^{(i)} = \pi(a^{(i)}|\boldsymbol{s}, \boldsymbol{m}) / \pi_{\text{old}}(a^{(i)}|\boldsymbol{s}, \boldsymbol{m})$ and $A^{(i)}$ is a factorized GAE [49] estimate. More details can be found in the supplementary.

3.3 Model Architecture

Our traffic model π_θ architecture uses common ideas from SOTA traffic agent motion forecasting literature in order to extract context and map features and predict agent actions (Figure 2). Recall that a state $\boldsymbol{s} = \{s^{(1)}, \ldots, s^{(N)}, \boldsymbol{m}\}$ consists of each individual agent's states $s^{(i)}$ that contain the agent's kinematic state over a history horizon H, and an HD map $\boldsymbol{m}$. From each agent's state history, a shared 1D CNN and GRU are used to extract agent history context features $h_a^{(i)} = f(s^{(i)})$. At the same time, a GNN is used to extract map features from a lane graph representation of the map input $h_m = g(\boldsymbol{m})$. A HeteroGNN [50] then jointly fuses all agent context features and map features before a shared MLP decodes actions for each agent independently.

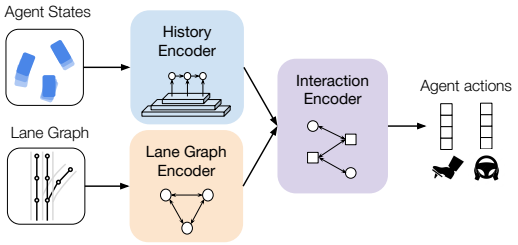

Figure 2: Our multiagent policy architecture. The value network uses the same architecture but regresses value targets instead of actions.

$$\{h^{(1)}, \ldots h^{(N)}\} = \text{HeteroGNN}(\{h_a^{(1)}, \ldots, h_a^{(N)}\}, h_m) \tag{6}$$

$$(\mu^{(i)}, \sigma^{(i)}) = \text{MLP}(h^{(i)}). \tag{7}$$

We use independent normal distributions to represent the joint agent policy, i.e. $\pi(a^{(i)}|\boldsymbol{s}) = \mathcal{N}(\mu^{(i)}, \sigma^{(i)})$, and thus $\pi(\boldsymbol{a}|\boldsymbol{s}) = \prod_{i=1}^{N} \pi(a^{(i)}|\boldsymbol{s})$. Note that agents are only independent *conditional* on their shared past context, and thus important interactive reasoning is still captured. Our value model uses the same architecture but does not share parameters with the policy; we compute $\{h_0^v, \ldots, h_N^v\}$ in a similar fashion, and decode per-agent value estimates $\hat{V}^{(i)} = \text{MLP}^v(h^{(i)})$.

3.4 Simulated Long-tail Scenarios

Nominal driving logs can be monotonous and provide weak learning signal when repeatedly used for training. In reality, most traffic infractions can be attributed to rare and long-tail scenarios belonging to a handful of scenario families [51] which can be difficult and dangerous to collect from the real world at scale. In this work, we procedurally generate long-tail scenarios to supplement nominal logs for training and testing. Following the self-driving industry standard, we use *logical scenarios* [52, 53] which vary in the behavioral patterns of particular hero agents with respect to an ego agent (e.g. cut-in, hard-braking, merging, etc.). Each logical scenario is parameterized by $\theta \in \Theta$ which controls lower-level aspects of the scenario such as behavioral characteristics of the hero agent (e.g. time-to-collision or distance triggers, aggressiveness, etc.), exact initial placement and kinematic states, and geolocation. A *concrete scenario* can then be procedurally generated in an automated fashion by sampling a logical scenario and corresponding parameters θ. While these scenarios are simulated and thus do not have associated human demonstrations, they provide rich learning signal due to interesting and rare interactions induced by the hero agents.

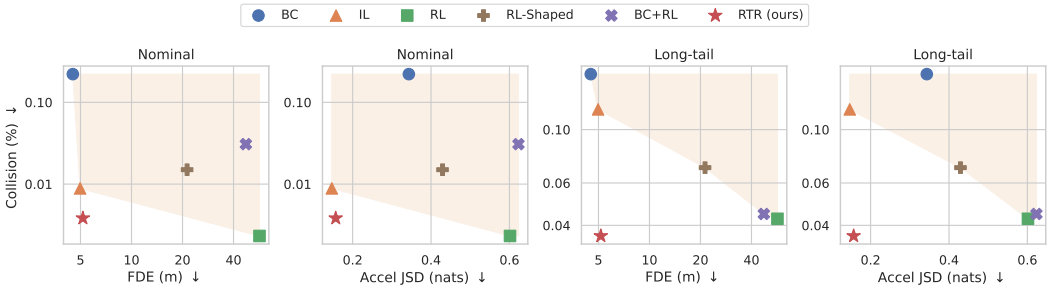

Figure 3: Metrics (lower is better) on held-out nominal and long-tail scenarios. Pareto frontier of baselines is shaded; RTR achieves the best tradeoff between infraction and other realism metrics.

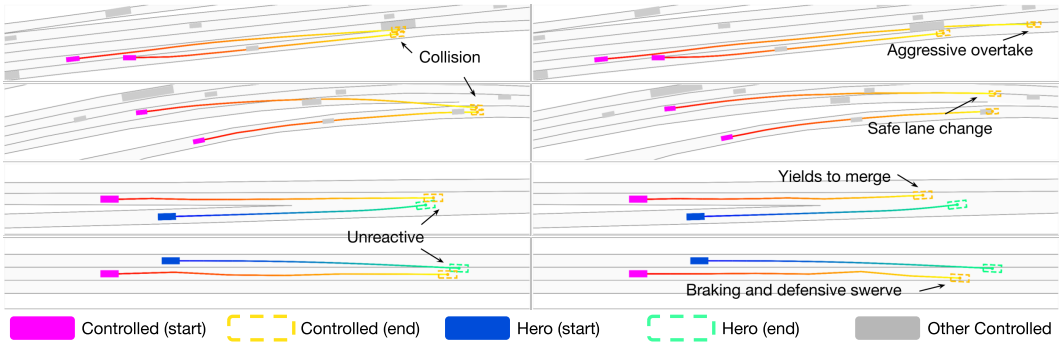

Figure 4: Qualitative examples comparing the baseline IL model (left) and RTR (right). Scenarios with hero agents (blue) are from the long-tail set. All other agents are controlled; pink is used for visual emphasis. RTR avoids infractions while maintaining diverse, human-like driving behavior.

4 Experiments

Scenario sets: Our experiments use two datasets that represent nominal and long-tail scenarios respectively. The NOMINAL dataset consists of a set of highway logs which capture varying traffic densities and road topologies while containing expert demonstrations. The dataset consists of 465 snippets for training and 115 for testing, where each snippet lasts for 20 seconds. We use LONGTAIL to denote the scenario set generated using the process outlined in Section 3.4 which contain rare scenarios like sudden cut-ins, etc. We use 25 logical scenarios to generate a total of 333 concrete scenarios, where 167 concrete scenarios are used for training and 166 are held-out for evaluation.

Metrics: We evaluate our traffic models' ability to 1) match human-like driving and 2) avoid infractions. For the former, we measure similarity to the demonstration data by computing the final displacement error (**FDE**) [31, 28], which measures the L2 distance between the agent's simulated and ground truth (GT) position after 5 seconds. Furthermore, we use Jensen-Shannon Divergence (**JSD**) [31, 29] between histograms of scenario features (agent acceleration) in order to measure distributional realism. Finally, to measure infraction rates, we consider **collision** and driving **off-road**. Results with more extensive metrics (and their definitions) can be found in the supplementary.

4.1 Benchmarking Traffic Models

Comparison to state-of-the-art: We evaluate RTR and several baselines on both nominal and long-tail scenarios. For comparability, we use the same input representation and model architecture as described in Section 3.3 for all methods. Our first two baselines are representative of state-of-the-art imitation learning approaches for traffic simulation. **BC** is our single-step behavior cloning baseline following [19]. The **IL** baseline is trained using closed-loop policy unrolling [28, 29]. Next, **RL** is trained using our proposed factorized version of PPO [33] with the reward in Equation 2. The

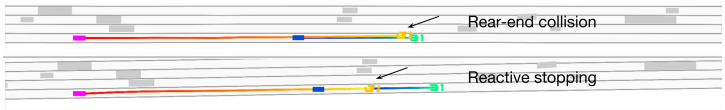

Figure 6: IL (top), RTR (bottom) on an out-of-distribution scenario where a hero agent (blue) comes to a complete stop on the highway.

Method	Col.	Off.
IL	11.81	1.02
RTR	**4.99**	**0.28**

Table 1: Results on out-of-distribution long-tail set.

RL-Shaped baseline includes an additional reward for driving at the speed-limit to encourage more human-like driving. Finally, **BC+RL** is an RL augmented BC baseline following [1].

Figures 3 and 4 show the results; a full table can be found in the supplementary. Firstly, the BC model achieves poor realism because it suffers from distribution shift during closed-loop evaluation as it encounters states unseen during training due to compounding error. Next, we see RL achieves low infraction rates but results in unhuman-like driving (Figure 5). This is because it is difficult for reward alone to capture realistic driving. Efforts in reward shaping result in improvements but are ultimately still insufficient. We see RL augmented BC improves upon BC infraction rates but still lacks realism, since open-loop BC does not provide supervision on states encountered during RL. On the other hand, closed-loop IL performs better as it is more robust to distribution shift, but still struggles on the long-tail scenario set without explicit supervision. Finally, the *holistic* closed-loop IL and RL approach of RTR improves infraction rates while maintaining reconstruction and JSD metrics.in We see RTR outperforms even pure RL in terms of infraction rate on long-tail scenarios, suggesting that including long-tail scenarios during training can help the model generalize to held-out evaluation long-tail scenarios.

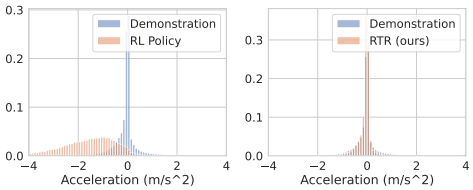

Figure 5: RL policy naively decelerates to avoid infractions. RTR learns to avoid collision more naturally without slowing down.

Out-of-distribution generalization: Recall from Section 3.4 that logical scenarios define a family of scenarios and concrete scenarios define variations within a family. While we have evaluated in-distribution generalization by using held-out concrete scenarios, we further evaluate on *held-out logical scenarios*. We use 11 held-out logical scenarios with new map topologies and behavioral patterns to procedurally generate an additional out-of-distribution set consisting of 84 concrete scenarios. Our results show that RTR generalizes to this set better than baselines (Figure 6, Table 1).

4.2 Downstream Evaluation

One downstream application of traffic simulation is generating synthetic data for training autonomy models. We evaluate if the improved realism of RTR transfers in this context. Each model is used to generate a synthetic dataset of 589 scenarios which we use to train a SOTA prediction model [50] before evaluating its performance on held-out real data. Besides FDE, the along-track error (ATE) and cross-track error (CTE) of predicted trajectories projected onto the GT are used as prediction metrics. Table 2 shows that using RTR to generate training data results in the best prediction model, showing that our approach can improve the application of traffic simulation in developing autonomous vehicles.

Method	FDE	ATE	CTE
BC	2.40	1.81	0.87
IL	1.80	1.44	**0.27**
RL	16.89	16.80	0.32
RL-Shp	6.80	6.52	0.33
BC+RL	9.30	9.13	0.44
RTR	**1.57**	**1.26**	**0.27**

Table 2: Pred. model trained on synthetic, evaluated on real.

4.3 Additional Analysis

Long-tail scenarios: We evaluate our approach of using procedurally generated scenarios against the alternative of mining hard scenarios from data [1, 54] by curating a set of logs from NOMINAL that the IL model commits an infraction on. Figure 7 shows that using only curated scenarios does

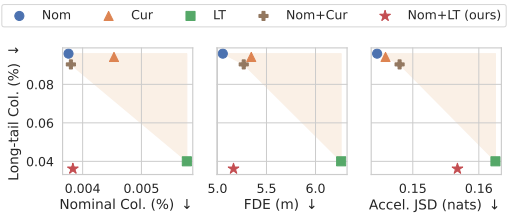

Figure 7: Using both nominal and long-tail yields the best tradeoff.

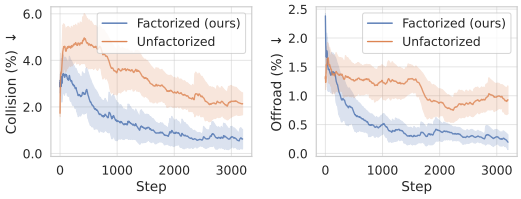

Figure 8: Our factorized PPO vs. standard PPO which uses a single scene-level reward.

not transfer well to the long-tail set, and in fact introduces a regression in the nominal scenarios, suggesting the model is overfitting to the curated scenarios. Up-sampling curated scenarios (Nom+Cur) also fails – relying purely on offline data may require prohibitively larger scale data collection.

Factorized multiagent RL: To ablate our factorized per-agent approach to multiagent PPO, we compare to a standard PPO implementation where the scene-level reward $R(\boldsymbol{s}, \boldsymbol{a})$ is used as supervision for the joint policy rather than each individual agent reward $R^{(i)}(\boldsymbol{s}, a^{(i)})$. Figure 8 shows that the factorized loss vastly outperforms the alternative, likely due to the fact that multiagent credit assignment is extremely difficult when using the scene-level reward, leading to poor sample efficiency.

Balancing the trade-off: Recall from Section 3.2 that RTR balances human-like driving and avoiding infractions by weighting the IL vs. RL loss with λ and nominal vs. long-tail training with α. We found increasing the relative weight of RL and long-tail scenarios generally improves infraction avoidance while increasing the relative weight of IL and nominal training generally improves other realism metrics as expected (Table 3). However, RTR is not particularly sensitive, and all configurations still dominate the baseline Pareto frontier.

		Nominal			LT
λ	α	Col.	FDE	Acc.	Col.
1.0	0.5	0.38	5.50	0.18	3.61
5.0	0.5	0.38	5.16	**0.16**	3.61
10	0.5	0.52	**5.10**	0.17	3.82
5.0	0.3	**0.35**	5.20	0.17	4.12
5.0	0.7	0.56	6.10	0.23	**3.51**

Table 3: Balancing realism and infraction avoidance.

Imitation learning signal: We consider the alternative of using a frozen pretrained IL policy as regularization [41] instead of our approach of using offline data. A frozen policy potentially provides more accurate closed-loop supervision, as a Euclidian-based distance loss with demonstration data may be inaccurate if the rollout has diverged. We evaluate two baselines: KL Reward and KL Loss, where the KL between the current and frozen policy is added to the reward or loss respectively. Our results in Table 4 show that using demonstration data is still the most performant, suggesting that the inaccuracy from an imperfect IL policy is larger than that of using a distance-based loss.

	Nominal			LT
Method	Col.	FDE	Acc.	Col.
KL-L	0.42	25.68	0.48	5.08
KL-R	**0.38**	15.19	0.44	4.97
RTR	**0.38**	**5.16**	**0.16**	**3.61**

Table 4: Comparing different alternatives to our proposed IL loss.

5 Conclusion and Limitations

We have presented RTR, a method for learning realistic traffic agents with closed-loop IL+RL using both real-world logs and procedurally generated long-tail scenarios. While we have shown substantial improvements over baselines in simulation realism and downstream tasks, we recognize some existing limitations. Firstly, while using logical scenarios as a framework for procedural generation exploits human prior knowledge and is currently an industry standard, manually designing scenarios can be a difficult process, and ensuring an adequate coverage of all possible scenarios is an open problem. Exploring automated alternatives like adversarial approaches to scenario generation would be an interesting future direction. Secondly, while we have explored the downstream task of generating an offline dataset to train prediction models, other applications like training and testing the entire autonomy stack end-to-end in closed-loop is a promising future direction.

References

[1] Y. Lu, J. Fu, G. Tucker, X. Pan, E. Bronstein, B. Roelofs, B. Sapp, B. White, A. Faust, S. Whiteson, et al. Imitation is not enough: Robustifying imitation with reinforcement learning for challenging driving scenarios. *arXiv preprint arXiv:2212.11419*, 2022.

[2] E. Vinitsky, N. Lichtlé, X. Yang, B. Amos, and J. Foerster. Nocturne: a scalable driving benchmark for bringing multi-agent learning one step closer to the real world. *arXiv preprint arXiv:2206.09889*, 2022.

[3] D. Chen, V. Koltun, and P. Krähenbühl. Learning to drive from a world on rails. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15590–15599, 2021.

[4] M. Toromanoff, E. Wirbel, and F. Moutarde. End-to-end model-free reinforcement learning for urban driving using implicit affordances. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7153–7162, 2020.

[5] C. Zhang, R. Guo, W. Zeng, Y. Xiong, B. Dai, R. Hu, M. Ren, and R. Urtasun. Rethinking closed-loop training for autonomous driving. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIX*, pages 264–282. Springer, 2022.

[6] X. Pan, Y. You, Z. Wang, and C. Lu. Virtual to real reinforcement learning for autonomous driving. *arXiv preprint arXiv:1704.03952*, 2017.

[7] S. Shalev-Shwartz, S. Shammah, and A. Shashua. Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295*, 2016.

[8] 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)*, pages 2575–2582. IEEE, 2018.

[9] M. Balmer, M. Rieser, K. Meister, D. Charypar, N. Lefebvre, and K. Nagel. Matsim-t: Architecture and simulation times. In *Multi-agent systems for traffic and transportation engineering*, pages 57–78. IGI Global, 2009.

[10] J. Casas, J. L. Ferrer, D. Garcia, J. Perarnau, and A. Torday. Traffic simulation with aimsun. *Fundamentals of traffic simulation*, pages 173–232, 2010.

[11] M. Ben-Akiva, H. N. Koutsopoulos, T. Toledo, Q. Yang, C. F. Choudhury, C. Antoniou, and R. Balakrishna. Traffic simulation with mitsimlab. *Fundamentals of traffic simulation*, pages 233–268, 2010.

[12] M. Treiber, A. Hennecke, and D. Helbing. Congested traffic states in empirical observations and microscopic simulations. *Physical Review E*, 62(2):1805–1824, aug 2000. doi:10.1103/physreve.62.1805. URL https://doi.org/10.1103%2Fphysreve.62.1805.

[13] K. Kreutz and J. Eggert. Analysis of the generalized intelligent driver model (gidm) for uncontrolled intersections. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 3223–3230, 2021. doi:10.1109/ITSC48978.2021.9564423.

[14] A. Kesting. Mobil : General lane-changing model for car-following models. 2007.

[15] D. A. Pomerleau. Alvinn: An autonomous land vehicle in a neural network. *Advances in neural information processing systems*, 1, 1988.

[16] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.

[17] F. Codevilla, M. Müller, A. López, V. Koltun, and A. Dosovitskiy. End-to-end driving via conditional imitation learning. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 4693–4700. IEEE, 2018.

[18] A. Prakash, K. Chitta, and A. Geiger. Multi-modal fusion transformer for end-to-end autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7077–7087, 2021.

[19] L. Bergamini, Y. Ye, O. Scheel, L. Chen, C. Hu, L. Del Pero, B. Osiński, H. Grimmett, and P. Ondruska. Simnet: Learning reactive self-driving simulations from real-world observations. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5119–5125. IEEE, 2021.

[20] S. Ross, G. Gordon, and D. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.

[21] M. Bansal, A. Krizhevsky, and A. Ogale. Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. *arXiv preprint arXiv:1812.03079*, 2018.

[22] K. Brantley, W. Sun, and M. Henaff. Disagreement-regularized imitation learning. In *International Conference on Learning Representations*, 2020.

[23] M. Henaff, A. Canziani, and Y. LeCun. Model-predictive policy learning with uncertainty regularization for driving in dense traffic. *arXiv preprint arXiv:1901.02705*, 2019.

[24] J. Ho and S. Ermon. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29, 2016.

[25] J. Fu, K. Luo, and S. Levine. Learning robust rewards with adversarial inverse reinforcement learning. *arXiv preprint arXiv:1710.11248*, 2017.

[26] S. K. S. Ghasemipour, R. Zemel, and S. Gu. A divergence minimization perspective on imitation learning methods. In *Conference on Robot Learning*, pages 1259–1277. PMLR, 2020.

[27] L. Ke, S. Choudhury, M. Barnes, W. Sun, G. Lee, and S. Srinivasa. Imitation learning as f-divergence minimization. In *Algorithmic Foundations of Robotics XIV: Proceedings of the Fourteenth Workshop on the Algorithmic Foundations of Robotics 14*, pages 313–329. Springer, 2021.

[28] S. Suo, S. Regalado, S. Casas, and R. Urtasun. Trafficsim: Learning to simulate realistic multi-agent behaviors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10400–10409, June 2021.

[29] M. Igl, D. Kim, A. Kuefler, P. Mougin, P. Shah, K. Shiarlis, D. Anguelov, M. Palatucci, B. White, and S. Whiteson. Symphony: Learning realistic and diverse agents for autonomous driving simulation, 2022. URL https://arxiv.org/abs/2205.03195.

[30] A. Ścibior, V. Lioutas, D. Reda, P. Bateni, and F. Wood. Imagining the road ahead: Multi-agent trajectory prediction via differentiable simulation. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 720–725. IEEE, 2021.

[31] S. Suo, K. Wong, J. Xu, J. Tu, A. Cui, S. Casas, and R. Urtasun. Mixsim: A hierarchical framework for mixed reality traffic simulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9622–9631, June 2023.

[32] V. Lioutas, A. Scibior, and F. Wood. Titrated: Learned human driving behavior without infractions via amortized inference. *Transactions on Machine Learning Research*, 2022.

[33] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[34] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.

[35] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

[36] I. Uchendu, T. Xiao, Y. Lu, B. Zhu, M. Yan, J. Simon, M. Bennice, C. Fu, C. Ma, J. Jiao, et al. Jump-start reinforcement learning. *arXiv preprint arXiv:2204.02372*, 2022.

[37] A. Nair, A. Gupta, M. Dalal, and S. Levine. Awac: Accelerating online reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*, 2020.

[38] N. Stiennon, L. Ouyang, J. Wu, D. Ziegler, R. Lowe, C. Voss, A. Radford, D. Amodei, and P. F. Christiano. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33:3008–3021, 2020.

[39] Y. Lu, K. Hausman, Y. Chebotar, M. Yan, E. Jang, A. Herzog, T. Xiao, A. Irpan, M. Khansari, D. Kalashnikov, et al. Aw-opt: Learning robotic skills with imitation and reinforcement at scale. In *Conference on Robot Learning*, pages 1078–1088. PMLR, 2022.

[40] Y. Zhu, Z. Wang, J. Merel, A. Rusu, T. Erez, S. Cabi, S. Tunyasuvunakool, J. Kramár, R. Hadsell, N. de Freitas, et al. Reinforcement and imitation learning for diverse visuomotor skills. *arXiv preprint arXiv:1802.09564*, 2018.

[41] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, et al. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*, 2022.

[42] T. Hester, M. Vecerik, O. Pietquin, M. Lanctot, T. Schaul, B. Piot, D. Horgan, J. Quan, A. Sendonaris, I. Osband, et al. Deep q-learning from demonstrations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

[43] M. Vecerik, T. Hester, J. Scholz, F. Wang, O. Pietquin, B. Piot, N. Heess, T. Rothörl, T. Lampe, and M. Riedmiller. Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards. *arXiv preprint arXiv:1707.08817*, 2017.

[44] A. Kumar, A. Zhou, G. Tucker, and S. Levine. Conservative q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 33:1179–1191, 2020.

[45] S. Fujimoto and S. S. Gu. A minimalist approach to offline reinforcement learning. *Advances in neural information processing systems*, 34:20132–20145, 2021.

[46] X. Liang, T. Wang, L. Yang, and E. Xing. Cirl: Controllable imitative reinforcement learning for vision-based self-driving. In *Proceedings of the European conference on computer vision (ECCV)*, pages 584–599, 2018.

[47] S. M. LaValle. *Planning algorithms*. Cambridge university press, 2006.

[48] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[49] 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.

[50] A. Cui, S. Casas, K. Wong, S. Suo, and R. Urtasun. Gorela: Go relative for viewpoint-invariant motion forecasting. *arXiv preprint arXiv:2211.02545*, 2022.

[51] W. G. Najm, J. D. Smith, M. Yanagisawa, et al. Pre-crash scenario typology for crash avoidance research. Technical report, United States. National Highway Traffic Safety Administration, 2007.

[52] T. Menzel, G. Bagschik, and M. Maurer. Scenarios for development, test and validation of automated vehicles. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 1821–1827. IEEE, 2018.

[53] H. Weber, J. Bock, J. Klimke, C. Roesener, J. Hiller, R. Krajewski, A. Zlocki, and L. Eckstein. A framework for definition of logical scenarios for safety assurance of automated driving. *Traffic injury prevention*, 20(sup1):S65–S70, 2019.

[54] E. Bronstein, S. Srinivasan, S. Paul, A. Sinha, M. O'Kelly, P. Nikdel, and S. Whiteson. Embedding synthetic off-policy experience for autonomous driving via zero-shot curricula. *arXiv preprint arXiv:2212.01375*, 2022.