Imitation Is Not Enough: Robustifying Imitation with Reinforcement Learning for Challenging Driving Scenarios

Yiren Lu[1], Justin Fu[1], George Tucker[2], Xinlei Pan[1], Eli Bronstein[1], Rebecca Roelofs[2], Benjamin Sapp[1], Brandyn White[1], Aleksandra Faust[2], Shimon Whiteson[1], Dragomir Anguelov[1], Sergey Levine[2,3]

Abstract— Imitation learning (IL) is a simple and powerful way to use high-quality human driving data, which can be collected at scale, to produce human-like behavior. However, policies based on imitation learning alone often fail to sufficiently account for safety and reliability concerns. In this paper, we show how imitation learning combined with reinforcement learning using simple rewards can substantially improve the safety and reliability of driving policies over those learned from imitation alone. In particular, we train a policy on over 100k miles of urban driving data, and measure its effectiveness in test scenarios grouped by different levels of collision likelihood. Our analysis shows that while imitation can perform well in low-difficulty scenarios that are well-covered by the demonstration data, our proposed approach significantly improves robustness on the most challenging scenarios (over 38% reduction in failures). To our knowledge, this is the first application of a combined imitation and reinforcement learning approach in autonomous driving that utilizes large amounts of real-world human driving data.

I. INTRODUCTION

Building an autonomous driving system that is deployable at scale presents many difficulties. First and foremost is the challenge of handling the numerous rare and challenging edge cases that occur in real-world driving. To this end, imitative learning based approaches have been proposed that allow the performance of the method to scale with the amount of data available [1], [2], [3]. While situations that are well represented in the demonstration data are likely to be handled correctly by such a policy, more unusual or dangerous situations that occur only rarely in the data might cause the imitation policy – which has not been explicitly instructed on what constitutes a risky or inappropriate response – to respond unpredictably. The problem is compounded by complex interactions, where human expert driving data in similar scenarios may be scarce and sub-optimal [4].

[1] Waymo Research. [2] Google Research, Brain Team [3] UC Berkeley. Point of contact: maxlu@waymo.com.
Webpage: waymo.com/research/imitation-is-not-enough-robustifying-imitation-with-reinforcement-learning
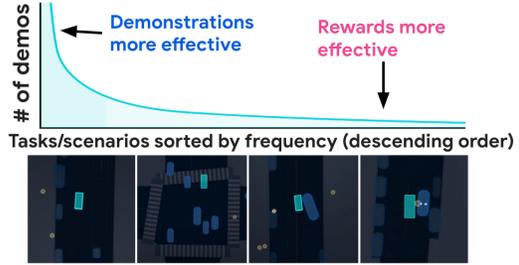
Fig. 1: The demonstration-reward trade-off. As the amount of data for a particular scenario decreases, reward signals become more important for learning. We show a few visual examples representing scenarios with different frequencies.

Reinforcement Learning (RL) has the potential to resolve this by leveraging explicit reward functions that tell the policy what constitutes safe or unsafe outcomes (e.g., collisions). Furthermore, because RL methods train in closed-loop, RL policies can establish causal relationships between observations, actions, and outcomes. This yields policies that are (1) less vulnerable to covariate shifts and spurious correlations commonly seen in open loop IL [5], [6], and (2) aware of safety considerations encoded in their reward function, but which are only implicit in the demonstrations.

However, relying on RL alone, e.g., [7], [8], [9], is also problematic because it heavily depends on reward design, which is an open challenge in autonomous driving [10]. Without accounting for imitation fidelity, driving policies trained with RL may be technically safe but unnatural, and may have a hard time making forward progress in situations that demand human-like driving behavior to coordinate with other agents and follow driving conventions. IL and RL offer complementary strengths: IL increases realism and eases the reward design burden and RL improves safety and robustness, especially in rare and challenging scenarios in the absence of abundant data (Fig. 1).

In this paper we focus on the driving scenarios that are most likely to exhibit safety and reliability concerns, leveraging the difficulty estimation from [11]. Our proposed method, BC-SAC, combines IL and RL with a

simple reward function, and trains on difficult driving scenarios. Difficulty is estimated via a classifier that estimates the likelihood of a collision or near-miss when re-simulated with a pre-trained planning policy. Our proposed reward function enforces safety of the agent, while natural driving behaviors are implicitly learned with IL. The training data comes from a subset of real-world human driving data (over 100k miles of real-world urban driving data) [11]. We demonstrate that this approach substantially improves the safety and reliability of policies learned over imitation alone without compromising on human-like behavior, showing 38% and 40% improvements over pure IL and RL baselines.

The main contributions of our work are: (1) We conduct the first large-scale application of a combined IL and RL approach in autonomous driving utilizing large amounts of real-world urban human driving data (over 100k miles) and a *simple* reward function. (2) We systematically evaluate its performance and baseline performance by slicing the dataset by difficulty, demonstrating that combining IL and RL improves safety and reliability of policies over those learned from imitation alone (over 38% reduction in safety events on the most difficult bucket).

II. RELATED WORK

Learning-based approaches in autonomous driving. We briefly summarize key properties of different learning-based algorithms for planning in Table I. IL was among the earliest and most popular learning-based approaches adopted for deriving driving policies [1], [2], [24], [25], [26], [27]. Controllable models trained with either IL [3], [28] or RL [8] allow the user to specify high-level commands in the form of goals or control signals (e.g., left, right, straight) to combine higher-level route planning with low-level control.

Two drawbacks of IL methods are: (1) open-loop IL (such as the widely used behavioral cloning approach [12], [14], [13], [29], [30]) suffers from covariate shift [5] (which can be addressed with closed-loop training [15], [16]), (2) IL methods lack *explicit* knowledge of what constitutes good driving, such as collision avoidance. RL methods have been proposed that allow the policy to learn from explicit reward signals with closed-loop training and have been applied to tasks such as lane-keeping [31], intersection traversal [32], and lane changing [33]. While these works show the efficacy of RL on specific scenarios, our work analyzes both the large-scale, aggregate performance *and* challenging and safety-critical edge cases that make autonomous driving difficult to deploy in a real-world system.

RL and other closed-loop methods for autonomous driving typically use simulation for training. There are a number of such public environments, which vary in how realistic they are, in particular what drives the simulated agents (e.g., expert-following/log playback [34], [35], [36], [37], intelligent driving model (IDM) [38], or other rule based systems [39] and ML-based agents [38], [40]), and whether scenarios are procedurally generated (e.g., [39], [41], [40]) or initialized from real-world driving scenes [42], [36], [34], [37]). In our experiments, we develop and evaluate in closed-loop on real-world data with other agents following logs.

Combining IL and RL. Methods such as DQfD [22], DDPGfD [43], and DAPG [44] have shown that IL can help RL overcome exploration challenges in domains with known sparse rewards. Offline RL approaches, such as TD3+BC [21] and CQL [20] combine RL objectives with IL ones to regularize Q-learning updates and avoid overestimating out-of-distribution values. Our goal is not to propose a novel algorithmic combination of IL and RL, but rather to leverage this general approach to address challenges in autonomous driving at scale.

Addressing challenging and safety-critical scenarios for autonomous vehicles. [4] learns policies that address long-tail scenarios in autonomous driving by using an ensemble of IL planners combined with model-predictive control. Another approach to improving safety is to augment a learned planner with a rule-based fallback layer that guarantees safety [45], [25]. Our work differs from these approaches, in that we directly incorporate safety awareness into the model learning process through a reward. Our method is also compatible with a fallback layer if needed, although we this is potential future work. Another way to improve robustness of polices is to increase the frequency of negative examples during training. [46] collects failure data that covers various ways an unmanned aerial vehicle can crash, and the combined negative and positive data helps to train more robust policies. [11] investigates the use of curriculum training to improve performance on challenging edge cases. While we also increase the exposure of the policy to challenging scenarios during training, we extend these findings by showing how RL yields outsized improvements on the hardest scenarios.

III. BACKGROUND

A. Markov Decision Processes (MDPs)

In this work, we cast the autonomous driving policies learning problem as a Markov decision process. Following standard formalism, we define an MDP as a tuple $\{\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma, \rho_0\}$. $\mathcal{S}$ and $\mathcal{A}$ denote the state

TABLE I: A comparison of different learning-based approaches to robotic control and autonomous driving.

	Offline Demo	Closed-loop	Rewards	Example Methods
Behavior Cloning (BC)	Expert Demos	No	No	Multipath [12], Precog [13], Trajectron++ [14]
Adversarial Imitation/IRL	Expert Demos	Yes	No	IRL [15], GAIL [16], MGAIL [17]
RL	No	Yes	Yes	DQN [18], SAC [19]
Offline RL	Behavioral Data	No	Yes	CQL [20], TD3+BC [21]
"Imitative" RL	Expert Demos	Yes	Yes	DQfD [22], DAPG [23], BC-SAC (ours)

and action spaces, respectively. $\mathcal{T}$ denotes to transition model. $\mathcal{R}$ represents the reward function, and γ represents the discount factor. ρ_0 represents the initial state distribution. The objective is to find a policy π, a (stochastic) mapping from $\mathcal{S}$ to $\mathcal{A}$, that maximizes the expected discounted sum of rewards, $\pi^* = \max_\pi \mathbb{E}_{\mathcal{T}, \pi, \rho_0} \left[\sum_{t=0}^\infty \gamma^t R(s_t, a_t)\right].$

B. Imitation Learning (IL)

IL constructs an optimal policy by mimicking an expert. We assume an expert (an optimal policy), denoted as π_β, produces a dataset of trajectories $\mathcal{D} = \{s_0, a_0, \cdots, s_N, a_N\}$ through interaction with the environment. The learner's goal is to train a policy π that imitates the π_β. In practice, we only observe the expert states, so we estimate expert actions using inverse dynamics. For example, behavioral cloning (BC) trains the policy via a log-likelihood objective, $\mathbb{E}_{s,a\sim\mathcal{D}}[\log\pi(a|s)]$. Alternatively, closed loop approaches include inverse RL (IRL) [15] and adversarial IL (GAIL [16], MGAIL [17]), which instead aim to more directly match the occupancy measure or state-action visitation distribution between the policy and the expert, rather than indirectly through the conditional action distribution. In principle, this can resolve the covariate shift issue that affects open loop imitation [5].

C. Reinforcement Learning (RL)

RL aims to learn an optimal policy through an iterative, online trial and error process. In this work we use off-policy, value-based RL algorithms such as Q-learning. These methods aim to learn the state-action value function, defined as the expected future return when starting from a particular state and action:

$$Q^\pi(s,a) = \mathbb{E}_{\mathcal{T}, \pi, \rho_0}\left[\sum_{t=0}^\infty \gamma^t R(s_t, a_t)|s_0 = s, a_0 = a\right].$$

In this work, we use an actor-critic method for training continuous control policies. Typical actor-critic methods alternate between training a critic Q to minimize the Bellman error and an actor π to maximize the value function. We use the entropy-regularized updates of Soft Actor-Critic (SAC) [19]:

$$\min_Q \ \mathbb{E}_{s,a,s'\sim\pi}[(Q(s,a) - \hat{Q}(s,a,s'))^2] \quad (1)$$

$$\max_\pi \ \mathbb{E}_{s,a\sim\pi}\left[Q(s,a) + \mathcal{H}(\pi(\cdot|s))\right], \quad (2)$$

where

$$\hat{Q}(s,a,s') = r(s,a) + \gamma\mathbb{E}_{a'\sim\pi}\left[\bar{Q}(s',a') - \log\pi(a'|s')\right] \quad (3)$$

and $\bar{Q}$ denotes a target network that is a copy of the critic through which gradients do not pass.

IV. LEARNING TO DRIVE WITH RL-AUGMENTED BC

We wish to design an approach that benefits from the complementary strengths of IL and RL. Imitation provides an abundant source of learning signal without the need for reward design, and RL addresses the weaknesses of IL in rare and challenging scenarios where data is scarce. Following this intuition, we formulate an objective that utilizes the learning signal from demonstrations where data is abundant and the reward signal where data is scarce. Specifically, we utilize a weighted mixture of the IL and RL objectives:

$$\max_\pi \ \mathbb{E}_{\mathcal{T}, \pi, \rho_0}\left[\sum_{t=0}^\infty \gamma^t R(s_t, a_t)\right] + \lambda\mathbb{E}_{s,a\sim\mathcal{D}}[\log\pi(a|s)]. \quad (4)$$

A. Behavior Cloned Soft Actor-Critic (BC-SAC)

While in principle a variety of RL methods could be combined with IL to optimize Eq. 4, a convenient choice for efficient training is to use actor-critic algorithms, in which case the policy can be optimized with respect to Eq. 4 simply by adding the imitation learning objective to the expected value of the Q-function (i.e., the critic), similarly to DAPG [23] or TD3+BC [21]. Building on the widely used SAC framework, which further adds an entropy regularization objective to the actor, we obtain our full actor objective:

$$\mathbb{E}_{s,a\sim\pi}[Q(s,a) + \mathcal{H}(\pi(\cdot|s))] + \lambda\mathbb{E}_{s,a\sim\mathcal{D}}[\log\pi(a|s)].$$

The critic update remains the same as in SAC, outlined in Eq 1. With the appropriate setting of λ, this objective encourages the policy to mimic the expert data when it is within the data distribution $\mathcal{D}$. However, in out-of-distribution states the policy primarily relies on reward to learn. Fig. 2 visualizes this concept.
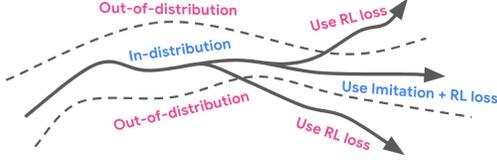
Fig. 2: Different objective influence. For in-distribution states, both IL and RL objectives provide learning signal. For out-of-distribution states, the RL objective dominates.

B. Reward Function

While designing a reward function to capture "good" driving behavior is an open-challenge [10], we can side-step this issue by relying on the imitation learning loss to primarily guide the policy, while the simple reward function only needs to encode safety constraints. To this end, we use a combination of collision and off-road distances as our reward signal. The collision reward is

$$R_{\text{collision}} = \min(d_{\text{collision}} - d_{\text{c\_offset}}, 0), \qquad (5)$$

where $d_{\text{collision}}$ is the Euclidean distance in meters of the closest points between the ego vehicle and a nearest bounding box of other vehicles; $d_{\text{c\_offset}}$ (default 1.0) is an offset added to encourage the vehicle to keep a distance from nearby objects. The off-road reward is

$$R_{\text{off-road}} = \text{clip}(-d_{\text{o\_offset}} - d_{\text{to-edge}}, -2.0, 0.0), \qquad (6)$$

where $d_{\text{to-edge}}$ is the distance in meters of the vehicle to the nearest road edge (negative being on-road, positive being off-road). $d_{\text{o\_offset}}$ (default 1.0) is an offset to encourage the vehicle to keep a distance to road edge. We combine the rewards additively, such that $R = R_{\text{collision}} + R_{\text{off-road}}$.

C. Forward and Inverse Vehicle Dynamics Models

We update the vehicle's state using the kinematic bicycle dynamics model [47], which computes the vehicle's next pose (x, y, θ) given a steering and acceleration action $a = (a_{\text{steer}}, a_{\text{accel}})$. In order to obtain expert actions for imitation learning, we use an inverse dynamics model to solve for the actions that would have achieved the same states as the logged trajectories in our dataset. These expert actions are found by minimizing the MSE of the corners' (x, y) positions between the inferred state $\mathcal{T}(s_t, a_t)$ and ground-truth next state s_{t+1}.

D. Model Architecture

We use a dual actor-critic architecture similar to TD3 and SAC [48], [19]: the main components are an actor network $\pi(a|s)$, a double Q-critic network $Q(s, a)$ and a target double Q-critic network $\bar{Q}(s, a)$. Each network has a separate Transformer observation encoder described in [49] that encodes features including all vehicle states, road-graph points, traffic lights signals, and route goals. The actor network outputs a tanh-squashed diagonal Gaussian distribution parameterized by a mean μ and variance σ.

E. Training on Difficult Examples

The performance of learning-based methods strongly depends on the training data distribution, especially in safety-critical settings with long-tail distributions [50], [51], [45], [52], [53]). Autonomous driving falls in this category: most scenarios are mundane, but a sizable minority of scenarios have critical safety concerns. Following [11], which demonstrated that training on more difficult examples results in better performance than using unbiased training distributions, we explore how the training distribution affects the method performance.

V. EXPERIMENTS

A. Experimental Setup

Datasets. We use a dataset (denoted **All**) consisting of over 100k miles of expert driving trajectories, split into 10 second segments, collected from a fleet of vehicles operating in San Francisco (SF) [11]. We divide these segments into 6.4 million for training and 10k for testing. Trajectories from the same vehicle operating on the same day are stored in the same partition to avoid train-test leakage. The trajectories, which are sampled at 15 Hz, contain features describing the autonomous vehicle (AV) state and the state of the environment as measured by the AV's perception system. We use the *difficulty model* described by [11] as a proxy for measuring the rarity of events, since it is difficult to directly construct a scenario-level out-of-distribution estimator, and challenging scenarios are generally less frequent. Given a run segment, the difficulty model predicts whether a segment will result in a collision or near-miss when re-simulated with an internal AV planner. We trained the difficulty model in a supervised manner using cross-entropy loss on a dataset consisting of 5.6k positive examples and 80k negative examples, with binary human labels. We create the **Top1**, **Top10**, and **Top50** subsets by selecting the top 1% (40k train, 1.2k test), 10% (400k train, 19k test), and 50% (2 million train, 66k test) percentiles of difficulty model scores from a chronologically separate dataset of 4 million segments, respectively.

Simulation. As mentioned in Sec. IV-C, vehicle dynamics are modeled using a 2D bicycle dynamics model. The behavior of other vehicles and pedestrians in the scene are replayed from the logs (log-playback), similarly to [34], [35], [36]. While this means that agents are
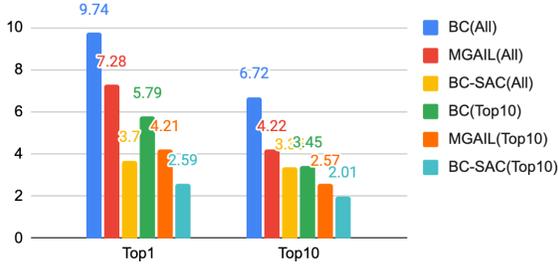
Fig. 3: Failure rates on the most challenging evaluation sets: Top1 and Top10 (lower is better, with training on All and Top10). BC-SAC consistently achieves the lowest error rates.
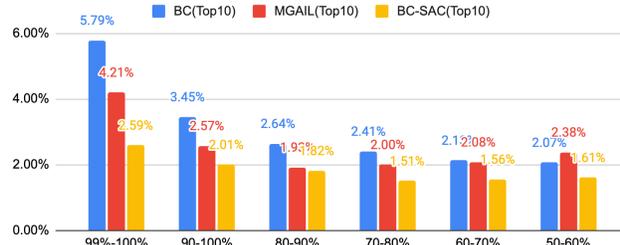
Fig. 4: Failure rates of BC, MGAIL, and BC-SAC across scenarios of varying difficulty levels (50%-100%, lower is better). While all methods perform worse as the evaluation dataset becomes more challenging, BC-SAC always performs best and shows the least degradation.

non-reactive, it ensures that the behavior of other agents is human-like, and the inclusion of imitative losses discourages the learned policy to deviate too far from the logs, which would cause the log-playback agents to become unrealistic. We also use short segments of 10s to mitigate pose divergence.

Baselines. We compare our method to both open-loop (*BC* [1]) and closed-loop (*MGAIL* [17]) imitative methods. The latter takes advantage of closed loop training and the differentiability of the simulator dynamics. For completeness, we also include a SAC baseline to represent an RL-only approach.

Metrics. We evaluate agents using two metrics:

1) *Failure Rate*: Percentage of the run segments that have at least one *Collision* or *Off-road* event at any timestep. *Collision* is true if the bounding box of the ego vehicle intersects with a bounding box of another object. *Off-road* is true if the bounding box of the ego vehicle deviates from the drivable surface according to the map.

2) *Route Progress Ratio*: Ratio of the distance traveled along the route by the policy compared to the expert demonstration. We project the ego vehicle's state onto the route and compute the total length from the start of the route.

B. Results

We evaluate the baseline methods (BC, MGAIL, SAC) and our method (BC-SAC) trained on several subsets of the training dataset (All, Top10, and Top1), and evaluate against subsets of the evaluation set (Top1, Top10, Top50, All) in Table II. All configurations are evaluated with three random seeds, reporting mean and standard deviation. Previously, [11] showed that training MGAIL on Top10 yields similar performance with training on All. Similarly, we find that all methods perform best when trained on Top10. Notably, BC trained on Top1 performs significantly worse compared to training on All or Top10, which reflects the fact that imitation learning methods rely on large amounts of data to implicitly infer driving preferences. In contrast, BC-SAC performs robustly when trained on Top1. Given that all methods perform best when trained on Top10, we focus on that setting in the following subsections.

BC-SAC comparison to imitation methods (BC, MGAIL) in the challenging scenarios. Figure 4 compares BC-SAC against BC and MGAIL across the evaluation dataset slices according to difficulty levels. BC-SAC achieves better performance overall, especially in the more challenging slices where the performance of both BC and MGAIL substantially degrade. Additionally, BC-SAC has the lowest variance across scenarios of varying difficulty in performance ($\sigma = 0.37$) vs. BC ($\sigma = 1.29$) and MGAIL ($\sigma = 0.78$).

BC-SAC comparison to RL-only training (SAC). In all configurations, BC-SAC outperforms SAC in terms of safety metrics (Table II), likely because BC-SAC also utilizes learning signal from large amount of demonstrations. SAC generates actions that deviate significantly from the demonstrations with more boundary action values yielding unnatural (more swerves) and uncomfortable (abrupt acceleration) driving behavior (Figure 5). With a BC loss, BC-SAC generates an action distribution similar to the logs.

Reward shaping and RL / IL weights. We conduct a set of ablation studies to answer how the form of the reward function and the weights on the RL and imitation components influence final performance. We use a smaller dataset constructed by sampling 10% of the Top10 data and compare: (1) our full reward vs. a discrete binary reward (Fig 6 Right), (2) off-road and collision reward term weights (Fig 6 Left), (3) off-road and collision offset parameters (Fig 7), and (4) the weight on the RL and IL terms in the objective (Fig 8). The results indicate that the proposed shaped reward improves overall performance over the simpler sparse reward with an appropriate choice of reward parameters,

Method	Training	Top1 (%)	Top10 (%)	Top50 (%)	All (%)	Route Progress Ratio, All(%)
BC	All	9.74±0.49	6.72±0.47	5.14±0.39	4.35±0.27	99.00±0.39
MGAIL	All	7.28±0.98	4.22±0.77	3.40±0.97	**2.48±0.29**	**99.55±1.91**
SAC	All	5.29±0.66	4.64±1.08	4.12±0.74	6.66±0.44	77.82±8.21
BC-SAC	All	**3.72±0.62**	**2.88±0.23**	**2.64±0.21**	3.35±0.31	95.26±8.64
BC	Top10	5.79±0.82	3.45±0.72	2.71±0.57	3.64±0.31	**98.06±0.18**
MGAIL	Top10	4.21±0.95	2.57±0.52	2.20±0.52	**2.45±0.35**	96.57±1.19
SAC	Top10	4.33±0.47	4.11±0.63	3.66±0.47	5.60±0.86	71.05±2.47
BC-SAC	Top10	**2.59±0.31**	**2.01±0.29**	**1.76±0.20**	2.81±0.26	87.63±0.58
BC	Top1	7.66±1.13	7.84±0.92	6.63±0.78	6.85±0.65	**94.10±1.00**
MGAIL	Top1	4.24±0.95	3.16±0.43	2.74±0.46	3.79±0.46	93.10±11.72
SAC	Top1	4.15±0.31	3.87±0.12	3.46±0.16	5.98±1.03	75.63±2.19
BC-SAC	Top1	**3.61±0.87**	**2.96±1.11**	**2.69±0.87**	3.38±0.48	75.00±17.21

TABLE II: Failure rates (lower is better) and progress ratios (higher is better) of BC-SAC and baselines on different training/evaluation subsets.

(a) SAC acceleration (b) SAC tire angles

(c) BC-SAC acceleration (d) BC-SAC tire angles

Fig. 5: Marginal action distributions. SAC/BC-SAC (orange) vs logs (blue).
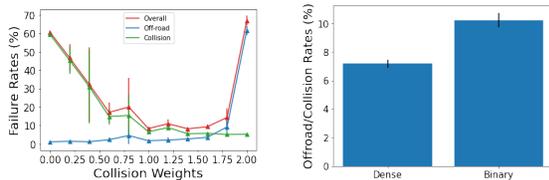
Fig. 6: **Left**: Off-road / collision weights. Off-road weight and collision weight add up to 2.0. The x-axis is the collision weight. A balanced choice of off-road and collision weights lead to the best performance. **Right**: Dense vs binary rewards. Binary reward is defined as -1 when a safety event happens and 0 otherwise. Dense rewards lead to fewer safety events.
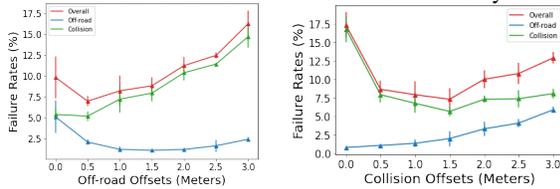
Fig. 7: Off-road offset $d_{\text{o\_offset}}$ and collision offset $d_{\text{c\_offset}}$ ablations. A small amount of offsets improves overall performance.

and a balance between imitation and RL terms leads to the best performance.

Progress-safety balance. While our work focuses on safety-critical scenarios, in Fig 8 Right, we show that introducing a small amount of a progress reward leads to significantly more progress without major regressions in safety metrics. However, large progress rewards lead to degradation in performance.

In-depth failure analysis. Table III presents a detailed analysis of failure modes on a set of 80 sampled scenar-

Method	CLIP	COLL	OFF	RED	LAN	DIV
BC-SAC	8	7	2	1	7	15
MGAIL	16	8	8	2	0	6

TABLE III: Failure frequency categorizations, per type, incurred by BC-SAC and MGAIL on a small sample set (N=80). BC-SAC generally has fewer direct collisions and off-road events, but has a greater frequency of being hit by other objects.

Fig. 8: **Left** Imitation weights (log-scale) vs failure rates. **Right** Progress reward weights (log-scale) vs policy evaluation performance: safety event rate and route progress ratio.

ios from the Top1 and Top10 buckets. We categorize failures into 6 broad buckets. CLIP (clipping): small collisions that occur when a vehicle collides with an object on the side while moving. OFF (off-road): failures when the agent drives off the road. LAN (bad lane): an agent encroaches into another lane, either the wrong lane or a bad merge, which results in a collision. COLL (collision): collision where the planning agent is at fault and drives into another vehicle. RED (red light): red light violations that result in collisions. Finally, DIV (log divergence): collisions where a sim agent collides with the planning agent due to divergence from the logs.

Overall, MGAIL tends to have more clipping collisions and off-road events. Fig 9 shows two of the cases where RL improves over IL. We hypothesize that our method improves in these cases because MGAIL, as an imitation method, lacks an explicit penalty for collisions, and thus is not sensitive to small collisions during otherwise realistic behavior. On the other hand, the collisions encountered by BC-SAC tend to be cases where the collision is not directly the result of the AV planner's action, but the planner diverges from the logs in a way such that it is hit by other vehicles. Because BC-SAC also is not explicitly rewarded for following traffic rules (though it inherits this behavior via imitation), we also see a small amount of failures due to that.
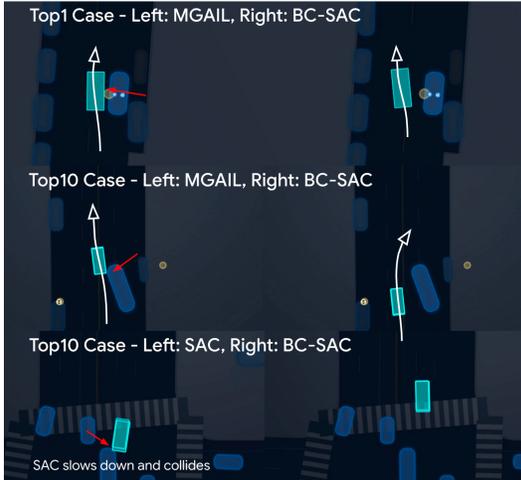
Fig. 9: Visualizations of a few scenarios where BC-SAC improves over imitation (MGAIL) and RL-only (SAC). The cyan car is controlled. **Example 1**: MGAIL collides with a pedestrian exiting a double parked car while BC-SAC leaves enough clearance. **Example 2**: MGAIL does not provide sufficient clearance and collides with the incoming vehicle. **Example 3**: SAC slows down in an intersection resulting in an rear collision. BC-SAC maintains a proper speed profile through the intersection without a collision.

VI. CONCLUSIONS

We presented a method for robust autonomous driving in challenging driving scenarios, that combines imitation learning with RL (BC-SAC), paired with a simple safety reward, and trained on large datasets of real-world driving. Overall, the method significantly improves safety and reliability in challenging scenarios, resulting in more than 38% reduction in safety events of the most difficult scenarios compared to IL-only and RL-only baselines. Our extensive experiments examined the roles of training datasets, reward shaping and IL / RL objective terms. BC-SAC inherits implicit human-like driving behaviors from imitation, while RL is a fail-safe for handling out-of-distribution safety scenarios. Similarly to the IL-only settings, training on the top 10% of the most challenging scenarios yields the most robust performance in the combined IL and RL setting. While this work mainly focused on optimizing safety-related rewards, a natural extension is to incorporate other factors into the objective, such as progress, traffic rule adherence, and passenger comfort. Besides the reward function, this approach does not account for unexpected behavior of other agents in response to out-of-distribution actions on the part of the ego vehicle, and it still requires heuristically choosing the tradeoff between the IL and RL objectives. A promising future work direction would be to enable reactive sim agents for training and evaluation and to extend the approach to enforce safety as an explicit constraint, perhaps in combination with methodology to mitigate distributional shift.

REFERENCES

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

[2] 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.

[3] F. Codevilla, M. Miiller, A. López, V. Koltun, and A. Dosovit-skiy, "End-to-end driving via conditional imitation learning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1–9.

[4] W. Zhou, Z. Cao, N. Deng, X. Liu, K. Jiang, and D. Yang, "Long-tail prediction uncertainty aware trajectory planning for self-driving vehicles," *arXiv preprint arXiv:2207.00788*, 2022.

[5] 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*. JMLR Workshop and Conference Proceedings, 2011, pp. 627–635.

[6] P. De Haan, D. Jayaraman, and S. Levine, "Causal confusion in imitation learning," *Advances in Neural Information Processing Systems*, vol. 32, 2019.

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

[8] 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)*, 2018, pp. 584–599.

[9] Z. Zhang, A. Liniger, D. Dai, F. Yu, and L. Van Gool, "End-to-end urban driving by imitating a reinforcement learning coach," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 15 222–15 232.

[10] W. B. Knox, A. Allievi, H. Banzhaf, F. Schmitt, and P. Stone, "Reward (mis) design for autonomous driving," *arXiv preprint arXiv:2104.13906*, 2021.

[11] 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," in *6th Annual Conference on Robot Learning*, 2022.

[12] Y. Chai, B. Sapp, M. Bansal, and D. Anguelov, "Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction," *arXiv preprint arXiv:1910.05449*, 2019.

[13] N. Rhinehart, R. McAllister, K. M. Kitani, and S. Levine, "PRECOG: prediction conditioned on goals in visual multi-agent settings," *CoRR*, vol. abs/1905.01296, 2019.

[14] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Multi-agent generative trajectory forecasting with heterogeneous data for control," *arXiv preprint arXiv:2001.03093*, 2020.

[15] A. Y. Ng and S. J. Russell, "Algorithms for inverse reinforcement learning," in *Proceedings of 17th International Conference on Machine Learning, 2000*, 2000, pp. 663–670.

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

[17] N. Baram, O. Anschel, and S. Mannor, "Model-based adversarial imitation learning," *arXiv preprint arXiv:1612.02179*, 2016.

[18] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, "Human-level control through deep reinforcement learning," *nature*, vol. 518, no. 7540, pp. 529–533, 2015.

[19] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," *CoRR*, vol. abs/1801.01290, 2018.

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

[21] S. Fujimoto and S. S. Gu, "A minimalist approach to offline reinforcement learning," *Advances in neural information processing systems*, vol. 34, pp. 20 132–20 145, 2021.

[22] 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*, vol. 32, no. 1, 2018.

[23] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning complex dexterous manipulation with deep reinforcement learning and demonstrations," in *Robotics: Science and Systems*, 2018.

[24] Z. Zhang, A. Liniger, D. Dai, F. Yu, and L. Van Gool, "End-to-end urban driving by imitating a reinforcement learning coach," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15 222–15 232.

[25] M. Vitelli, Y. Chang, Y. Ye, A. Ferreira, M. Wołczyk, B. Osiński, M. Niendorf, H. Grimmett, Q. Huang, and A. Jain, "Safetynet: Safe planning for real-world self-driving vehicles using machine-learned policies," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 897–904.

[26] N. Nayakanti, R. Al-Rfou, A. Zhou, K. Goel, K. S. Refaat, and B. Sapp, "Wayformer: Motion forecasting via simple & efficient attention networks," *arXiv preprint arXiv:2207.05844*, 2022.

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

[28] N. Rhinehart, R. McAllister, and S. Levine, "Deep imitative models for flexible inference, planning, and control," *arXiv preprint arXiv:1810.06544*, 2018.

[29] M. Liang, B. Yang, R. Hu, Y. Chen, R. Liao, S. Feng, and R. Urtasun, "Learning lane graph representations for motion forecasting," *arXiv preprint arXiv:2007.13732*, 2020.

[30] J. Ngiam, B. Caine, V. Vasudevan, Z. Zhang, H. L. Chiang, J. Ling, R. Roelofs, A. Bewley, C. Liu, A. Venugopal, D. Weiss, B. Sapp, Z. Chen, and J. Shlens, "Scene transformer: A unified multi-task model for behavior prediction and planning," *CoRR*, vol. abs/2106.08417, 2021.

[31] A. Kendall, J. Hawke, D. Janz, P. Mazur, D. Reda, J.-M. Allen, V.-D. Lam, A. Bewley, and A. Shah, "Learning to drive in a day," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8248–8254.

[32] D. Isele, R. Rahimi, A. Cosgun, K. Subramanian, and K. Fujimura, "Navigating occluded intersections with autonomous vehicles using deep reinforcement learning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 2034–2039.

[33] P. Wang, C.-Y. Chan, and A. de La Fortelle, "A reinforcement learning based approach for automated lane change maneuvers," in *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2018, pp. 1379–1384.

[34] 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.

[35] P. Kothari, C. Perone, L. Bergamini, A. Alahi, and P. Ondruska, "Drivergym: Democratising reinforcement learning for autonomous driving," *arXiv preprint arXiv:2111.06889*, 2021.

[36] Q. Li, Z. Peng, L. Feng, Q. Zhang, Z. Xue, and B. Zhou, "Metadrive: Composing diverse driving scenarios for generalizable reinforcement learning," *IEEE transactions on pattern analysis and machine intelligence*, 2022.

[37] V. Lioutas, J. W. Lavington, J. Sefas, M. Niedoba, Y. Liu, B. Zwartsenberg, S. Dabiri, F. Wood, and A. Scibior, "Critic sequential monte carlo," *arXiv preprint arXiv:2205.15460*, 2022.

[38] H. Caesar, J. Kabzan, K. S. Tan, W. K. Fong, E. Wolff, A. Lang, L. Fletcher, O. Beijbom, and S. Omari, "nuplan: A closed-loop ml-based planning benchmark for autonomous vehicles," *arXiv preprint arXiv:2106.11810*, 2021.

[39] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Conference on robot learning*. PMLR, 2017, pp. 1–16.

[40] K. Ramamohanarao, H. Xie, L. Kulik, S. Karunasekera, E. Tanin, R. Zhang, and E. B. Khunayn, "Smarts: Scalable microscopic adaptive road traffic simulator," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 8, no. 2, pp. 1–22, 2016.

[41] E. Leurent, "An environment for autonomous driving decision-making," https://github.com/eleurent/highway-env, 2018.

[42] W. Zhan, L. Sun, D. Wang, H. Shi, A. Clausse, M. Naumann, J. Kummerle, H. Konigshof, C. Stiller, A. de La Fortelle *et al.*, "Interaction dataset: An international, adversarial and cooperative motion dataset in interactive driving scenarios with semantic maps," *arXiv preprint arXiv:1910.03088*, 2019.

[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. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning complex dexterous manipulation with deep reinforcement learning and demonstrations," *arXiv preprint arXiv:1709.10087*, 2017.

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

[46] D. Gandhi, L. Pinto, and A. Gupta, "Learning to fly by crashing," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 3948–3955.

[47] R. Rajamani, *Vehicle dynamics and control*. Springer Science & Business Media, 2011.

[48] S. Fujimoto, H. Hoof, and D. Meger, "Addressing function approximation error in actor-critic methods," in *International conference on machine learning*. PMLR, 2018, pp. 1587–1596.

[49] A. Jaegle, F. Gimeno, A. Brock, O. Vinyals, A. Zisserman, and J. Carreira, "Perceiver: General perception with iterative attention," in *International conference on machine learning*. PMLR, 2021, pp. 4651–4664.

[50] J. Frank, S. Mannor, and D. Precup, "Reinforcement learning in the presence of rare events," in *Proceedings of the 25th international conference on Machine learning*, 2008, pp. 336–343.

[51] N. Kalra and S. M. Paddock, *Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability?* RAND Corporation, 2016.

[52] S. Paul, K. Chatzilygeroudis, K. Ciosek, J.-B. Mouret, M. Osborne, and S. Whiteson, "Alternating optimisation and quadrature for robust control," in *AAAI Conference on Artificial Intelligence*, 2018.

[53] S. Paul, M. A. Osborne, and S. Whiteson, "Fingerprint policy optimisation for robust reinforcement learning," in *International Conference on Machine Learning*, 2019.

[54] L. Espeholt, H. Soyer, R. Munos, K. Simonyan, V. Mnih, T. Ward, Y. Doron, V. Firoiu, T. Harley, I. Dunning *et al.*, "Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures," in *International conference on machine learning*. PMLR, 2018, pp. 1407–1416.

[55] E. Bronstein, M. Palatucci, D. Notz, B. White, A. Kuefler, Y. Lu, S. Paul, P. Nikdel, P. Mougin, H. Chen *et al.*, "Hierarchical model-based imitation learning for planning in autonomous driving," *arXiv preprint arXiv:2210.09539*, 2022.

A. IL + RL Distributed Actor-Learner Training Architecture
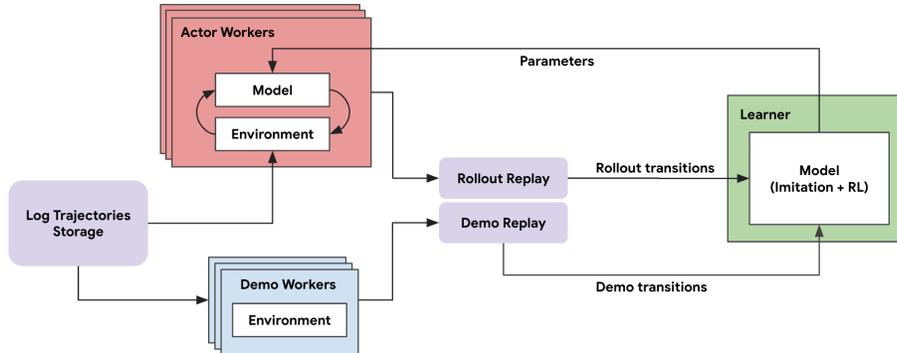
Fig. 10: IL + RL distributed actor-learner training architecture. We extend the distributed IMPALA architecture [54] with additional demo rollout workers and a demo replay buffer, which produce rollout transitions in the same format as the actor workers. The learner worker samples from both the rollout replay buffer and the demo replay buffer to perform training updates in an off-policy manner.

B. Additional Details on Model Architectures and Hyper-parameters Settings

We use a dual actor-critic architecture similar to TD3 and SAC [48], [19]: each of the main components, actor network $\pi(a|s)$, double Q-critic network $Q(s,a)$ and target double Q-critic network $\bar{Q}(s,a)$, has a separate Transformer observation encoder described in [55], and the encoder embedding is fed to a $(256, 256)$ fully connected head. The actor network outputs a $\tanh$-squashed diagonal Gaussian distribution parameterized by a mean μ and variance σ.

We train the BC-SAC algorithm with the following hyper-parameters: the actor learning rate is 1e-4, the critic learning rate is 1e-4, the imitation learning rate is 5e-5, the batch size is 64, and the reward discount ratio is 0.92. The sample-to-insert ratio for replay is 8, which is the average number of times the learner should sample each item in the replay buffer during the item's entire lifetime. In practice, instead of performing a combined gradient step of both the IL and RL objectives, we alternate the training steps between IL and RL with different update frequencies. For every 8 RL updates, we update with IL loss for one time. The hyper-parameters are found by performing grid-search.

For SAC, we use the same network design and hyper-parameters as in BC-SAC, except that it does not perform IL step.

For BC, we discretize the 2d action space (steer, acceleration) in to $31 \times 7 = 217$ actions with the same underlying dynamics model. We use a similar network design for BC as in BC-SAC's actor network with a Softmax prediction head representing probabilities of the discrete actions. We use the cross-entropy loss with a learning rate of 1e-4 and batch size of 256 for training.

For MGAIL, we follow the network design and hyper-parameters setting presented in [55].