

PRIBOOT: A New Data-Driven Expert for Improved Driving Simulations

Daniel Coelho, Miguel Oliveira, Vítor Santos, and Antonio M. López

Abstract—The advancement of Autonomous Driving (AD) systems in simulated environments like CARLA is essential for real-world automotive technology. CARLA’s new Leaderboard 2.0 poses greater challenges, and current AD methods have struggled due to insufficient ground truth data. Human driving logs from CARLA and previous expert agents like Autopilot and Roach have proven inadequate under these challenging conditions. To address this, we introduce PRIBOOT, an expert agent that leverages limited human logs with privileged information. We developed a novel BEV representation tailored for these complex driving scenarios, processed as an RGB image to enable transfer learning, rather than using masks. Additionally, we propose the Infraction Rate Score (IRS), a new metric for a balanced assessment of driving performance over extended routes. PRIBOOT is the first model to achieve a 75% Route Completion (RC) in Leaderboard 2.0, along with a Driving Score (DS) of 20% and an IRS of 45%. PRIBOOT enables researchers to generate extensive datasets, potentially resolving data availability issues in this benchmark.

I. INTRODUCTION

Autonomous Driving (AD) is a key technological advancement with the potential to transform transportation, improve road safety, and redefine urban environments [1], [2]. Despite its potential, developing fully autonomous vehicles involves significant challenges. These include integrating diverse sensors, processing complex data, making real-time decisions, and addressing ethical issues. Such vehicles must operate reliably in unpredictable conditions, requiring advanced systems capable of handling a wide range of scenarios [3]. Real-world testing of autonomous vehicles, while necessary, is often expensive, risky, and encumbered by ethical dilemmas.

Simulations serve as a critical complement to real-world testing, providing a safe and controlled environment that replicates complex real-world scenarios without the associated costs and risks [3], [4]. This enhances the development of autonomous driving technologies by allowing preliminary

Daniel Coelho is with the Department of Mechanical Engineering, University of Aveiro, 3810-193 Aveiro, Portugal, and with the Intelligent System Associate Laboratory (LASI), Institute of Electronics and Informatics Engineering of Aveiro (IEETA), University of Aveiro, 3810-193 Aveiro, Portugal (e-mail: danielsilveira@ua.pt).

Miguel Oliveira is with the Department of Mechanical Engineering, University of Aveiro, 3810-193 Aveiro, Portugal, and with the Intelligent System Associate Laboratory (LASI), Institute of Electronics and Informatics Engineering of Aveiro (IEETA), University of Aveiro, 3810-193 Aveiro, Portugal (e-mail: mriem@ua.pt)

Vítor Santos is with the Department of Mechanical Engineering, University of Aveiro, 3810-193 Aveiro, Portugal, and with the Intelligent System Associate Laboratory (LASI), Institute of Electronics and Informatics Engineering of Aveiro (IEETA), University of Aveiro, 3810-193 Aveiro, Portugal (e-mail: vitor@ua.pt)

Antonio M. López is with the Department of Computer Science, Computer Vision Center (CVC), Universitat Autònoma de Barcelona (UAB), Spain. (e-mail:antonio@cvc.uab.cat)

testing and refinement in simulations, reserving real-world trials for the final stages of development. Moreover, in these simulated environments, it is possible to leverage privileged information, otherwise not available in the real-world, to create expert systems that can provide demonstrations, further enriching the development process.

Among various open-source AD simulators, CARLA [3] is often listed as the premier choice [4], [5]. CARLA offers a suite of essential features for realistic and effective simulation of driving scenarios. These include comprehensive environmental conditions, detailed vehicle models, and a wide array of sensors, making it an ideal platform for advanced AD research and development.

To accelerate innovation, CARLA introduced the CARLA Leaderboard 1.0¹ benchmark, designed to assess the driving proficiency of autonomous agents within realistic traffic scenarios. Despite the complex scenarios presented in the benchmark, various methods such as ReasonNet [6], InterFuser [7], and TCP [8] have consistently reported high performance over the years. Notably, the CARLA Autopilot, a rule-based agent, achieved near-perfect performance. This underscores the benchmark’s capacity to be effectively mastered using current technologies. Building on this foundation, CARLA Leaderboard 2.0² introduces even more complex and challenging scenarios, such as obstacles in the lane and parking exits. These novel scenarios significantly increase the difficulty level, challenging the limits of existing autonomous driving systems. To this date, all approaches tested on the Leaderboard 2.0 benchmark have shown very poor performance, with the highest Route Completion (RC) reaching 15%, and the highest DS reaching 1% [9]. We believe that the primary reason for this notable decline in performance can be attributed to the insufficiency of available training data. CARLA provides a set of human driving logs from a few route scenarios, but these are insufficient for training models that rely on sensor information. In Leaderboard 1.0, researchers could leverage online experts like CARLA Autopilot or Roach [10] to generate demonstrations. However, in Leaderboard 2.0, these experts are either markedly less effective or completely ineffective, as we will show in Section II.

This paper presents a method to address the challenges posed by the limited training data availability. The driving logs from CARLA, while insufficient alone for models requiring sensor inputs, become significantly more useful when

¹<https://leaderboard.carla.org/#leaderboard-10>

²<https://leaderboard.carla.org/>

combined with privileged information from the simulator, specifically Bird’s Eye View (BEV). This integration effectively simplifies the complexity of the benchmark. Employing this strategy, we apply Imitation Learning techniques to develop PRIBOOT (**P**rileged **I**nformation **B**ootstrapping), an expert agent capable of navigating the demanding scenarios presented in Leaderboard 2.0. PRIBOOT utilizes privileged information to master these scenarios, subsequently enabling the generation of extensive datasets or providing online demonstrations. Although PRIBOOT was designed to address the challenges of Leaderboard 2.0, it is also applicable to any other CARLA benchmark.

Overall, we summarize our main contributions as follows:

- Introduce PRIBOOT, an expert agent that effectively leverages privileged information and limited data for model training, marking the first instance of achieving significant performance milestones on the CARLA Leaderboard 2.0;
- Develop a tailored Bird’s Eye View (BEV) representation to effectively address the complex driving scenarios encountered in CARLA Leaderboard 2.0.
- Process the BEV as an RGB image rather than a set of masks. This facilitates the application of transfer learning techniques, which significantly enhance model performance and efficiency, particularly in the context of limited data availability;
- Introduce Infraction Rate Score (**IRS**), a novel evaluation metric that considers infractions per kilometer rather than the total number of infractions. This metric is designed to complement the Driving Score (**DS**) by providing a more detailed assessment of driving behavior over long routes.

The source code of PRIBOOT is available at <https://github.com/DanielCoelho112/priboot>.

II. RELATED WORK

This section is divided into two topics: the application of expert agents in AD, and an overview of all expert agents utilized in CARLA.

A. Application of Experts in Autonomous Driving

In recent years, the field of AD has seen significant advancements through the application of online experts [8], [10]–[12]. A notable example of this is showcased in [11], where the utility of online experts is demonstrated in real-world, high-speed off-road driving scenarios. In their approach, an initial expert system equipped with expensive sensors is developed using a combination of hand-engineered components. This expert system then serves as a reference model, providing high-quality driving demonstrations to train a student model, which operates using more affordable sensors.

Building on the foundational use of online experts, the transfer of knowledge from the expert to the student model can be accomplished through various methodologies. One prevalent method involves the creation of offline datasets,

which are subsequently employed for offline Imitation Learning (IL) [6], [13] or Reinforcement Learning from Demonstrations (RLfD) [14]. These approaches are particularly valuable in scenarios where direct interaction with the environment is either too costly or filled with risks. However, a significant challenge with using offline datasets is the potential for a distribution shift [15]. To address the issue of distribution shift, an alternative strategy is online IL [16], where the student actively explores the environment while the teacher provides on-demand supervision. This method helps to align the student’s learning experience more closely with the actual operational environment, thereby reducing the impact of distribution shift. Nevertheless, this approach still relies heavily on the quality of the data provided by the expert. If the expert’s behavior is not optimal, the student is likely to inherit these imperfections [10]. To further refine this process and overcome the limitations of potentially suboptimal expert data, another innovative approach is Reinforcement Learning from Online Demonstrations (RLfOLD) [17]. This technique merges the benefits of Online IL with the principles of RL, tackling both the issue of distribution shift and the challenge of learning from a suboptimal expert.

B. Experts in CARLA

CARLA incorporates a built-in expert system known as Autopilot, which relies on a series of handcrafted rules that utilize the internal state of the simulator for navigation [3]. In Leaderboard 1.0, Autopilot demonstrated commendable performance, contributing significantly to data collection for top-ranked methods such as ReasonNet [6], which relies on datasets generated by this expert. However, the transition to Leaderboard 2.0 reveals a stark contrast in the efficacy of the Autopilot system. As detailed in Section IV, the performance of Autopilot is markedly diminished in the more demanding scenarios of this updated benchmark. The primary challenge lies in the inherent limitations of a rule-based navigation framework, which struggles to adapt to the complex and dynamic driving conditions presented in Leaderboard 2.0, such as yielding to emergency vehicles or overtaking obstacles in the lane.

While Autopilot has shown competent performance in earlier benchmarks, the adoption of learning-based experts presents distinct advantages [10], [16], [18]. These methods usually decouple the perception from planning, which simplifies the training process. Typically, such methods leverage privileged information from the simulator to bypass the need for complex perception systems, focusing instead on training the planning module directly. For example, LBC [16] and SAM [18] replace the perception module with simulator-derived privileged information, and then train the planning component using IL based on demonstrations provided by the Autopilot. To ensure effective knowledge transfer from the expert to the student, LBC aims to minimize the output differences between them, whereas SAM focuses on aligning the latent representations of both models. These learning-based experts have been assessed in straightforward benchmarks, such as the NoCrash benchmark [13]. This




Fig. 1. BEV used in PRIBOOT. This representation was built upon Roach and LBC, with critical adaptations to tailor it for the complexities of Leaderboard 2.0. (a) Differentiates emergency vehicles (dark blue) from regular vehicles (blue). (b) Introduces an additional class for construction objects, illustrated in orange. Additionally, both images depict a simple method to represent motion with directional arrows, illustrated in green. The images with a white frame provide a zoomed-in view to highlight specific details.

benchmark, with its relatively simple navigation and collision avoidance tasks, represents a significantly lesser challenge than even the earlier Leaderboard 1.0, and far less demanding than the complexities encountered in Leaderboard 2.0.

More recently, the Roach expert [10] was introduced and has since become the most utilized expert in Leaderboard 1.0 [8], [19]. Roach processes inputs using a Bird’s Eye View (BEV) image that encapsulates roads, lanes, routes, vehicles, pedestrians, traffic lights, and stop signs. This information is then processed using a model-free Reinforcement Learning (RL) algorithm to generate vehicle control commands. While Roach has demonstrated impressive results in Leaderboard 1.0, its applicability to Leaderboard 2.0 is questionable without significant modifications. Several challenges hinder the transition of Roach to the more demanding Leaderboard 2.0. Firstly, their BEV implementation struggles with scalability issues in the larger towns of Leaderboard 2.0, primarily due to memory constraints when computing the cache for the roads and lanes. Secondly, the existing classes in the BEV representation fall short in capturing complex new scenarios introduced in the updated leaderboard, such as construction zones or the presence of emergency vehicles. Lastly, there is uncertainty regarding the effectiveness of Roach’s model-free RL approach when faced with the heightened complexity and dynamic requirements of Leaderboard 2.0. It is important to note that Roach was trained for about one week on an Nvidia RTX 2080 Ti to achieve its results on Leaderboard 1.0. Considering the increased difficulty and complexity of scenarios in Leaderboard 2.0, adapting and retraining Roach could potentially require significantly more time, further complicating its deployment in this new benchmark.

Recognizing the limitations of existing expert agents for Leaderboard 2.0, CARLA has made available a set of driving logs that showcase human-driven routes in various scenarios. However, these logs alone do not suffice to train a system capable of processing sensor inputs and generating control commands. In response, and inspired by the approaches of LBC and Roach, we propose PRIBOOT, a method that simplifies the perception component by employing a Bird’s

Eye View (BEV) as the primary input. However, instead of using BEV as independent mask channels for training a CNN from scratch, PRIBOOT converts the mask into an RGB image and leverages transfer learning techniques using pre-trained networks from the ImageNet dataset [20]. This adaptation is crucial, particularly given the limited data available.

III. METHOD

PRIBOOT (**P**rileged **I**nformation **B**ootstrapping) leverages the limited logs available in Leaderboard 2.0 to establish the first expert agent capable of achieving satisfactory results within this demanding benchmark, as we show in Section IV. The development of PRIBOOT was structured in two phases: First, we focused on generating the most effective input representation tailored to the unique challenges of Leaderboard 2.0, as detailed in Section III-A. Following this, we designed and implemented a neural network architecture that is specifically optimized for handling the constraints of limited data, described in Section III-B.

A. Generation of Bird’s Eye View

Building on the approach used by Roach [10] and LBC [16], we employ a BEV to model the environment. However, adaptations were necessary to tailor it for the complexities of Leaderboard 2.0. Roach’s and LBC’s BEV include various classes such as roads, desired routes, lane boundaries, vehicles, pedestrians, traffic lights, and stop signs. While these classes were adequate for Leaderboard 1.0, they proved insufficient for the expanded scope of Leaderboard 2.0. Our enhancements to the BEV are outlined below:

a) Scalable Cache: Current BEV approaches utilize a caching mechanism to store the roads and lanes for the entire town, a process completed once per town to facilitate real-time BEV generation. However, applying this method to the larger towns in Leaderboard 2.0 caused memory overflows due to the use of Pygame. We addressed this by adopting a more efficient caching technique inspired

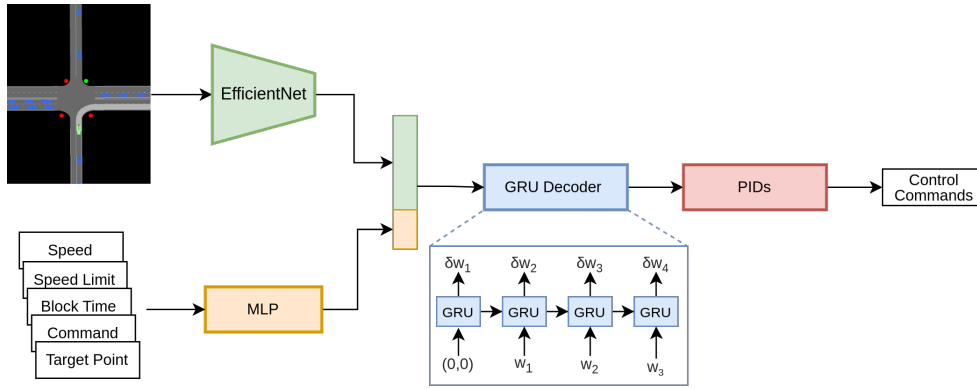


Fig. 2. Architecture of PRIBOOT. The system receives two types of inputs: a BEV image and a vector of vehicle measurements. These inputs are processed independently—the BEV through a pretrained EfficientNet model and the vehicle measurements via an MLP. The resultant feature vectors from both models are concatenated to form a comprehensive feature vector, which is then fed into a GRU-based waypoint decoder, similar to the approach used by Transfuser [21]. The final stage involves processing the waypoints through both longitudinal and lateral PID controllers to generate the vehicle control commands.

by deepsense.ai³, implementing the cache with NumPy for enhanced performance.

b) Decomposition of the Vehicles Class: In current BEV representations, all vehicle types are aggregated under a single class. We refined this by segmenting the Vehicles class into three distinct categories: Bikes, Emergency Vehicles, and Regular Vehicles. This differentiation is crucial as the driving behavior varies significantly based on the type of nearby vehicle, especially in emergency situations (see Figure 1a).

c) Simplified Motion Representation: Roach’s BEV uses multiple temporal masks to capture movement, which increases significantly the computational load. Instead, we introduced a single additional mask featuring an arrow for each actor, as illustrated with green arrows in Figure 1. This arrow indicates both the direction (orientation) and the speed (length) of the actor, simplifying the representation while reducing memory and computational demands.

d) Incorporation of a New Class: Leaderboard 2.0 introduces scenarios requiring interaction with new environmental elements not covered by existing classes. For instance, construction zones that necessitate slight route deviations were not previously accounted for. To accommodate this, we added a new class named Construction, which encompasses all pertinent elements like traffic cones and street barriers, represented in orange in 1b.

e) RGB Format Instead of Masks: Roach and LBC process the BEV using independent mask channels, requiring the training of a CNN from scratch. Given the limited data in Leaderboard 2.0, we found that converting these masks into an RGB format and utilizing pre-trained visual encoders not only saves training time but also enhances the model’s performance and efficiency.

B. Architecture

The architecture of PRIBOOT is depicted in Figure 2. Our system takes as input a BEV image and a vector of vehicle measurements. The vehicle measurement vector

encompasses several key parameters: current speed and the road speed limit, block time, a target point, and a navigation command. The “block time” parameter denotes the duration during which the vehicle has been stationary, aiding the system in determining whether to overtake or maintain its position due to typical traffic conditions. The “target point” is a waypoint located 30 meters ahead on the desired trajectory provided by a global planner, and the “navigation command” provides high-level directional indication from the global planner, encoded as a one-hot vector.

We utilize an EfficientNet [22] for processing the BEV image and an MLP for handling vehicle measurements. Given the constraint of limited available data, we adopt transfer learning by employing the pretraining of EfficientNet with the ImageNet dataset. Subsequently, the feature vectors extracted from both the EfficientNet and the MLP are merged and inputted into an autoregressive GRU decoder. This decoder is tasked with predicting the subsequent $T=4$ waypoints $\{w_t\}_{t=1}^T$ within the ego-vehicle coordinate framework, drawing inspiration from the methodology applied in Transfuser [21].

To convert the predicted waypoints into control commands, we employ two PID controllers—one for lateral control and another for longitudinal control—following methodologies from [16], [21]. The longitudinal controller uses the magnitude of the average vector between consecutive waypoints, while the lateral controller relies on their orientation. Additionally, similar to Transfuser, we integrate a creeping behavior and a safety heuristic mechanism utilizing information from the simulator.

This system is trained end-to-end using an L_1 loss between the predicted waypoints and the ground truth waypoints from the human logs. Let w_t^* represent the ground truth waypoint at timestep t , the the loss function is defined as:

$$\mathcal{L} = \sum_{t=1}^T \|w_t - w_t^*\|_1. \quad (1)$$

The human demonstrations typically exhibit minimal de-

³<https://github.com/deepsense-ai/carla-birdeye-view>




Fig. 3. Data augmentation techniques used to expose the agent to a broader range of driving scenarios. In these illustrations, white dots indicate the future waypoints that were followed by the human driver, and the black dot represents the target point. (a) Displays a sample with no augmentation, showing the standard scenario. (b) Shows a sample where both translation and rotation augmentations have been applied to the ego vehicle, illustrating a situation where the agent needs to recover to the center of the lane.

violation from the center of the lane, resulting in noise-free data. However, this adherence to the centerline causes a distribution shift between the training distribution and inference distribution. During inference, due to planning or controller inaccuracies, the agent may find itself in scenarios that deviate from the center of the lane. These instances are encountered as out-of-distribution events, presenting difficulties for the agent to navigate. To address this issue, inspired by the LBC approach, we use data augmentation techniques regarding the position and orientation of the vehicle. By varying the position and orientation of the vehicle, we expose the agent to diverse configurations, enabling it to learn effective recovery strategies. Figure 3 provides a visual representation of this augmentation process.

IV. EXPERIMENTS

This section starts with an overview of the setup used for collecting the experiments, followed by a comparative analysis of expert agents. It concludes with a presentation of an ablation study.

A. Setup

1) *Benchmark*: CARLA Leaderboard 2.0 builds upon CARLA Leaderboard 1.0, increasing the complexity of the benchmark in three distinct ways: a) by extending the route lengths approximately tenfold, b) by incorporating a new set of intricate driving scenarios derived from the NHTSA typology [23], and c) by increasing the frequency of these scenarios along each route. Additionally, this new benchmark introduces larger and more complex environments, as exemplified by Town12 and Town13. Town 12 is a 10×10 km² map that features a mix of urban, residential, and rural areas, offering varied types of challenges. Town13, while sharing many characteristics with Town12, distinguishes itself with different architectural styles, road and pavement textures,

TABLE I

COMPARISON OF RUN TIME INFERENCE USING THE EXPERT AGENTS.

	Run Time ↓ s
Autopilot	0.007
PRIBOOT	0.011

TABLE II

ABBREVIATION AND THE CORRESPONDING FULL NAME OF THE METRICS USED IN LEADERBOARD 2.0.

Abbreviation	Full Name
DS	Driving Score
IRS	Infraction Rate Score
RC	Route Completion
IP	Infraction Penalty
C.P	Collisions Pedestrians
C.V	Collisions Vehicles
C.L	Collisions Layout
R.L	Red Light Infractions
Stop	Stop Sign Infractions
O.R	Off-road Infractions
R.D	Route Deviation
Block	Agent Blocked
Y.E	Yield Emergency Infractions
S.T	Scenario Timeouts
M.S	Min Speed Infractions

and vegetation types. These enhancements aim to rigorously test the adaptability and resilience of autonomous driving systems under varied and challenging conditions.

The benchmark uses different metrics to assess different aspects of driving performance. The Route Completion (**RC**) indicates the percentage of the route completed by the agent. The Infraction Penalty (**IP**) quantifies the severity of infractions and is calculated using the following formula:

$$\mathbf{IP} = \prod_{i=1}^q p_i^{n_i}, \quad (2)$$

where q denotes the total number of different infraction types, p_i is the penalty associated with the infraction type i , and n_i is the number of infractions of type i . The main metric of the benchmark, Driving Score (**DS**), is calculated by multiplying **RC** and **IP**, providing a composite score that reflects both route completion success and adherence to driving regulations.

2) *Infraction Rate Score*: While **DS** provides valuable insights into agent performance, it inherently biases against longer routes due to its cumulative penalty for infractions, which are statistically more likely to occur over extended distances. To address this discrepancy and promote fairness, we introduce the Infraction Rate Score (**IRS**). This metric accounts for the infraction rate per kilometer, adjusting for route length and providing a balanced evaluation across varying driving conditions. The **IRS** is defined as:

$$\mathbf{IRS} = \mathbf{RC} \cdot \prod_{i=1}^q e^{-\lambda \cdot \frac{n_i}{L} \cdot (1-p_i)}, \quad (3)$$

TABLE III

DRIVING PERFORMANCE AND INFRACTION ANALYSIS OF EXPERT AGENTS ON CARLA LEADERBOARD 2.0 IN TOWN12 AND TOWN13.

		DS \uparrow	IRS \uparrow	RC \uparrow	IP \uparrow	C.P \downarrow	C.V \downarrow	C.L \downarrow	R.L \downarrow	Stop \downarrow	O.R \downarrow	R.D \downarrow	Block \downarrow	Y.E \downarrow	S.T \downarrow	M.S \downarrow
		%	%	%	%	#/Km	#/Km	#/Km	#/Km	#/Km	#/Km	#/Km	#/Km	#/Km	#/Km	#/Km
Town12	Autopilot	1.22	0.51	5.97	0.26	1.26	4.59	0.58	0.11	1.84	0.62	0.66	1.26	0.00	0.34	0.00
	PRIBOOT	22.80	42.75	76.46	0.30	0.00	0.31	0.06	0.01	0.02	0.05	0.00	0.06	0.04	0.03	0.11
Town13	Autopilot	0.99	0.22	5.55	0.20	0.83	3.06	0.83	0.00	0.02	0.35	0.69	0.69	0.00	0.10	0.00
	PRIBOOT	18.84	46.97	74.29	0.24	0.01	0.34	0.05	0.00	0.01	0.05	0.00	0.04	0.02	0.02	0.06

where L represents the length of the route in kilometers, and λ is a tunable exponent set to 4 based on empirical testing to optimize the metric’s sensitivity to infractions per distance traveled.

3) *Training Details*: We utilized the human driving logs provided by CARLA to train PRIBOOT. These logs correspond to 10 routes in Town12 and 10 routes in Town13 and amount to approximately 700,000 samples collected at a frequency of 20Hz. Each sample contains all the information required at each training step, including the BEV image, the vector of vehicle measurements, and the global location of the agent on the map. For our experiments, we used CARLA version 0.9.15. PRIBOOT was trained on a single NVIDIA A40 GPU. During the training phase, we used a batch size of 256 and the Adam optimizer [24] with a learning rate of 0.0001.

B. Comparative Analysis

This section outlines a comparative analysis conducted on Leaderboard 2.0, focusing exclusively on two expert agents: Autopilot and PRIBOOT. An attempt was made to adapt the Roach system to this benchmark; however, it was unsuccessful. The benchmark currently cannot support running a RL algorithm like Roach due to memory leaks that prevent the execution of millions of steps without causing server crashes.

Table I provides a comparison of the runtime between Autopilot and PRIBOOT. In this evaluation, Autopilot achieves a runtime of 0.007 seconds, while PRIBOOT records a runtime of 0.011 seconds. This difference in performance is expected, given that Autopilot operates based on a predefined set of rules, whereas PRIBOOT processes high-dimensional inputs.

For the Leaderboard 2.0 results, 15 metrics were utilized to assess the performance of the models. These metrics are detailed in Table II, where each abbreviation is associated with its full metric name.

As demonstrated in Table III, we conducted performance comparisons of the agents in two distinct Towns: Town12 and Town13. The evaluations are based on averages derived from 90 routes in Town12 and 20 routes in Town13, as stipulated in Leaderboard 2.0. PRIBOOT consistently outperformed Autopilot across nearly all metrics in both towns, often by substantial margins. In Town12, for instance, PRIBOOT’s **DS** was approximately 19 times higher than that of Autopilot, and its **IRS** was 84 times better. Similar trends were observed in Town13, with PRIBOOT achieving 19 times

higher **DS** and 214 times higher **IRS** than Autopilot. Notably, PRIBOOT recorded zero collisions with pedestrians per kilometer in Town12 and only 0.01 collisions per kilometer in Town13, underscoring its effectiveness in minimizing accidents involving pedestrians.

In contrast, Autopilot demonstrated superior performance in two specific metrics: yielding to emergency vehicles and maintaining minimum speed. The former was due to its lower **RC** score, which resulted in zero scenarios requiring yielding to an emergency vehicle. The latter stems from Autopilot operating at a fixed target speed consistently above the minimum speed requirement for the roads where the agent drove.

PRIBOOT stands out as the first agent to achieve a **RC** of approximately 75% in both towns, coupled with a satisfactory **DS** and **IRS**. This marks a significant milestone, positioning PRIBOOT as a pioneering agent capable of navigating the complexities of the benchmark, which can be used for data collection or online demonstrations.

To enhance the quantitative comparison presented earlier, we also include a qualitative evaluation. Our analysis of all routes in the benchmark reveals that Autopilot struggles with the novel scenarios introduced by Leaderboard 2.0, particularly those requiring slight deviations from the global planner’s trajectory. These scenarios include instances like parking exits and lane obstacles. Figure 4 illustrates a sequence of keyframes in a parking exit scenario, first showing Autopilot’s performance and then PRIBOOT’s. As shown, Autopilot immediately exits the park without considering the vehicles in the lane, leading to a collision. Conversely, PRIBOOT waits for a moment when the lane is clear before exiting, as expected.

Additionally, we provide access to a series of demonstration videos that illustrate the performance of PRIBOOT on Leaderboard 2.0. These can be accessed here.

C. Ablation Study

To explore the individual contributions of key components within PRIBOOT, particularly under conditions of limited data, we performed an ablation study focusing on two crucial elements: data augmentation and the utilization of RGB BEV in conjunction with transfer learning. This study involved training two variants of PRIBOOT: the first variant (referred to as "w/o aug") was developed without the data augmentations depicted in Figure 3b, and the second variant (referred to as "w/ masks") employed the BEV as a set

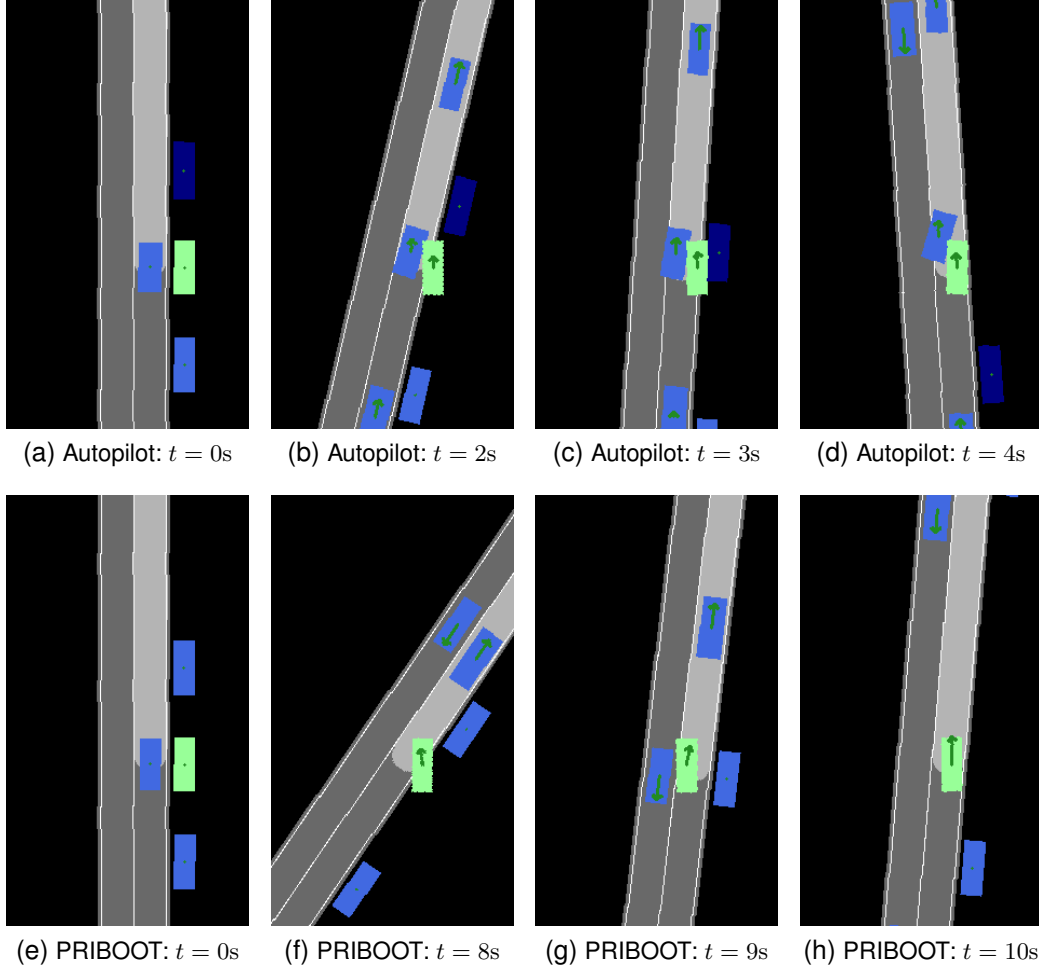


Fig. 4. Qualitative comparison in a parking exit scenario between Autopilot and PRIBOOT. The first row depicts a sequence of keyframes from Autopilot, while the second row shows the keyframes from PRIBOOT.

TABLE IV

ABLATION STUDY: DRIVING PERFORMANCE AND INFRACTION ANALYSIS OF PRIBOOT VARIANTS ON CARLA LEADERBOARD 2.0 IN TOWN13.

		DS \uparrow	IRS \uparrow	RC \uparrow	IP \uparrow	C.P \downarrow	C.V \downarrow	C.L \downarrow	R.L \downarrow	Stop \downarrow	O.R \downarrow	R.D \downarrow	Block \downarrow	Y.E \downarrow	S.T \downarrow	M.S \downarrow
		%	%	%	%	#/Km	#/Km	#/Km	#/Km	#/Km	#/Km	#/Km	#/Km	#/Km	#/Km	#/Km
Town13	w/o aug	2.66	5.22	15.92	0.29	0.01	1.22	0.35	0.00	0.02	0.45	0.02	0.40	0.00	0.17	0.02
	w/ masks	5.55	21.08	54.02	0.18	0.05	0.52	0.23	0.01	0.00	0.08	0.03	0.06	0.05	0.11	0.14
	PRIBOOT	18.84	46.97	74.29	0.24	0.01	0.34	0.05	0.00	0.01	0.05	0.00	0.04	0.02	0.02	0.06

of masks and training a CNN from scratch, consistent with methodologies reported in the literature [10].

The comparative analysis of driving performance and infractions for these variants is presented in Table IV. The w/o aug variant exhibited a significant decline in performance, as evidenced by a RC of approximately 16%, which adversely affected all other performance metrics. The reason for this is that due to planning or control inaccuracies, the agent encounters situations where it deviates from the center of the lane and lacks the capability to effectively recover. On the other hand, the w/ masks variant demonstrated improved results compared to the w/o aug variant. However, it still fell short of the full PRIBOOT system’s capabilities. Specifically, the w/ masks variant scored three times lower in terms of

DS and two times lower in terms of IRS.

Figure 5 illustrates the validation loss across epochs for the ablations considered. While the w/o aug variant achieves results similar to PRIBOOT during training, it recurrently faces out-of-distribution events, as detailed in Table IV. In contrast, the w/ masks variant displays a distinct pattern in validation loss: it requires 20% more epochs to converge and converges at a loss value that is twice that of PRIBOOT. This performance deficit underscores the critical role of utilizing RGB BEV and transfer learning techniques in cases where data availability is limited.

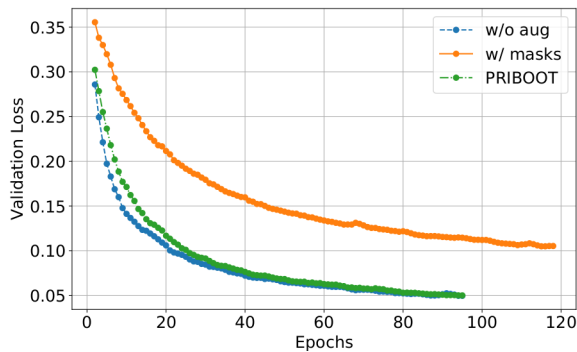


Fig. 5. Validation loss across epochs of PRIBOOT variants.

V. CONCLUSION

In this paper, we introduced PRIBOOT, a system that utilizes privileged information alongside limited human driving logs to establish the first expert with satisfactory driving performance on the CARLA Leaderboard 2.0. Our results demonstrate that PRIBOOT significantly outperforms Autopilot across nearly all benchmark metrics, highlighting its superior capability in complex and challenging autonomous driving scenarios. Additionally, we presented an ablation study that evaluates the impact of using augmentations to aid recovery processes. Furthermore, we demonstrated the benefits of employing RGB BEV images with transfer learning, which proved more efficient in terms of training speed and performance than using masks and training a CNN from scratch. While our work has focused on the CARLA simulator, it is important to note that the idea behind PRIBOOT can eventually be applied in other simulators. In the future, we plan to employ PRIBOOT to generate large datasets that can be instrumental in training student models that receive sensor information as input.

ACKNOWLEDGMENTS

This work has been supported by FCT - Foundation for Science and Technology, in the context of Ph.D. scholarship 2022.10977.BD and by National Funds through the FCT - Foundation for Science and Technology, in the context of the project UIDB/00127/2020.

REFERENCES

- [1] P. Andersson and P. Ivehammar, "Benefits and costs of autonomous trucks and cars," *Journal of Transportation Technologies*, vol. 9, no. 2, 2019.
- [2] D. Coelho and M. Oliveira, "A review of end-to-end autonomous driving in urban environments," *IEEE Access*, vol. 10, pp. 75 296–75 311, 2022.
- [3] 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.
- [4] Y. Li, W. Yuan, S. Zhang, W. Yan, Q. Shen, C. Wang, and M. Yang, "Choose your simulator wisely: A review on open-source simulators for autonomous driving," *IEEE Transactions on Intelligent Vehicles*, 2024.
- [5] P. Kaur, S. Taghavi, Z. Tian, and W. Shi, "A survey on simulators for testing self-driving cars," in *2021 Fourth International Conference on Connected and Autonomous Driving (MetroCAD)*. IEEE, 2021, pp. 62–70.
- [6] H. Shao, L. Wang, R. Chen, S. L. Waslander, H. Li, and Y. Liu, "Reasonnet: End-to-end driving with temporal and global reasoning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 13 723–13 733.
- [7] H. Shao, L. Wang, R. Chen, H. Li, and Y. Liu, "Safety-enhanced autonomous driving using interpretable sensor fusion transformer," in *Conference on Robot Learning*. PMLR, 2023, pp. 726–737.
- [8] P. Wu, X. Jia, L. Chen, J. Yan, H. Li, and Y. Qiao, "Trajectory-guided control prediction for end-to-end autonomous driving: A simple yet strong baseline," *Advances in Neural Information Processing Systems*, vol. 35, pp. 6119–6132, 2022.
- [9] W. Zhang, M. Elmahgiubi, K. Rezaee, B. Khamidehi, H. Mirkhani, F. Arasteh, C. Li, M. A. Kaleem, E. R. Corral-Soto, D. Sharma *et al.*, "Analysis of a modular autonomous driving architecture: The top submission to carla leaderboard 2.0 challenge," *arXiv preprint arXiv:2405.01394*, 2024.
- [10] 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.
- [11] Y. Pan, C.-A. Cheng, K. Saigol, K. Lee, X. Yan, E. Theodorou, and B. Boots, "Agile autonomous driving using end-to-end deep imitation learning," *arXiv preprint arXiv:1709.07174*, 2017.
- [12] X. Jia, Y. Gao, L. Chen, J. Yan, P. L. Liu, and H. Li, "Driveadapter: Breaking the coupling barrier of perception and planning in end-to-end autonomous driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 7953–7963.
- [13] F. Codevilla, E. Santana, A. M. López, and A. Gaidon, "Exploring the limitations of behavior cloning for autonomous driving," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 9329–9338.
- [14] R. Chekroun, M. Toromanoff, S. Hornauer, and F. Moutarde, "Gri: General reinforced imitation and its application to vision-based autonomous driving," *Robotics*, vol. 12, no. 5, p. 127, 2023.
- [15] 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.
- [16] D. Chen, B. Zhou, V. Koltun, and P. Krähennühl, "Learning by cheating," in *Conference on Robot Learning*. PMLR, 2020, pp. 66–75.
- [17] D. Coelho, M. Oliveira, and V. Santos, "Rlfold: Reinforcement learning from online demonstrations in urban autonomous driving," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 10, 2024, pp. 11 660–11 668.
- [18] A. Zhao, T. He, Y. Liang, H. Huang, G. Van den Broeck, and S. Soatto, "Sam: Squeeze-and-mimic networks for conditional visual driving policy learning," in *Conference on Robot Learning*. PMLR, 2021, pp. 156–175.
- [19] Y. Xiao, F. Codevilla, D. Porres, and A. M. López, "Scaling vision-based end-to-end autonomous driving with multi-view attention learning," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 1586–1593.
- [20] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 248–255.
- [21] K. Chitta, A. Prakash, B. Jaeger, Z. Yu, K. Renz, and A. Geiger, "Transfuser: Imitation with transformer-based sensor fusion for autonomous driving," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [22] M. Tan and Q. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *International conference on machine learning*. PMLR, 2019, pp. 6105–6114.
- [23] W. G. Najm, J. D. Smith, and M. Yanagisawa, "Pre-crash scenario typology for crash avoidance research," John A. Volpe National Transportation Systems Center (U.S.), Technical Report DOT-VNTSC-NHTSA-06-02; DOT HS 810 767, 4 2007. [Online]. Available: <https://rosap.ntl.bts.gov/view/dot/6281>
- [24] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.