Event-Aware Relabeling for Addressing Future Leakage in End-to-End Autonomous Driving

1st Haechul Chang
The Robotics Program
Korea Advanced Institute of Science and Technology (KAIST)
Daejeon, Rep. of Korea, 34051
haechulj@kaist.ac.kr

2nd Siewoo Kim
The CCS Graduate School of Mobility
Korea Advanced Institute of Science and Technology (KAIST)
Daejeon, Rep. of Korea, 34051
siewoo.kim@kaist.ac.kr

3rd Adeeb M. Islam
The Robotics Program
Korea Advanced Institute of Science and Technology (KAIST)
Daejeon, Rep. of Korea, 34051
adeeb.islam8@kaist.ac.kr

4th Seong-Jun Kim
The CCS Graduate School of Mobility
Korea Advanced Institute of Science and Technology (KAIST)
Daejeon, Rep. of Korea, 34051
sjun328@kaist.ac.kr

5th Seung-Hyun Kong*
The CCS Graduate School of Mobility
Korea Advanced Institute of Science and Technology (KAIST)
Daejeon, Rep. of Korea, 34051
skong@kaist.ac.kr

***Abstract*—Recent advances in autonomous driving have shifted the field from modular pipelines to End-to-End (E2E) frameworks. E2E frameworks directly map raw sensor inputs to control commands, leveraging large-scale data to learn robust driving policies. However, the conventional data generation process relies on future actions and trajectory information as Ground Truth (GT). This procedure introduces a "Future Leakage" problem, where information unavailable at inference time unintentionally enters the training process. This leads to causal confusion, where the model learns spurious correlations instead of true causality, thus compromising driving safety. In this paper, we propose an Event-Aware Relabeling technique to address this problem. The proposed method detects interaction events with obstacles and vehicles by utilizing traffic lights and the ego-vehicle's state information. Then it effectively eliminates the leakage of future information by performing appropriate relabeling on the data prior to the event occurrence. Experimental results demonstrate that our method effectively mitigates the causal confusion problem and achieves a driving score improvement of approximately 14% compared to baseline. This study highlights that a high-quality data curation process is essential to ensure the safety of E2E autonomous driving systems and validates the potential of data-centric performance optimization.**

***Index Terms*—End-to-end autonomous driving, dataset refinement, causal confusion**

I. Introduction

Autonomous driving research is transitioning from traditional modular pipelines to End-to-End (E2E) paradigms. Traditional approaches are often limited by error propagation and a heavy reliance on manually engineered components. In contrast, E2E models directly map raw sensor inputs to control commands. Leveraging large-scale data-driven learning, these systems have demonstrated robust generalization capabilities [1]–[3].

*Corresponding author

Most E2E models rely on Imitation Learning (IL), where expert future trajectories are used as Ground Truth (GT). However, directly adopting expert demonstrations induces the future leakage problem. This occurs because the GT implicitly embeds information unavailable at the current time step, such as future traffic light changes. This leakage prevents the model from learning true causal relationships and leads to causal confusion [4]. As a result, the model may capture spurious correlations rather than understanding actual driving contexts. This issue often remains hidden in open-loop evaluations but results in critical safety failures in real-world closed-loop environments.

These limitations indicate that ensuring the safety of E2E autonomous driving requires efforts beyond architectural improvements. A more fundamental solution lies in enhancing data quality and restoring proper causal structure within training labels. To address this challenge, we propose an Event-Aware Relabeling framework. This method identifies critical driving events, such as traffic stops or object interactions, and rectifies the preceding data segments where future information is unintentionally embedded. This process enables the model to learn decisions grounded solely in information observable at each moment, thereby establishing a causally consistent training signal.

The main contributions of this study are summarized as follows:

- First, we provide a concrete analysis of the causal confusion problem caused by future leakage in E2E autonomous driving datasets and identify it as a cause of driving performance degradation.
- Second, we propose an Event-Aware Relabeling algorithm that utilizes traffic light and ego-vehicle status information to eliminate future information and preserve

causality at the current time step.

- Third, we demonstrate that our method improves driving scores by approximately 14% compared to baselines. This significant gain underscores the critical role of data curation in enhancing the safety and performance of autonomous driving systems.

The remainder of this paper is organized as follows. Section II reviews the latest research trends in E2E autonomous driving and related works on dataset curation. Section III details the specific methodology of the Event-Aware Relabeling algorithm proposed in this study. Section IV presents a quantitative and qualitative analysis of the proposed method's impact on driving safety and performance. Finally, Section V concludes the paper by reaffirming the importance of addressing the future leakage problem based on the experimental results.

II. Related Works

A. End-to-End Autonomous Driving

E2E autonomous driving directly maps raw sensor inputs such as camera or LiDAR data to control signals including steering, throttle, and brake or to future trajectories through a unified neural network. Beginning with ALVINN [1] in 1989, NVIDIA's PilotNet [2] further demonstrated the feasibility of the E2E paradigm by successfully performing lane-keeping with a CNN-based architecture. More recent works such as Uni-AD [3], VAD [5], and Gen-AD [6] focus on improving situational understanding and strengthening interaction modeling with surrounding agents by explicitly integrating perception and planning within a single framework. In addition, with the rapid advancement of Large Language Models (LLMs), E2E autonomous driving research has begun to explore LLM-inspired architectures. For example, SimLingo [7] incorporates LLMs to combine semantic reasoning with driving control. However, imitation learning approaches are inherently limited by two critical challenges: covariate shift and causal confusion. Covariate shift arises from the distribution mismatch between expert demonstrations and real-world environments. Meanwhile, causal confusion occurs when models learn incorrect causal associations due to spurious correlations within the dataset.

B. Causal Confusion in Imitation Learning

Causal confusion occurs when a model learns spurious correlations instead of identifying the true causal factors behind expert actions [4]. De Haan et al. [4] point out that high-dimensional sensory data such as images can unintentionally intensify this issue by distracting the model from causally meaningful cues. In autonomous driving, this often appears as the inertia problem, in which models tend to imitate previous actions or rely heavily on speed-related cues instead of responding appropriately to visual information such as traffic lights or obstacles [8]. The future leakage phenomenon, which is a central focus of this study, further aggravates causal confusion. Specifically, future events such as traffic light transitions or obstacle interactions are often inadvertently embedded in GT labels due to temporal misalignment. Consequently, the model attempts to infer nonexistent causal cues from the current input. This reduces generalization performance and undermines the model's ability to respond reliably to unexpected situations in real-world environments.

C. Data Curation and Relabeling

Achieving robust E2E autonomous driving requires not only improvements in architectural design but also the construction of high-quality datasets. Early work such as DAgger [9] attempted to mitigate distribution mismatch by allowing expert intervention during online training. However, this approach is costly and difficult to scale. As a result, more recent efforts have shifted toward offline data curation and relabeling strategies to enhance learning efficiency. Some studies [10], [11] focus on removing noise from expert demonstrations or reducing ambiguity in supervisory signals by aggregating behaviors from multiple experts. Nonetheless, these approaches mainly address human error or sensor noise and do not directly resolve the temporal inconsistencies that cause future leakage. In contrast, the Event-Aware Relabeling method proposed in this paper detects future events in advance and selectively relabels segments that contain causally invalid information at the current time. This process eliminates future-dependent bias in the GT labels and enables the model to learn causally sound driving policies, thereby overcoming the limitations of existing data curation techniques.

III. Methodology

A. Overview

This section describes the overall framework designed to address the future leakage problem that commonly arises during the training of E2E autonomous driving models. The proposed framework is composed of three core components: (1) Model Architecture that processes multimodal inputs to perceive the environment and generate control commands; (2) Event Detection module that identifies driving scenarios within expert demonstrations prone to future leakage; and (3) Event-Aware Relabeling module that rectifies GT labels in the detected segments to ensure causal consistency. We adopt the SimLingo-base [7] model as a baseline and modify it with structures tailored to autonomous driving. In addition, we enhance training stability through a data curation process that eliminates causally invalid supervision.

B. Model Architecture

The model proposed in this study is based on SimLingo-base and is designed to jointly process visual information, navigation cues, and ego-vehicle state signals. The overall architecture is illustrated in Figure 2. The visual inputs consist of images captured from five vehicle-mounted cameras that provide front, left, right, rear-left, and rear-right views. To incorporate temporal context, the model uses both the current frame at time t and a previous frame at time $t-\delta$, resulting in a total of ten input images. A route map image that provides global navigation and planning information is also included.

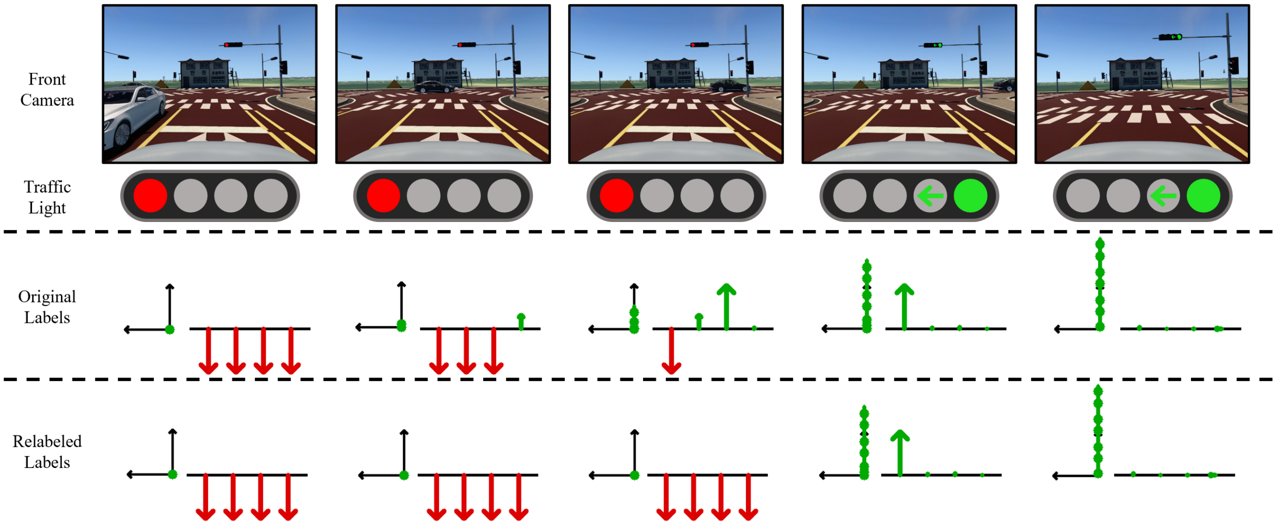

Fig. 1. **Event-Aware Relabeling applied to trajectory and action labels during stop scenes.** This figure compares the original trajectory and action labels generated before applying relabeling with those modified to maintain a stationary state over the same interval. In each label visualization, the left plots show the trajectory labels derived from future positions, while the right plots present the future action labels extracted at 0.5-second intervals over a 2-second horizon. In the action labels, red arrows pointing downward indicate deceleration, whereas green arrows pointing upward represent acceleration.

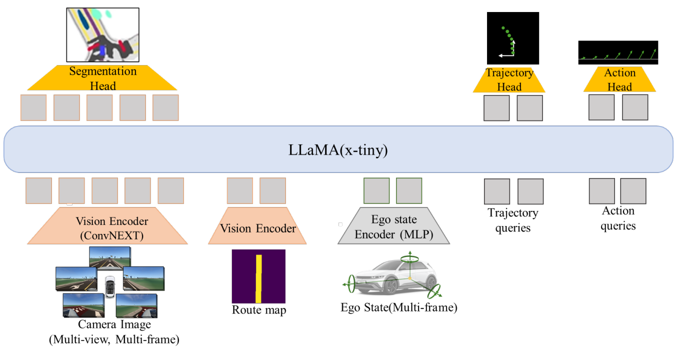

Fig. 2. **Overview of the proposed architecture.** The model encodes multi-modal inputs using a LLaMA (x-tiny) backbone to jointly predict BEV segmentation, trajectories, and control actions.

All visual inputs are passed through a Vision Encoder, which converts them into high-dimensional feature tokens that are then forwarded to the backbone network. For ego-state information, we use the most recent two seconds of vehicle motion data. This includes linear velocities along the three axes (v_x, v_y, v_z) as well as angular velocities $(\omega_x, \omega_y, \omega_z)$, which enable the model to precisely capture rotational dynamics. The encoded visual tokens and ego-state features are processed by a transformer backbone built upon the Llama (x-tiny) architecture. The output layer of the model is composed of two heads designed for perception and planning. The Segmentation Head receives tokens generated from camera images and predicts road layouts and surrounding objects in a Bird's Eye View (BEV) representation, which encourages the model to explicitly learn the geometric structure of the driving environment. The Planning Head takes ten trajectory queries and ten action queries as inputs. These queries are processed through a transformer decoder, after which the model predicts future trajectory points and control actions for the next ten time steps.

C. Event Detection

Future leakage primarily occurs in transitional situations in which the vehicle is either stopping or starting, and future information is implicitly pulled into the labeling process. Therefore, we define these situations as events and detect them using traffic light information and ego-vehicle behavior analysis. Examples of the detection process and results are shown in Figure 1. First, events are detected based on the status of the traffic light that the autonomous vehicle is currently facing. Rather than considering all traffic lights visible in the scene, we selectively identify valid traffic lights by examining the geometric relationship between the vehicle heading and the spatial position of each traffic light. Traffic lights that are too distant or excessively close to the vehicle are excluded, and only those located at a meaningful distance along the driving path are considered. An event occurrence is defined at the moment when a selected traffic light transitions from a stop signal (Red or Yellow) to a go signal (Green). Second, to complement cases where traffic light information is uncertain, we directly analyze changes in the motion state of the ego-vehicle. When the vehicle speed converges to zero and the vehicle is in a stationary state ($v \approx 0$), abrupt transitions between the current action and a future action are detected. Specifically, we identify the time step at which sustained braking transitions into acceleration and define this point as a potential departure event. This allows the detection of starting intentions even in the absence of reliable traffic light signals.

D. Event-aware Relabel

In segments where events are detected, future information may influence the GT labels recorded during data collection and cause causal confusion. To prevent this, we perform Event-Aware Relabeling. When an event such as a departure

TABLE I
COMPARISON OF DRIVING SCORES BY SCENARIO ACCORDING TO RELABELING STRATEGIES

Model	Scenario 1	Scenario 2	Scenario 3	Scenario 4	Scenario 5	Scenario 6	Scenario 7	Scenario 8	Scenario 9	Scenario 10	Avg.
Baseline	86.67	73.33	76.67	70.40	78.33	86.67	70.22	73.15	37.9	76.56	72.99
Relabel	**100.0**	**98.33**	**86.67**	**83.42**	**90.00**	**98.33**	**76.67**	**98.33**	**55.09**	**85.00**	**87.18**

after a traffic light change is detected, we trace back a predefined interval, referred to as the pre-event horizon, and refine the data within this segment. If the expert demonstrated a consistently maintained behavior before the abrupt change, such as remaining stationary, all labels within this interval are replaced with this consistent behavior. For the action labels, if the vehicle maintains a braking state above a certain threshold, all action labels prior to the event are modified to indicate full braking. This ensures that, even if the expert begins to apply acceleration at a future time step, the GT labels instruct the model to keep the vehicle fully stationary until the event actually occurs. The same principle applies to trajectory labels. Instead of using trajectory points that move toward a future position, the trajectory points are modified to remain at the current position of the stationary vehicle. Figure 1 illustrates the difference before and after relabeling. In the original data, future departure behavior was embedded before the traffic signal changed, whereas the proposed method enforces a full-stop state until the moment of the signal change and thereby restores causal consistency in the training labels.

IV. EXPERIMENTS

In this study, we validate the effectiveness of the proposed methodology using MORAI, a digital twin simulator that precisely models road environments in Korea. The training dataset was collected within MORAI and consists of approximately 2.6 million frames, corresponding to a total size of about 2.2 TB.

The evaluation was conducted across ten diverse driving scenarios constructed on K-City, a standardized autonomous driving testbed. These scenarios include tasks ranging from basic traffic rule compliance such as traffic signal obedience and lane keeping to more complex maneuvers that require high-level decision making, including unprotected left turns and stopping for unexpected pedestrian crossings, as summarized in Table II. The evaluation was performed in a real-time closed-loop setting, in which the model directly controls the vehicle and interacts with the environment, instead of relying on a conventional open-loop evaluation based on pre-recorded datasets. To ensure fairness, each scenario was executed three times and the final performance score was computed as the average over these trials.

Driving performance was measured according to the evaluation metrics defined in the CARLA Leaderboard [12]. The final Driving Score (DS) is calculated as the product of the Route Completion (RC), which measures the proportion of the planned route successfully driven by the agent, and a Penalty term that accounts for safety violations such as collisions, traffic signal violations, and lane boundary intrusions.

TABLE II
SCENARIO CONFIGURATIONS AND TEST ITEMS

Scenario	Key Test Items
Scenario 1	Lane Keeping, Car Following
Scenario 2	Unprotected Left Turn, Pedestrian Yielding
Scenario 3	Flashing Yellow Intersection Response, Right Turn
Scenario 4	Obstacle Avoidance, Traffic Light Compliance
Scenario 5	Traffic Light Compliance, Merging
Scenario 6	Traffic Light Compliance, Merging
Scenario 7	Pedestrian Yielding, Obstacle Avoidance
Scenario 8	Roundabout Driving
Scenario 9	Highway Driving, Highway Merging, Lane Change
Scenario 10	High-Density Traffic, Toll Gate Passage

A. Quantitative Results

To evaluate the effectiveness of the proposed Event-Aware Relabeling technique, we compared a baseline model trained without any data refinement (Baseline) with a model that applies the proposed method to both action and trajectory labels (Relabel). Table I presents the driving scores for each scenario as well as the overall average score.

The results show that the Baseline model consistently achieved lower driving scores across all scenarios, indicating that future leakage-induced causal confusion significantly degrades decision-making performance in real driving situations. In contrast, the Relabel model exhibited stable performance improvements in every scenario, achieving an average driving score approximately 14% higher than the Baseline.

These findings demonstrate that eliminating future leakage during the label generation process alone can substantially enhance the driving stability of E2E autonomous driving systems. This suggests that data curation plays a critical role in ensuring causally consistent learning signals, thereby enabling more reliable decision-making in real-world driving environments.

B. Qualitative Results

Figure 3 provides a visual comparison of model predictions with and without the proposed relabeling method. The baseline model, which suffers from future leakage, often exhibits unstable action outputs. For example, even when the traffic signal is still red, the model prematurely anticipates a future green signal and oscillates between actions or attempts to initiate acceleration when it should remain stationary. In contrast, the model trained with Event-Aware Relabeling produces significantly more stable and consistent predictions during event-critical situations involving traffic lights or surrounding vehicles. Notably, during red-signal waiting periods, the model reliably outputs a stationary action regardless of future signal changes. Furthermore, in interactions with leading vehicles, it generates smooth deceleration trajectories while maintaining

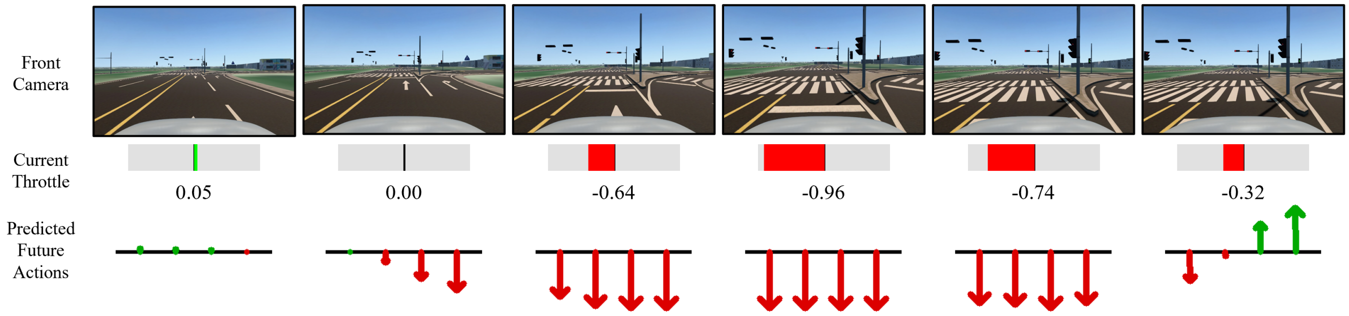

(a) Predictions from the model trained on the original labels

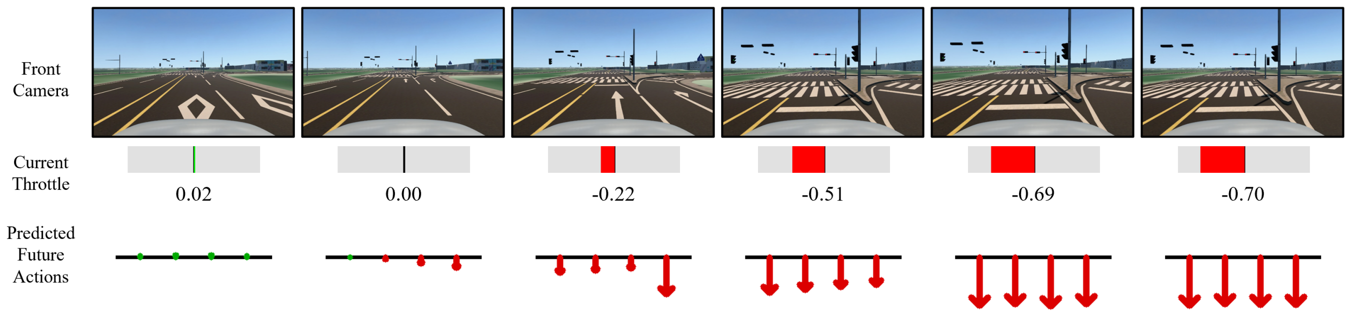

(b) Predictions from the model trained on the Event-Aware Relabeled labels

Fig. 3. Qualitative comparison of model predictions with and without Event-Aware Relabeling.

a safe following distance. These results confirm that the proposed method encourages the model to perform correct causal reasoning based solely on currently observable information rather than memorizing future states embedded in the dataset.

V. Conclusions

In this paper, we identified the future leakage problem that arises during E2E autonomous driving training and proposed an Event-Aware Relabeling technique to address it. By detecting events through analysis of traffic light states and ego-vehicle behavior, and by causally redefining the action and trajectory labels within the corresponding segments, the proposed method prevents the model from memorizing future information and guides it toward learning valid driving logic.

Experimental results show that the proposed technique improves driving scores by approximately 14% compared to the baseline, particularly enhancing stability in scenarios involving traffic light waiting and interactions with leading vehicles. Overall, this study demonstrates that high-quality data curation is essential for ensuring the safety of autonomous driving systems, extending beyond architectural improvements, and highlights the importance of data-centric approaches for achieving reliable full autonomy.

Acknowledgment

This work was supported by the Ministry of Trade, Industry and Resources (MOTIR) grant funded by the Korea government (No. N02250174).

References

[1] D. A. Pomerleau, "Alvinn: An autonomous land vehicle in a neural network," in *Advances in neural information processing systems*, 1989, pp. 305–313.

[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] Y. Hu, J. Yang, L. Chen, K. Li, C. Sima, X. Zhu, S. Chai, S. Du, T. Lin, W. Wang *et al.*, "Planning-oriented autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 17 853–17 862.

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

[5] B. Jiang, S. Chen, Q. Xu, B. Liao, J. Chen, H. Zhou, Q. Zhang, W. Liu, C. Huang, and X. Wang, "Vad: Vectorized scene representation for efficient autonomous driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 8340–8350.

[6] Z. Weng, Z. Wang, and C. McCarthy, "End-to-end autonomous driving with generative modeling," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.

[7] K. Renz, L. Chen, E. Arani, and O. Sinavski, "Simlingo: Vision-only closed-loop autonomous driving with language-action alignment," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 11 993–12 003.

[8] 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.

[9] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proceedings of the twenty-fourth international conference on artificial intelligence and statistics*, 2011, pp. 627–635.

[10] M. Bansal, A. Krizhevsky, and A. Ogale, "Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst," in *Robotics: Science and Systems (RSS)*, 2019.

[11] F. Codevilla, M. Müller, A. López, V. Koltun, and A. Dosovitskiy, "End-to-end driving via conditional imitation learning," *2018 IEEE*

International Conference on Robotics and Automation (ICRA), pp. 4693–4700, 2018.

[12] 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.