DriveSafer: End-to-End Autonomous Driving with Safety Guidance

Shounak Sural
Carnegie Mellon University
Pittsburgh, PA
iamshounak@gmail.com

Ragunathan (Raj) Rajkumar
Carnegie Mellon University
Pittsburgh, PA
rajkumar@andrew.cmu.edu

Abstract

End-to-End (E2E) autonomous driving models have shown growing capability in recent years, with performance improving on increasingly challenging benchmarks. However, modern generative E2E planners still suffer from a substantial number of catastrophic failures in safety-critical scenarios. We find that many such failures arise from violations of physical constraints and safety requirements, leading to unsafe behavior. Motivated by this finding, in this paper, we focus on improving safety outcomes in generative end-to-end driving with a targeted reduction of catastrophic planning failures, instead of enhancing average planning quality. Towards this end, we propose DriveSafer, a failure-aware safety framework for end-to-end planners. DriveSafer explicitly steers generative planners towards safe behaviors leveraging both training-time safety constraints and inference-time safety guidance. Compared to the state-of-the-art DiffusionDrive model from CVPR 2025, on the NAVSIM benchmark, DriveSafer reduces the number of catastrophic failures (`PDMS=0`) by 48%, with over 65% reduction in drivable-area compliance failures.

1. Introduction

End-to-End (E2E) autonomous driving models have recently become increasingly popular. Compared to traditional modular and rule-based architectures, these models show promise in addressing more complex scenarios with a large number of interacting agents, and in long-tailed scenarios aided by the fusion of vision and language with internet-scale datasets. While these E2E models handle many situations well, they can also be surprisingly brittle in many commonly seen situations which are straightforward for humans and even rule-based systems to handle. These misbehaviors typically show up in the long tail in end-to-end autonomous driving benchmarks such as NuPlan [2] and NAVSIM [8] and long-tail-focused datasets such as KIT-E2E [30] and WOD-E2E [37]. To identify many such failure cases, some benchmarks also employ CARLA [9] simulation allowing closed-loop testing with benchmarks such as Bench2Drive [14] and Longest6 [6]. Importantly, strong average benchmark performance for models, however, still masks a long tail of catastrophic failures.

While E2E models built using popular benchmarks have improved planning over the years, even the best-performing models such as DiffusionDrive [23] still suffer from catastrophic failures that simple rule-based planners can easily avoid. To reduce the failure modes of existing E2E models, a careful evaluation of safety is necessary. Many existing E2E models are benchmarked using the L2 distance error metric from human-driven ground truth trajectories as the primary metric on datasets such as NuScenes. While imitation learning minimizing the L2 error from ground truth offers useful insights, such metrics can be surprisingly easy to imitate *without* adding actual functional value. Methods such as AD-MLP [39] and Ego-MLP [22] have shown that a surprisingly simple MLP model which does *not* even consume sensor inputs can outperform state-of-the-art methods on the NuScenes benchmark which optimizes only for trajectory error. Such "blind" models that achieve state-of-the-art performance clearly underscore the need for better benchmarks and different evaluation metrics. NAVSIM [8] is one such benchmark that has standardized evaluation metrics which are aligned with driving safety. Their proposed Predictive Driver Model Score (PDMS) metric comprises collision checks, drivable area compliance, comfort, time-to-collision and driving direction compliance. These metrics are much better indicators for a safe planner. NAVSIMv2 [3] introduces the Extended PDMS metric that further incorporates traffic-light compliance, lane-keeping and extended comfort metrics for a comprehensive safety-conscious evaluation. Earlier proposed "blind" MLP models which can exhibit unsafe behaviors do not perform favorably under extended evaluation.

Many modern state-of-the-art methods report their performance on NAVSIM in terms of the PDMS score, but there is less emphasis on analyzing catastrophic failure cases in these models. In this paper, we specifically consider a PDMS score of zero in the NAVSIM E2E planning benchmark as a catastrophic failure, since its occurrences

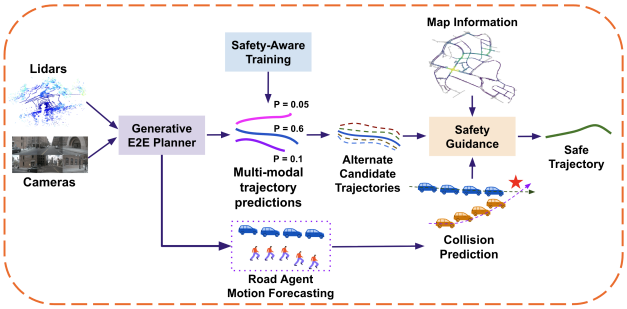

Figure 1. Overview of our *DriveSafer* framework

can potentially be hazardous and life-threatening in real-world driving. A PDMS score of 0 is assigned in the following cases: (i) the ego vehicle collides with another moving object such as a vehicle or a pedestrian, (ii) the ego vehicle drives off the permitted drivable region possibly onto a sidewalk or a solid divider, and (iii) the ego vehicle motion differs significantly from the intended driving direction, such as going straight when explicitly directed to turn right. Thus, we take a failure-centric view of safe driving behavior and focus on reducing these catastrophic failure cases as opposed to average benchmark performance. We carefully evaluate such failure cases and propose a framework called *DriveSafer* that strategically reduces these failures. Correspondingly, our contributions are as follows:

- We propose *DriveSafer*, a failure-centric safety enforcement framework for E2E planning that directly targets reducing catastrophic failures using synergistic training- and inference-time safety enforcement.
- We introduce a safety-aware training strategy with specialized physics-informed and safety-aware losses along with lightweight and integrated forecasting and planning.
- We develop an inference-time safety guidance method that rejects infeasible trajectories and steers generated trajectories towards safety-compliant ones.
- *DriveSafer* reduces catastrophic failures (`PDMS=0`) on the NAVSIM benchmark by 48% over the state-of-the-art DiffusionDrive method while also improving overall PDMS.

2. Related Work

End-to-End Autonomous Driving: Early E2E studies focused on generating driving commands such as steering angles and acceleration/brake commands directly from raw sensor inputs through imitation and reinforcement learning [4, 7]. However, these methods do not scale across different platforms due to varying vehicle dynamics. It is now well understood that low-level actuation commands are best generated through specialized controllers for each vehicle given a local motion plan. Hence, recent end-to-end autonomous driving literature focuses on generating local motion plans from sensor data. E2E methods use strategies such as learning scene representations in the BEV space [6, 11, 29, 34], learning vectorized architectures [5, 15], vision-language action model-based planning [10, 25, 33, 40, 44], world model-based planners [18, 31, 32, 42] and diffusion-based generative planners [23, 43].

Figure 2. A complex left-turn in NAVSIM [8]

Datasets for Safety in Autonomous Driving: Another direction of research focuses on enforcing safety in autonomous driving via datasets and simulation that critically evaluates safety, especially in rare and long-tailed scenarios. For example, NuPlan evaluates full-scale E2E behavior using closed-loop simulation. NAVSIMv2 [3] uses neural reconstruction to generate alternate scenarios complete with generated sensor data for pseudo simulation. WOD-E2E [37] uses millions of miles of Waymo driving data to mine long-tailed scenarios that occur rarely in daily life ($< 0.03\%$ of cases). Bench2Drive [14] models challenging interactions and harsh weather conditions in CARLA [9] for safety evaluation of models.

Methods for Safe Autonomous Driving: Another branch of work aims to make driving policies safer through methods such as CAT [41] which turns simple real-world scenarios into alternate plausible safety-critical ones to perform closed-loop adversarial training. DriveAdapter [13] injects hand-crafted safety rules into a teacher-student paradigm for training. Control Barrier Functions are sometimes used to enforce safety in planning [1, 16, 36]. LEAD [24] narrows the sim2real gap by minimizing unrealistic privileged information available to planning experts in the CARLA simulator [9].

Similar to LLM safety research, Reinforcement Learning-based alignment ideas have been applied to E2E autonomous driving with methods such as TrajHF [17], Carl [12] and DriveDPO [26]. Unlike prior work that focuses on such RL-based methods, our proposed *DriveSafer* framework focuses on guiding a generative planner towards safer choices by adhering to certain rules of the road. This guidance significantly reduces catastrophic failures and yields acceptable motion plans in a lot of cases.

3. Methodology

An overview of our *DriveSafer* framework is shown in Figure 1. We will next describe our framework.

3.1. Base End-to-end Planning Model

Our *DriveSafer* framework is based on the truncated diffusion model called DiffusionDrive [23]. We choose this

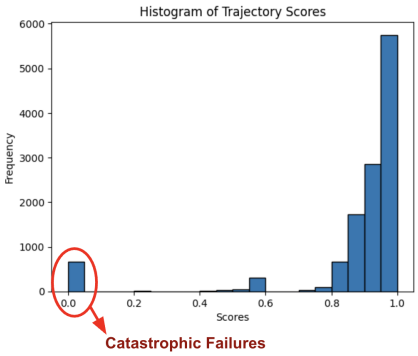

Figure 3. PDM Score Distribution for DiffusionDrive showing catastrophic failure cases with `PDMS=0`

primary baseline since it is a state-of-the-art method on the NAVSIM benchmark and performs inference at 45 fps, making it suitable for *DriveSafer* to be deployed on a real-world AV.

DiffusionDrive consumes lidar and camera sensor data and learns intermediate representations necessary for object detection, tracking and map segmentation. It then trains a 2-step diffusion model conditioned on the learnt intermediate representation. Instead of sampling from pure noise like most diffusion models, DiffusionDrive starts with noise centered around a few typical trajectory modes and performs trajectory planning with mode selection followed by denoising for fine adjustments in only 2 steps. While performing well in many situations, its effectiveness for real-world use is limited due to a large number of catastrophic failures as highlighted in Figure 3. We address such drawbacks with *DriveSafer*.

3.2. Safety-Aware Training in *DriveSafer*

Safety-aware training in *DriveSafer* is intended to make the generated trajectories acceptable for driving safety. We incorporate safety-aware training in *DriveSafer* using additional losses defined as follows.

$$L_{\text{total}} = L_{\text{base}} + \lambda_{\text{DAC}} L_{\text{DAC}} + \lambda_{\text{Col}} L_{\text{Col}} + \lambda_{\text{Comf}} L_{\text{Comf}}. \quad (1)$$

$$L_{\text{DAC}} = \frac{1}{BT} \sum_{b=1}^{B} \sum_{t=1}^{T} \Big(1 - M_b(\hat{\mathbf{p}}_{b,t})\Big) \quad (2)$$

$$d_{b,t,n} = \|\hat{\mathbf{p}}_{b,t} - \mathbf{a}_{b,t,n}\|_2 \quad (3)$$

$$L_{\text{Col}} = \frac{1}{BT|\mathcal{N}_a|} \sum_{b=1}^{B} \sum_{t=1}^{T} \sum_{n \in \mathcal{N}_a} \max(0, d_{\text{col}} - d_{b,t,n}) \quad (4)$$

$$\mathbf{v}_{b,t} = \frac{\hat{\mathbf{p}}_{b,t+1} - \hat{\mathbf{p}}_{b,t}}{\Delta t}, \mathbf{j}_{b,t} = \frac{\mathbf{v}_{b,t+2} - 2\mathbf{v}_{b,t+1} + \mathbf{v}_{b,t}}{\Delta t^2} \quad (5)$$

$$L_{\text{Comf}} = \frac{1}{B(T-3)} \sum_{b=1}^{B} \sum_{t=1}^{T-3} \max(0, \ \|\mathbf{j}_{b,t}\|_2 - j_{\text{th}}) \quad (6)$$

Here, L_{base} is the original training loss, B is the batch size denoting the number of scenes, and T is the planning horizon. Among the additional loss terms, L_{DAC} penalizes predicted trajectories from leaving the drivable area, L_{Col} penalizes trajectories that come too close to surrounding agents, and L_{Comf} reduces sudden motion changes to ensure comfort. Additionally, $\hat{\mathbf{p}}_{b,t} \in \mathbb{R}^2$ is the predicted ego position at time step t in scene b, M_b is a binary drivable-area mask for scene b obtained from the underlying map, $\mathcal{N}_a$ is the set of valid surrounding agents and $\mathbf{a}_{b,t,n} \in \mathbb{R}^2$ is the position of agent n at time step t. Furthermore, $d_{b,t,n}$ denotes the Euclidean distance between the ego and agent n, d_{col} is the collision-distance threshold, Δt is the time interval between consecutive trajectory points, $\mathbf{v}_{b,t}$ is the ego velocity, $\mathbf{j}_{b,t}$ is the ego jerk, and j_{th} is the jerk threshold.

Table 1. Comparison of *DriveSafer* with previous end-to-end autonomous driving methods on the `navtest` split of NAVSIM.

Method	Sensor Input	NC↑	DAC↑	TTC↑	Comf.↑	EP↑	PDMS↑
Constant Velocity [8]	-	68.0	57.8	50.0	**100**	19.4	20.6
Ego Status MLP [8]	-	93.0	77.3	83.6	**100**	62.8	65.6
UniAD [11]	Camera	97.8	91.9	92.9	**100**	78.8	83.4
LTF [6]	Camera	97.4	92.8	92.4	**100**	79.0	83.8
PARA-Drive [34]	Camera	97.9	92.4	93.0	99.8	79.3	84.0
LAW [19]	Camera	97.9	92.4	93.0	99.8	79.3	84.0
DriveX [27]	Camera	97.5	94.0	93.0	**100**	79.7	84.5
DRAMA [38]	C & L	98.0	93.1	94.8	**100**	80.1	85.5
Transfuser [6]	C & L	97.7	92.8	92.8	**100**	79.2	84.0
VADv2 [5]	C & L	97.2	89.1	91.6	**100**	76.0	80.9
Hydra-MDP [21]	C & L	98.3	96.0	94.6	**100**	78.7	86.5
WoTE[20]	C & L	**98.5**	96.8	**94.9**	99.9	81.9	88.3
DiffusionDrive [23]	C & L	98.2	96.2	94.7	**100**	82.2	88.1
DriveSafer (Ours)	**C & L**	98.3	**98.7**	94.6	**100**	**84.4**	**90.3**

During training, collision losses are calculated based on ground truth. For inference-time collision awareness, the motion of other road agents must be estimated well. We observe from Figure 3 that many of the catastrophic failures for the baseline DiffusionDrive model involve collisions. However, DiffusionDrive only predicts agent locations for the current time step and the model is not optimized for forecasting the future locations of other agents. Hence, we add a lightweight adapter layer to predict the motion of other agents for future timesteps, enabling joint prediction and planning.

3.3. Inference-Time Safety Guidance

Once the model is trained with safety-aware losses, it will ideally learn to better avoid generating motion plans that violate safety. However, due to the stochastic nature of diffusion-based models, in addition to limitations in modeling sensor noise and complex scenes with many agents, enforcing safety at inference time requires attention.

We start inference-time safety guidance with alternative hypothesis generation. Since a motion plan for 4 seconds can have the ego vehicle moving as much as 100 m at highway speeds, an unsafe motion plan can potentially be made safe with a slight nudge in heading/progress which results in a relatively minor deviation of say 1 m over this time horizon. Motivated by this idea, we generate alternate trajectories that have a variation of 5% in heading, speed and sub-meter lateral trajectory shifts to the top-most predicted mode from the diffusion model. Furthermore, we also use the second- and third-highest ranked modes of the diffusion

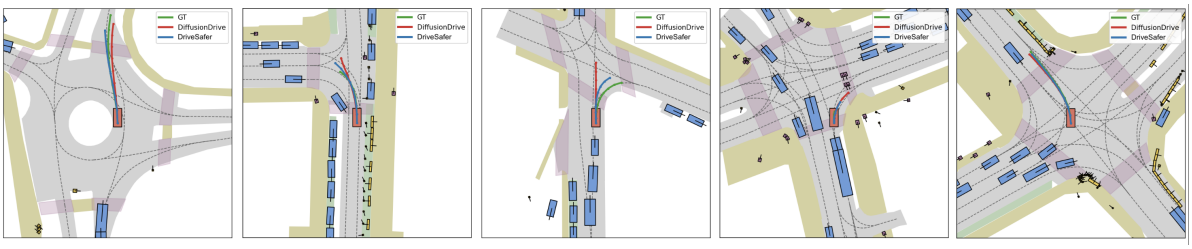

Figure 4. Sample instances where DiffusionDrive predictions had a PDMS score of 0 but were fixed with *DriveSafer* for a safer trajectory.

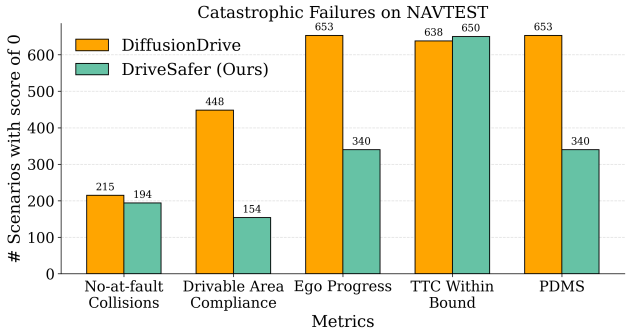

Figure 5. Analysis of catastrophic failure cases (PDM Score = 0) on NAVSIM navtest split. Lower values are better.

model as alternate hypotheses. Once these are generated, we use the predicted positions of agents and the underlying map to assign a score to these alternative trajectories using safety metrics. The safest trajectory is then chosen as the final output, enabling safety guidance during inference.

Note that we assume ego-vehicle localization is correct to sub-lane level accuracy, to be able to use the map information, but this is a reasonable assumption for many modern AVs. Visual localization methods such as MapLocNet [35] and BEVMapMatch [28] have shown that such localization accuracy can be obtained reliably over time using standard datasets such as NuScenes even when GNSS fails.

4. Experiments and Results

We use NAVSIM [8] as our benchmark for evaluating planner outputs due to its variety of complex scenarios in dense urban traffic such as the left turn shown in Figure 2 which requires accurate planning. Table 1 shows a comparative analysis of *DriveSafer* against other state-of-the-art E2E planning methods on the NAVSIM benchmark. While our focus is on improving catastrophic failures, our model also improves the PDMS score by 2.2 points upon the primary SOTA baseline model DiffusionDrive pushing the metric to 90.3. Moreover, we see marked improvements in the Drivable Area Compliance (DAC) and Ego Progress (EP) sub-metrics. These improvements can be attributed to the use of the underlying map for enforcing safety and generating alternative trajectories when the primary trajectory violates drivable area requirements. Additionally, training with safety-aware losses also improves these metrics.

Figure 5 provides a more in-depth analysis of the catastrophic failure cases of DiffusionDrive which have `PDMS=0`. We see a marked improvement with the number of cases obtaining `PDMS=0` being reduced from 653 to 340 ($\approx 48\%$). Moreover, on Drivable Area Compliance (DAC), the improvement is even better with only 154 cases with a DAC of 0 using *DriveSafer*, as compared to 448 using DiffusionDrive, showing over 65% reduction in this category of catastrophic failures.

Among the 313 samples out of 12146 in `navtest` that yield an increase of PDMS score from 0 to a positive value, a few are illustrated in Figure 4. The first instance shows the ground truth ego center trajectory in green and the trajectory produced by DiffusionDrive in red. The red planned trajectory results in the vehicle driving off the paved region, which is unsafe behavior. With *DriveSafer*, an alternate safe path, shown in green, is found. The next two samples show a violation of the intended direction of travel with DiffusionDrive, which also gets corrected with *DriveSafer*. The fourth sample shows a DiffusionDrive-generated path that collides with crossing pedestrians, but *DriveSafer* stops the vehicle early avoiding the collision. In the last example, the red trajectory comes too close to an oncoming vehicle resulting in a collision, but *DriveSafer* steers the trajectory away from such unsafe behavior with a slight nudge in the correct direction.

5. Conclusion

End-to-end autonomous driving models can be powerful, but their practicality is limited by the presence of a significant number of catastrophic failures that limit their real-world usability. We have presented *DriveSafer*, a framework to systematically combat such failures. *DriveSafer* deals with catastrophic failures through a combination of safety-aware losses during training and safety-guided inference. *DriveSafer* reduces the number of catastrophic failures (`PDMS=0`) on the NAVSIM benchmark by 48%, obtaining a PDMS score of 90.3 on the NAVSIM benchmark, outperforming existing state-of-the-art baselines on these metrics. More broadly, our results show that explicit safety guidance in E2E autonomous driving should be geared towards the direct mitigation of catastrophic failure cases. In the future, we will extend *DriveSafer* for real-world deployment on an AV and assess its impact on safety in the real world, aiming to eventually bring catastrophic failures down to zero.

References

[1] Jean Pierre Allamaa, Panagiotis Patrinos, Toshiyuki Ohtsuka, and Tong Duy Son. Real-time mpc with control barrier functions for autonomous driving using safety enhanced collocation. *IFAC-PapersOnLine*, 58(18):392–399, 2024. 2

[2] Holger Caesar, Juraj Kabzan, Kok Seang Tan, Whye Kit Fong, Eric Wolff, Alex Lang, Luke Fletcher, Oscar Beijbom, and Sammy Omari. nuplan: A closed-loop ml-based planning benchmark for autonomous vehicles. *arXiv preprint arXiv:2106.11810*, 2021. 1

[3] Wei Cao, Marcel Hallgarten, Tianyu Li, Daniel Dauner, Xunjiang Gu, Caojun Wang, Yakov Miron, Marco Aiello, Hongyang Li, Igor Gilitschenski, et al. Pseudo-simulation for autonomous driving. *arXiv preprint arXiv:2506.04218*, 2025. 1, 2

[4] Dian Chen, Brady Zhou, Vladlen Koltun, and Philipp Krähenbühl. Learning by cheating. In *Conference on robot learning*, pages 66–75. PMLR, 2020. 2

[5] Shaoyu Chen, Bo Jiang, Hao Gao, Bencheng Liao, Qing Xu, Qian Zhang, Chang Huang, Wenyu Liu, and Xinggang Wang. Vadv2: End-to-end vectorized autonomous driving via probabilistic planning. *arXiv preprint arXiv:2402.13243*, 2024. 2, 3

[6] Kashyap Chitta, Aditya Prakash, Bernhard Jaeger, Zehao Yu, Katrin Renz, and Andreas Geiger. Transfuser: Imitation with transformer-based sensor fusion for autonomous driving. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(11):12878–12895, 2023. 1, 2, 3

[7] Felipe Codevilla, Matthias Müller, Antonio López, Vladlen Koltun, and Alexey Dosovitskiy. End-to-end driving via conditional imitation learning. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 4693–4700. IEEE, 2018. 2

[8] Daniel Dauner, Marcel Hallgarten, Tianyu Li, Xinshuo Weng, Zhiyu Huang, Zetong Yang, Hongyang Li, Igor Gilitschenski, Boris Ivanovic, Marco Pavone, et al. Navsim: Data-driven non-reactive autonomous vehicle simulation and benchmarking. *Advances in Neural Information Processing Systems*, 37:28706–28719, 2024. 1, 2, 3, 4

[9] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *Conference on robot learning*, pages 1–16. PMLR, 2017. 1, 2

[10] Haoyu Fu, Diankun Zhang, Zongchuang Zhao, Jianfeng Cui, Dingkang Liang, Chong Zhang, Dingyuan Zhang, Hongwei Xie, Bing Wang, and Xiang Bai. Orion: A holistic end-to-end autonomous driving framework by vision-language instructed action generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 24823–24834, 2025. 2

[11] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhai Wang, et al. Planning-oriented autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 17853–17862, 2023. 2, 3

[12] Bernhard Jaeger, Daniel Dauner, Jens Beißwenger, Simon Gerstenecker, Kashyap Chitta, and Andreas Geiger. Carl: Learning scalable planning policies with simple rewards. *arXiv preprint arXiv:2504.17838*, 2025. 2

[13] Xiaosong Jia, Yulu Gao, Li Chen, Junchi Yan, Patrick Langechuan Liu, and Hongyang 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*, pages 7953–7963, 2023. 2

[14] Xiaosong Jia, Zhenjie Yang, Qifeng Li, Zhiyuan Zhang, and Junchi Yan. Bench2drive: Towards multi-ability benchmarking of closed-loop end-to-end autonomous driving. *Advances in Neural Information Processing Systems*, 37:819–844, 2024. 1, 2

[15] Bo Jiang, Shaoyu Chen, Qing Xu, Bencheng Liao, Jiajie Chen, Helong Zhou, Qian Zhang, Wenyu Liu, Chang Huang, and Xinggang Wang. Vad: Vectorized scene representation for efficient autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8350, 2023. 2

[16] Rohit Konda, Eric Squires, Pietro Pierpaoli, Magnus Egerstedt, and Samuel Coogan. Provably-safe autonomous navigation of traffic circles. In *2019 IEEE Conference on control technology and applications (CCTA)*, pages 876–881. IEEE, 2019. 2

[17] Derun Li, Jianwei Ren, Yue Wang, Xin Wen, Pengxiang Li, Leimeng Xu, Kun Zhan, Zhongpu Xia, Peng Jia, Xianpeng Lang, et al. Finetuning generative trajectory model with reinforcement learning from human feedback. *arXiv e-prints*, pages arXiv–2503, 2025. 2

[18] Qifeng Li, Xiaosong Jia, Shaobo Wang, and Junchi Yan. Think2drive: Efficient reinforcement learning by thinking with latent world model for autonomous driving (in carla-v2). In *European conference on computer vision*, pages 142–158. Springer, 2024. 2

[19] Yingyan Li, Lue Fan, Jiawei He, Yuqi Wang, Yuntao Chen, Zhaoxiang Zhang, and Tieniu Tan. Enhancing end-to-end autonomous driving with latent world model. *arXiv preprint arXiv:2406.08481*, 2024. 3

[20] Yingyan Li, Yuqi Wang, Yang Liu, Jiawei He, Lue Fan, and Zhaoxiang Zhang. End-to-end driving with online trajectory evaluation via bev world model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 27137–27146, 2025. 3

[21] Zhenxin Li, Kailin Li, Shihao Wang, Shiyi Lan, Zhiding Yu, Yishen Ji, Zhiqi Li, Ziyue Zhu, Jan Kautz, Zuxuan Wu, et al. Hydra-mdp: End-to-end multimodal planning with multi-target hydra-distillation. *arXiv preprint arXiv:2406.06978*, 2024. 3

[22] Zhiqi Li, Zhiding Yu, Shiyi Lan, Jiahan Li, Jan Kautz, Tong Lu, and Jose M. Alvarez. Is ego status all you need for open-loop end-to-end autonomous driving? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14864–14873, 2024. 1

[23] Bencheng Liao, Shaoyu Chen, Haoran Yin, Bo Jiang, Cheng Wang, Sixu Yan, Xinbang Zhang, Xiangyu Li, Ying Zhang, Qian Zhang, et al. Diffusiondrive: Truncated diffusion model for end-to-end autonomous driving. In *Proceedings of the*

Computer Vision and Pattern Recognition Conference, pages 12037–12047, 2025. 1, 2, 3

[24] Long Nguyen, Micha Fauth, Bernhard Jaeger, Daniel Dauner, Maximilian Igl, Andreas Geiger, and Kashyap Chitta. Lead: Minimizing learner-expert asymmetry in end-to-end driving. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026. 2

[25] Katrin Renz, Long Chen, Elahe Arani, and Oleg Sinavski. Simlingo: Vision-only closed-loop autonomous driving with language-action alignment. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 11993–12003, 2025. 2

[26] Shuyao Shang, Yuntao Chen, Yuqi Wang, Yingyan Li, and Zhaoxiang Zhang. Drivedpo: Policy learning via safety dpo for end-to-end autonomous driving. *arXiv preprint arXiv:2509.17940*, 2025. 2

[27] Chen Shi, Shaoshuai Shi, Kehua Sheng, Bo Zhang, and Li Jiang. Drivex: Omni scene modeling for learning generalizable world knowledge in autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 28599–28609, 2025. 3

[28] Shounak Sural and Ragunathan Rajkumar. Bevmapmatch: Multimodal bev neural map matching for robust re-localization of autonomous vehicles. *arXiv preprint arXiv:2603.25963*, 2026. 4

[29] Yingqi Tang, Zhuoran Xu, Zhaotie Meng, and Erkang Cheng. Hip-ad: Hierarchical and multi-granularity planning with deformable attention for autonomous driving in a single decoder. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 25605–25615, 2025. 2

[30] Royden Wagner, Omer Sahin Tas, Jaime Villa, Felix Hauser, Yinzhe Shen, Marlon Steiner, Dominik Strutz, Carlos Fernandez, Christian Kinzig, Guillermo S Guitierrez-Cabello, et al. Longtail driving scenarios with reasoning traces: The kitscenes longtail dataset. *arXiv preprint arXiv:2603.23607*, 2026. 1

[31] Xiaofeng Wang, Zheng Zhu, Guan Huang, Xinze Chen, Jiagang Zhu, and Jiwen Lu. Drivedreamer: Towards real-world-drive world models for autonomous driving. In *European conference on computer vision*, pages 55–72. Springer, 2024. 2

[32] Yuqi Wang, Jiawei He, Lue Fan, Hongxin Li, Yuntao Chen, and Zhaoxiang Zhang. Driving into the future: Multiview visual forecasting and planning with world model for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14749–14759, 2024. 2

[33] Yan Wang, Wenjie Luo, Junjie Bai, Yulong Cao, Tong Che, Ke Chen, Yuxiao Chen, Jenna Diamond, Yifan Ding, Wenhao Ding, et al. Alpamayo-r1: Bridging reasoning and action prediction for generalizable autonomous driving in the long tail. *arXiv preprint arXiv:2511.00088*, 2025. 2

[34] Xinshuo Weng, Boris Ivanovic, Yan Wang, Yue Wang, and Marco Pavone. Para-drive: Parallelized architecture for real-time autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15449–15458, 2024. 2, 3

[35] Hang Wu, Zhenghao Zhang, Siyuan Lin, Xiangru Mu, Qiang Zhao, Ming Yang, and Tong Qin. Maplocnet: Coarse-to-fine feature registration for visual re-localization in navigation maps. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 13198–13205. IEEE, 2024. 4

[36] Jianye Xu, Chang Che, and Bassam Alrifaee. A real-time control barrier function-based safety filter for motion planning with arbitrary road boundary constraints. *arXiv preprint arXiv:2505.02395*, 2025. 2

[37] Runsheng Xu, Hubert Lin, Wonseok Jeon, Hao Feng, Yuliang Zou, Liting Sun, John Gorman, Ekaterina Tolstaya, Sarah Tang, Brandyn White, et al. Wod-e2e: Waymo open dataset for end-to-end driving in challenging long-tail scenarios. *arXiv preprint arXiv:2510.26125*, 2025. 1, 2

[38] Chengran Yuan, Zhanqi Zhang, Jiawei Sun, Shuo Sun, Zefan Huang, Christina Dao Wen Lee, Dongen Li, Yuhang Han, Anthony Wong, Keng Peng Tee, et al. Drama: An efficient end-to-end motion planner for autonomous driving with mamba. *arXiv preprint arXiv:2408.03601*, 2024. 3

[39] Jiang-Tian Zhai, Ze Feng, Jinhao Du, Yongqiang Mao, Jiang-Jiang Liu, Zichang Tan, Yifu Zhang, Xiaoqing Ye, and Jingdong Wang. Rethinking the open-loop evaluation of end-to-end autonomous driving in nuscenes. *arXiv preprint arXiv:2305.10430*, 2023. 1

[40] Dapeng Zhang, Fei Shen, Rui Zhao, Yinda Chen, Peng Zhi, Chenyang Li, Rui Zhou, and Qingguo Zhou. Cocvla: Delving into adversarial domain transfer for explainable autonomous driving via chain-of-causality visual-language-action model. *arXiv preprint arXiv:2511.19914*, 2025. 2

[41] Linrui Zhang, Zhenghao Peng, Quanyi Li, and Bolei Zhou. Cat: Closed-loop adversarial training for safe end-to-end driving. In *Conference on Robot Learning*, pages 2357–2372. PMLR, 2023. 2

[42] Wenzhao Zheng, Weiliang Chen, Yuanhui Huang, Borui Zhang, Yueqi Duan, and Jiwen Lu. Occworld: Learning a 3d occupancy world model for autonomous driving. In *European conference on computer vision*, pages 55–72. Springer, 2024. 2

[43] Yinan Zheng, Ruiming Liang, Kexin Zheng, Jinliang Zheng, Liyuan Mao, Jianxiong Li, Weihao Gu, Rui Ai, Shengbo Eben Li, Xianyuan Zhan, et al. Diffusion-based planning for autonomous driving with flexible guidance. *arXiv preprint arXiv:2501.15564*, 2025. 2

[44] Zewei Zhou, Tianhui Cai, Seth Z Zhao, Yun Zhang, Zhiyu Huang, Bolei Zhou, and Jiaqi Ma. Autovla: A vision-language-action model for end-to-end autonomous driving with adaptive reasoning and reinforcement fine-tuning. *arXiv preprint arXiv:2506.13757*, 2025. 2