

# Data Scaling Laws for Imitation Learning-Based End-to-End Autonomous Driving

Yupeng Zheng<sup>1,2,3</sup>, Pengxuan Yang<sup>2</sup>, Zhongpu Xia<sup>1\*</sup>, Qichao Zhang<sup>1,2</sup>✉, Yuhang Zheng<sup>4</sup>, Bu Jin<sup>2</sup>, Teng Zhang<sup>3</sup>, Ben Lu<sup>3</sup>, Chao Han<sup>3</sup>, Weize Li<sup>6</sup>, Songen Gu<sup>7</sup>, Xianpeng Lang<sup>3</sup>, Xiangyuan Lan<sup>5</sup>, and Dongbin Zhao<sup>1,2</sup>  
<sup>1</sup>CASIA, <sup>2</sup>UCAS, <sup>3</sup>Li Auto, <sup>4</sup>NUS, <sup>5</sup>PCL, <sup>6</sup>Clemson, <sup>7</sup>FDU  
 ✉ Corresponding Author, \* Project Leader

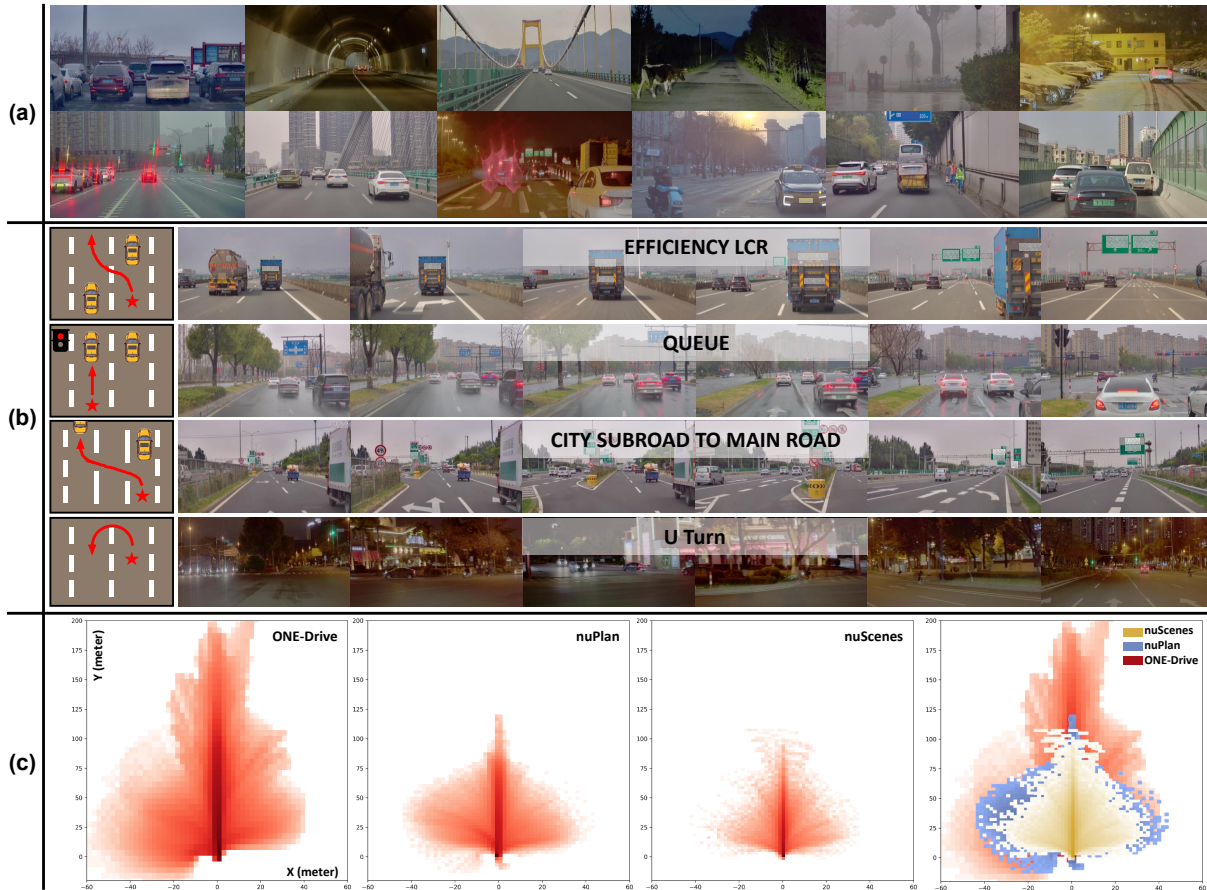


Fig. 1: We have collected and utilized a large-scale, diverse real-world dataset, enabling us to investigate data scaling laws of end-to-end driving. Figure (a) illustrates the diversity of our dataset, encompassing various weather conditions, road types, and traffic scenarios. Figure (b) presents 23 scenario types that we have identified to conduct in-depth analyses of the impact of data scale on generalization and the importance of data distribution. Figure (c) compares the trajectory distributions of existing datasets nuScenes [1] and nuPlan [2] with ours. Our trajectory distribution exhibits greater diversity, including a higher proportion of high-speed driving, turning, and lane changes.

**Abstract**—The end-to-end autonomous driving paradigm has recently attracted lots of attention due to its scalability. However, existing methods are constrained by the limited scale of real-world data, which hinders a comprehensive exploration of the scaling laws associated with end-to-end autonomous driving. To address this issue, we collected substantial data from various driving scenarios and behaviors and conducted an extensive study on the scaling laws of existing imitation learning-based end-to-end autonomous driving paradigms. Specifically, approximately 4 million demonstrations from 23 different scenario types were gathered, amounting to over 30,000 hours of driving demonstrations. We performed open-loop evaluations and closed-loop simulation evaluations in 1,400

diverse driving demonstrations (1,300 for open-loop and 100 for closed-loop) under stringent assessment conditions. Through experimental analysis, we discovered that (1) the performance of the driving model exhibits a power-law relationship with the amount of data, but this is not the case in closed-loop evaluation. The inconsistency between the two assessments shifts our focus toward the distribution of data rather than merely expanding its volume. (2) A small increase in the quantity of long-tailed data can significantly improve the performance for the corresponding scenarios; (3) appropriate scaling of data enables the model to achieve combinatorial generalization in novel scenes and actions. Our results highlight the critical role of data scaling in improving the generalizability of models across

diverse driving scenarios, assuring safe deployment in the real world. Project repo: <https://github.com/ucaszyp/Driving-Scaling-Law>

## I. INTRODUCTION

End-to-end autonomous driving has gained significant attention in recent years. It typically employs a differentiable model that takes raw sensor data as input and generates a potential planning trajectory as output. This paradigm allows for the direct optimization of the entire system in a data-driven manner, offering scalability where performance improvements can be achieved by increasing training data, as exemplified by scaling laws [3]–[5].

Some previous works [6], [7] have tried to investigate the effect of scaling up. However, they share the same challenge: insufficient real-world data, resulting in data scaling laws in end-to-end autonomous driving remaining underexplored. As shown in the table I, the open-source datasets [1], [2], [8] are typically thousand-scale, far less than the million-scale or billion-scale vision-language data in language models or generative models. Although simulators like CARLA [9] offer a promising solution to the data lack, the huge domain gap hinders their real-world application. Consequently, it remains under-explored whether end-to-end autonomous driving has such common data scaling laws and how autonomous driving vehicles can benefit from the laws.

In this paper, we delve into the data scaling laws of end-to-end autonomous driving in the real world. We aim to investigate the three critical questions:

- *Is there a data scaling law in the field of end-to-end autonomous driving?*
- *How does data quantity influence model performance when scaling training data?*
- *Can data scaling endow autonomous driving cars with the generalization to new scenarios?*

To answer the questions, we collect and annotate a million-scale dataset named ONE-Drive, which contains over 4 million driving demonstrations (about 30,000 hours) of real-world data in diverse cities and road conditions, as shown in the figure. 1. Based on the dataset, we first conduct a detailed analysis of the relationship between training data and model performance. Further, we split the scenarios into 23 types based on the traffic condition and agent behavior to delve into the generalization of the data scaling law. Lastly, we set the training data with different distributions on scenario types to analyze the relationship between data scaling and data distribution. For a comprehensive analysis, we conduct the open-loop test, closed-loop test, and real-world deployment to evaluate the planning results.

Throughout our experiments, we surprisingly find:

- **Scaling law.** The data scaling law in open-loop evaluation exhibits a power-law relationship with the amount of data, but this is not the case in closed-loop evaluation. The inconsistency between the two evaluations shifts our focus toward the distribution of data rather than merely expanding its volume. (Section V-A).

- **Importance of increasing the quantity of targeted data.** Targeted augmentation of long-tailed scenarios exhibits superior efficiency compared to indiscriminate expansion of the overall dataset. (Section V-B).
- **Combinatorial generalization.** With the scaling of training data, models gradually acquire combinatorial generalization capabilities, enabling them to combine known information to achieve generalizable planning for new scenarios and actions (Section V-C).

TABLE I: Comparison between the previous dataset and our introduced ONE-Drive for the end-to-end autonomous driving task.

Dataset	Setting		Source	Scale (demonstrations)
	Open-loop	Visual closed-loop		
Bench2drive [8]		✓	Simulator	10k
nuScenes [1]	✓		Real-world	1k
nuPlan [2]	✓		Real-world	20k
ONE-Drive (Ours)	✓	✓	Real-world	<b>4M</b>

## II. RELATED WORK

### A. End-to-end Autonomous Driving

Recent years have witnessed the development of autonomous driving vehicles [6], [7], [11]–[15]. Among them, end-to-end approaches catch great attention for their direct optimization of the whole system, offering a potential solution for the information loss and cascade error of traditional module-based approaches. They typically take raw sensor data as input and generate a potential path or plan as output. Some previous works focus on end-to-end solutions in simulators like CARLA [9]. For example, TransFuser [13] and InterFuser [15] combine information from different sensors to enhance the robustness and performance of the autonomous driving model in complex scenarios by fusing perception and planning. ReasonPlan [16] and SimLingo [7] achieve exceptional performance with vision-language reasoning. However, the domain gap between the simulator and the real world hinders their industrial practice. To avoid the domain gap, some work explores end-to-end autonomous planning in the real world. For example, UniAD [11], VAD [12], and UncAD [17] utilize Bird’s Eye View (BEV) queries to unify different tasks with the transformer. More recently, World4Drive [18] and WorldRFT [19] propose world model-based policies that enable end-to-end planning from the perspective view. Such data-driven optimization enables the ability to improve the system by simply scaling training resources. The end-to-end approaches offer a potential solution for the information loss and cascade error of traditional module-based approaches.

### B. Scaling Laws

In recent years, the development of foundation models has highlighted the significance of scaling laws in the field of large language models [3], vision generative models [5], [20]–[24], and robotics [25]–[29]. These laws elucidate the relationship between dataset size, model size, and performance, showing that the effectiveness of transformer-based

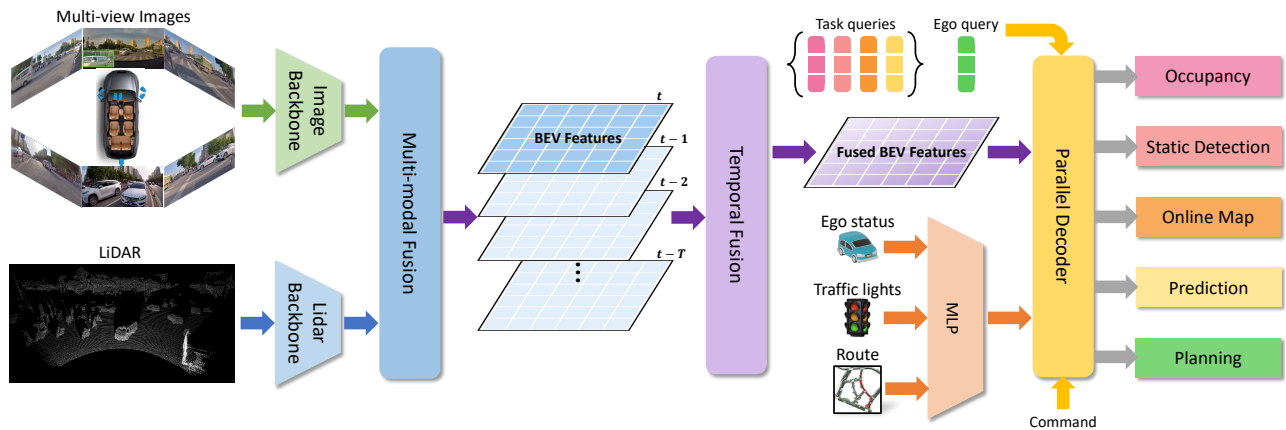


Fig. 2: To better investigate the data scaling laws of end-to-end algorithms, we implement a parallel baseline methodology. This approach, inspired by the PARA-Drive [10] framework, offers enhanced training stability and computational efficiency.

models follows a power-law function relative to the number of model parameters, the volume of training data, and the computational resources used for training. This concept has also been extended to other domains, such as image and video synthesis [4], allowing for a more efficient trade-off between resource allocation and performance outcomes. However, there is limited research on the application of scaling laws to end-to-end autonomous driving, primarily due to constraints related to data availability. It still remains uncertain whether any scaling law exists that links the scalability of neural networks to planning results in autonomous driving, which is the main focus of our work.

### III. METHOD

Following PARA-Drive [10], we implement a parallel modular end-to-end autonomous driving algorithm named the ONE model, as shown in the figure. 2. The input is LiDAR point clouds and panoramic RGB images, and the expected output is a planning trajectory.

#### A. BEV Encoder

Given the input of multi-view images and point clouds, we generate a unified BEV representation through the View Transformation module and the Temporal Fusion module.

**View Transformation.** For multi-view camera images, we utilize an image backbone to extract image features. Then we utilize LSS [30] to project these image features into a 3D frustum form. For LiDAR point clouds, we extract voxelized LiDAR features with 3D sparse convolution layers. The features are then flattened along the height dimension, leading to features in a BEV plane. Finally, the image features and point cloud features are fused with squeeze-and-excitation (SE) blocks [31].

**Temporal Fusion.** To effectively leverage historical information, we employ a temporal fusion module. We set up a FIFO queue to sequentially extract and store past BEV features and relative poses. During training and inference, the features stored in the queue are transformed using the pose information to align with the current pose. The transformed features are then fused with the current BEV features by

SE blocks [31]. Finally, the queue is updated by adding the current frame features to the end of the queue and dequeuing the front element of the queue.

#### B. Parallel Decoder

Our decoder comprises **five tasks**: online map prediction, 4D occupancy prediction, static detection, motion prediction, and planning. Inspired by [11], [12], [14], we design each module with learnable query features tailored specifically for its corresponding task.

**Online Mapping.** For online map prediction, we utilize a set of map queries to estimate a vector map from BEV features, along with the class scores corresponding to each map vector. The road map labels include lane dividers, lane centerlines, crosswalks, and stop lines. We use the loss function defined in MapTR [32].

**4D Occupancy Prediction.** For occupancy, we establish a set of occupancy queries to estimate the 3D occupancy and the corresponding flow for the next 3 seconds from BEV features. The 3D occupancy prediction loss is the same as [33] and the flow loss employed in Cam4DOcc [34].

**Static Detection.** We specifically design an additional static object detection task to address the challenge of capturing small static objects in occupancy prediction. We employ Hungarian matching to match the query features with ground truth values.

**Motion Prediction.** For dynamic objects, we follow [12] by setting a set of agent queries. Each agent query interacts with the BEV features through attention layers to capture environmental information. Similarly, we supervise the motion prediction of dynamic objects using the L1 loss.

**Ego Motion Planning.** In the planning module, to fulfill the requirements of real-world planning for precise adherence to navigation, we employ a multi-layer perceptron (MLP) to encode the traffic light status, road-level route (represented by a collection of dense waypoints at the road level), and the motion state of the ego-vehicle, which serve as contextual information for planning. Following [10], [12], we leverage learnable embeddings to construct ego-vehicle




Fig. 3: Distribution of the 23 scenario types.

TABLE II: Training steps and resources of different training dataset sizes (demonstrations).

Dataset Size (Demonstrations)	Training Resource	Training Steps
10 thousand	64 A100	$6.0 \times 10^5$
50 thousand	64 A100	$8.6 \times 10^5$
0.2 million	128 A100	$2.9 \times 10^6$
0.7 million	320 A100	$9.3 \times 10^6$
2 million	512 A100	$2.2 \times 10^7$
4 million	512 A100	$3.1 \times 10^7$

queries. These queries are passed as input to the cross-attention layers to interact with BEV features. Then the queries are used to predict multi-modal future planning trajectories with corresponding scores. During inference, we execute the trajectory with the highest confidence score.

## IV. EXPERIMENTS SETUP

### A. Implementation Details

**Data Collection.** The scene images were captured horizontally by seven cameras covering a  $360^\circ$  field of view (FOV), with each image having a resolution of  $3840 \times 2160$ . Additionally, we employed a 128-beam LiDAR sensor to acquire point cloud of the environment.

**Baseline.** Our method predicts a 6-second future trajectory with 2 seconds of historical information. It uses ResNet34 [35] as the default image backbone. The image resolution is downsampled to  $512 \times 960$  for training. The default BEV resolution is  $232 \times 80$  for a perception range of  $139.2m \times 48m$  longitudinally and laterally. The default hidden state and embedding dimension are 256. We construct five training datasets with varying numbers of demonstrations: **10 thousand**, **50 thousand**, **0.2 million**, **0.7 million**, **2 million**, and **4 million**. We train each model to full convergence during training for a fair comparison. This allows models trained on different dataset sizes to undergo different numbers of training steps. Training steps and computational resources are shown in the table. II

### B. Data Mining

To facilitate more flexible adjustment of data distribution, we define 23 scenario types according to the agent behavior and traffic conditions. They can be acquired based on the following meta-information:

- **Navigation Information:** This includes upcoming road names, road types, number of lanes, and distance to intersections.
- **Static Perception:** This includes lane markings, driveable areas, intersection areas, and pedestrian walkways.
- **Dynamic Perception:** This covers road obstacle information, including position, velocity, dimensions, historical trajectories, and predicted future trajectories.
- **Ego Vehicle Status:** This comprises the ego vehicle’s position, velocity, acceleration, historical trajectory, and future trajectory.

The distribution of these scenario types in the training dataset (4 million demonstrations) is shown in the figure. 3. For more details about the specific characteristics of the 23 scenario types, please see the supplementary materials.

### C. Evaluation

To comprehensively analyze the scaling laws of data, we evaluate our model using three methodologies: open-loop evaluation and closed-loop simulation. It is important to acknowledge that, owing to constraints in the reconstruction process, the closed-loop evaluation environment has not been fully aligned with the open-loop. Future work will align the two evaluation settings to make a fair comparison and analysis of the data scaling law in these two settings.

**Open-loop Evaluation.** The open-loop evaluation computes the Average Displacement Error (ADE) metric between the predicted trajectories and the ground truth trajectories, which can evaluate the network’s ability to fit the ground truth trajectory.

**Closed-loop Simulation.** For closed-loop simulation, we employ 3D-GS [36] to reconstruct partial test scenarios,




Fig. 4: Results of open-loop [(a)] and closed-loop [(b) and (c)] evaluation. (a): Data scaling laws on ADE. Dashed lines represent power-law fits, with the equations and coefficient  $r$  provided in the eq. 1. All axes are shown on a logarithmic scale. (b): Scores of each metric in closed-loop evaluation. (c): Total score of closed-loop evaluation. In (b) and (c), the X-axis (number of demonstrations) is displayed on a logarithmic scale, and the Y-axis is displayed on a linear scale.

enabling end-to-end simulation with visual closed-loop feedback. This simulation evaluates the driving trajectory in five dimensions: **safety**, **comfort**, **rule**, **efficiency**, and **navigation**. For each metric, a higher score indicates better performance. Finally, the total score is weighted by these five scores.

$$\begin{aligned} \text{score} = & 0.25 \times \text{safety} + 0.15 \times \text{comfort} + 0.2 \times \text{rule} \\ & + 0.25 \times \text{efficiency} + 0.15 \times \text{navigation} \end{aligned}$$

## V. EXPERIMENTS

In this section, we investigate the relationship between the scale of training data and the ability of end-to-end driving models to fit expert trajectories (Section V-A.1). Subsequently, we discover the inconsistency in data scaling laws between open-loop and closed-loop evaluations, revealing that merely increasing data volume does not indefinitely improve autonomous driving planning performance in real-world scenarios (Section V-A.2). This finding motivates us to explore data scaling strategies beyond simply augmenting data quantity. We examine how data distribution affects the efficiency of the data scaling law (Section V-B). Finally, we demonstrate that end-to-end autonomous driving systems, when scaled appropriately with data, exhibit combinatorial generalization capabilities.

### A. Unveiling of Data Scaling Laws

1) *Power-law Relationship*: Inspired by the power-law data scaling laws observed in large-scale language models [3], we aim to investigate the existence of power-law scaling relationships in imitation learning-based end-to-end autonomous driving. We begin our analysis by examining the Average Displacement Error (ADE) between ground truth trajectories and predicted trajectories in different training dataset scales. Specifically, we utilize open-loop evaluation results from five models with varying training dataset scales, as detailed in Section IV. We designate the amount of training data (quantified by the number of demonstrations used in training) as a variable  $X$  and the normalized trajectory ADE as a variable  $Y$ . After applying the logarithmic

transformation to both  $X$  and  $Y$ , the presence of a linear relationship would indicate a power-law scaling law in end-to-end autonomous driving data.

As shown in the figure. 4 (a), we conduct the linear model fitting of the log-transformed data. The fitting results yield:

$$Y = 0.6833 \cdot X^{-0.188}, r = -0.963 \quad (1)$$

The correlation coefficient  $r$  of the fit is -0.963, strongly suggesting the existence of a power-law data scaling law in trajectory fitting for imitation learning-based end-to-end autonomous driving.

2) *Differences Between Open-loop And Closed-loop*: Given that open-loop evaluation cannot fully reflect a vehicle’s planning performance in real-world scenarios, we further assess the data scaling laws in a closed-loop condition. Specifically, we employ a visual closed-loop simulator to evaluate the vehicle’s closed-loop performance of the models trained on variant scales of data. Following the description in Sec IV-C, we obtain the five closed-loop scores and the total score as illustrated in the figure. 4 (b) and (c).

From figure (b), we observe distinct relationships between various closed-loop evaluation metrics and the expansion of data volume. Metrics such as navigation and efficiency demonstrate a clear and stable increasing trend, as they are directly related to trajectory and path training. In contrast, the score for the driving rule, while showing an overall upward trend, exhibits instability during its growth. This occurs because driving rules are not explicitly modeled as loss functions in imitation learning—instead, the model must implicitly learn the patterns from large volumes of image-trajectory pairs. Due to inherent limitations in imitation learning, such as causal confusion, performance in rule compliance remains suboptimal. Similarly, metrics including safety and comfort experience rapid improvement initially, but their growth rates gradually slow within the range of 200k to 700k data samples. We hypothesize that in the early stage, increasing data volume enhances the model’s ability to fit trajectories, leading to swift gains in closed-loop safety. In later stages, however, minor discrepancies

in trajectory fitting have a diminishing impact on safety in most scenarios. Thus, further improving safety may depend more on optimizing data distribution to enhance performance across diverse situations, rather than simply increasing the amount of data.

### B. Data Quantity and Model Performance

Based on the foundation of the data scaling law, we explore the relationship between data quantity and model performance when increasing the scale of training data. Specifically, we categorized our dataset of 2 million demonstrations into 23 scenario types, as defined in Section IV-B, and analyzed the proportion of each type. To simplify the process, we identify two scenario types that exhibit poor performance (DOWN\_SUBROAD\_LC and SUBROAD\_TO\_MAINROAD\_HIGHWAY) from the 23 categories for in-depth study. By incrementally increasing the representation of these scenario types, we observe the consequent effects on planning performance.

As illustrated in the table. III, we maintained a constant overall volume of training data while increasing the quantity of these two scenario types in each iteration. Subsequently, we evaluated the model’s open-loop trajectory error after training. The results reveal that for these long-tailed scenarios, doubling the specific training data while keeping the total data volume constant led to improvements in planning performance ranging from 9.7% to 16.9%. Furthermore, quadrupling the specific training data yielded even more substantial gains, with performance improvements between 22.8% and 32.9%. Our findings suggest that even a relatively small increase in scenario-specific data (comprising only several hundred or thousands of demonstrations) can lead to significant enhancements in planning performance for these challenging scenarios.

### C. Combinatorial Generalization

In this section, we investigate the relationship between data scaling laws and the generalization capability, which is considered essential for deploying autonomous driving in real-world scenarios.

1) *Quantitative Analysis:* To demonstrate the existence of generalization, we split two categories (HIGHWAY\_NUDGE\_OBS and WAIT\_TURN) from the 23 scenario types defined in Sec IV-B as test categories, while using the remaining 21 categories to form the training set. To be consistent with previous settings, we trained models using three scales of the number of training data: **50 thousand**, **0.7 million**, and **2 million** demonstrations. Notably, we applied a strict filtering strategy when selecting test data for these two categories to ensure that each scenario does not overlap with other scenario types. The open-loop test results are shown in the figure. 5. For a convenient comparison, we select scenarios similar to these two scenario types.

We observed that (1) the model trained on **50 thousand** demonstrations showed larger displacement error from expert trajectories in the two test scenarios compared to similar scenarios, indicating insufficient generalization capability

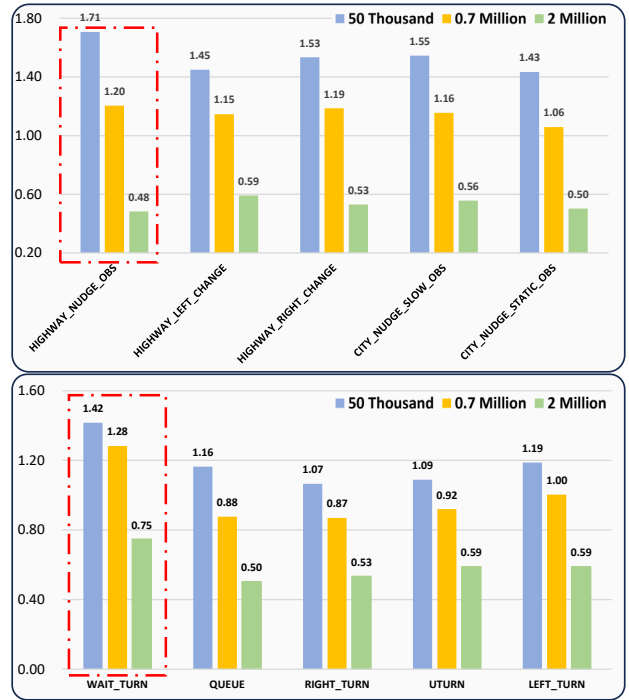


Fig. 5: To investigate the relationship between data scaling and generalization ability, we select data from 2 scenario types, HIGHWAY\_NUDGE\_OBS and WAIT\_TURN as test data. Models trained with **50 thousand**, **0.7 million**, and **2 million** demonstrations are represented in **blue**, **yellow**, and **green**, respectively.

with small-scale training data. (2) As the training data gradually increased to **green** demonstrations, the displacement error between the trajectories of the two test scenarios and other scenarios rapidly narrowed. The performance on HIGHWAY\_NUDGE\_OBS even surpasses the performance on other scenarios that participated in the training. (3) By learning high-speed driving and low-speed nudging obstacles separately from the training data, the model acquired the ability to generalize to high-speed HIGHWAY\_NUDGE\_OBS scenarios; through learning to turn and queuing at red lights, the model developed the capability to generalize to WAIT\_TURN scenarios. Based on these observations, we argue that **the data scaling law endows models with the ability to perform combinatorial generalization.**

2) *Qualitative Analysis:* To demonstrate the generalization ability more intuitively, we also provide a qualitative visualization of the predicted trajectory and ground truth map in the figure. 6, which shows the planning results of HIGHWAY\_NUDGE\_OBS and WAIT\_TURN scenarios. The BEV maps in the **green boxes** represent the planning result of the model trained using 50 thousand demonstrations without HIGHWAY\_NUDGE\_OBS and WAIT\_TURN scenarios (**50K + Unseen** in the legend), the BEV maps in the **blue boxes** represent the planning result of the model trained using 2 million demonstrations without HIGHWAY\_NUDGE\_OBS and WAIT\_TURN scenarios (**2M + Unseen** in the legend), and the BEV maps in the **orange boxes** represent the planning result of the model trained using 2 million demon-

TABLE III: Analysis of different scene distribution.

Scenario Type	Data Quantity			ADE ↓		
DOWN_SUBROAD_LC	4972	7218 (+45.1%)	21836 (+339.2%)	0.788	0.711 (-9.7%)	0.529 (-32.9%)
SUBROAD_TO_MAINROAD_HIGHWAY	643	1485 (+130.9%)	2569 (+299.5%)	0.786	0.653 (-16.9%)	0.593 (-22.8%)

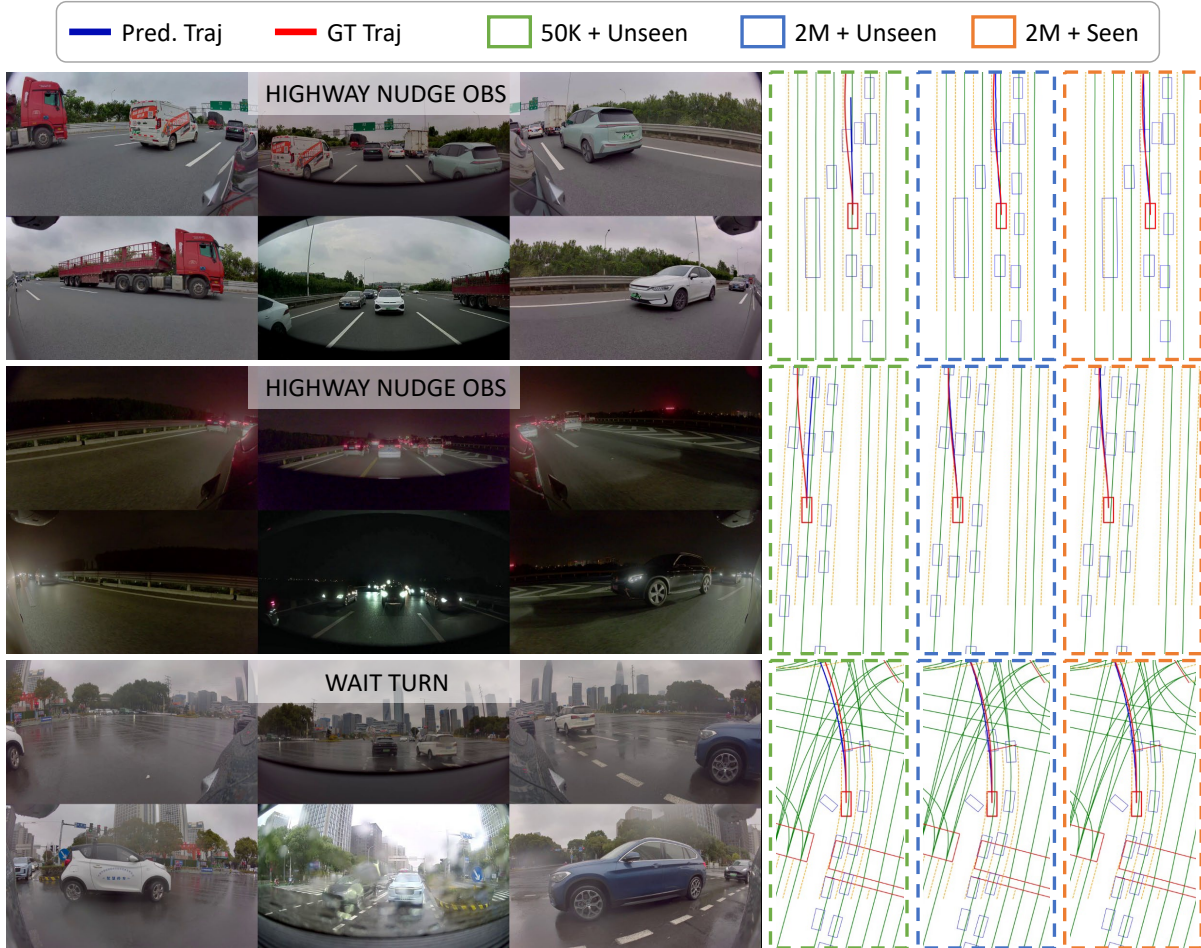


Fig. 6: Qualitative results on scenarios of HIGHWAY\_NUDGE\_OBS and WAIT\_TURN scenario types. We compare the planning trajectories generated by three different models, demonstrating the combinatorial generalization in new driving scenarios.

strations with HIGHWAY\_NUDGE\_OBS and WAIT\_TURN scenarios ( $2M + Seen$  in the legend).

Visual analysis reveals that appropriately increasing the scale of training data enables the model to achieve combinatorial generalization to novel scenarios. This enhanced generalization capability allows the model to perform competitively with counterparts specifically trained in these new scenarios. Our findings underscore the critical role of data scaling in improving model adaptability and robustness across diverse autonomous driving contexts.

## VI. CONCLUSION

### A. Conclusion

In this paper, we delve into the data scaling laws of the imitation learning-based end-to-end autonomous driving framework. Upon further investigation, we uncovered three intriguing findings:

- A power-law data scaling law in the open-loop metric, but different in the closed-loop metric.
- Data distribution plays a key role in the data scaling law of end-to-end autonomous driving.
- The data scaling law endows the model with combinatorial generalization, which powers the model’s zero-shot ability for new scenarios.

## VII. LIMITATION AND FUTURE WORK

While our work provides significant insights, we acknowledge several limitations. Due to computational constraints, our investigation primarily focused on imitation learning methods based on mean regression with BEV representation. To address these limitations and further enrich our understanding of data scaling laws in autonomous driving, we propose the following directions for future research:

- exploring scaling laws for a broader range of model

architecture (visuomotor policies and vision-language-action models);

- investigating how self-supervised pre-training changes the scaling law;

These avenues could provide valuable insights into the scalability and generalizability of autonomous driving systems across diverse scenarios and model paradigms.

## VIII. ACKNOWLEDGMENTS

This work is supported by the National Key Research and Development Program of China under Grant 2022YFA1004000 and the Beijing Natural Science Foundation under Grant 4242052.

## REFERENCES

- [1] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020.
- [2] H. Caesar, J. Kabzan, K. S. Tan, W. K. Fong, E. Wolff, A. Lang, L. Fletcher, O. Beijbom, and S. Omari, "nuPlan: A closed-loop ml-based planning benchmark for autonomous vehicles," *arXiv preprint arXiv:2106.11810*, 2021.
- [3] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, "Scaling laws for neural language models," *arXiv preprint arXiv:2001.08361*, 2020.
- [4] K. Tian, Y. Jiang, Z. Yuan, B. Peng, and L. Wang, "Visual autoregressive modeling: Scalable image generation via next-scale prediction," *Advances in neural information processing systems*, 2024.
- [5] T. Henighan, J. Kaplan, M. Katz, M. Chen, C. Hesse, J. Jackson, H. Jun, T. B. Brown, P. Dhariwal, S. Gray *et al.*, "Scaling laws for autoregressive generative modeling," *arXiv preprint arXiv:2010.14701*, 2020.
- [6] Z. Li, K. Li, S. Wang, S. Lan, Z. Yu, Y. Ji, Z. Li, Z. Zhu, J. Kautz, Z. Wu *et al.*, "Hydra-mdp: End-to-end multimodal planning with multi-target hydra-distillation," *arXiv preprint arXiv:2406.06978*, 2024.
- [7] X. Jia, J. You, Z. Zhang, and J. Yan, "Drivetransformer: Unified transformer for scalable end-to-end autonomous driving," *International Conference on Learning Representations*, 2025.
- [8] X. Jia, Z. Yang, Q. Li, Z. Zhang, and J. Yan, "Bench2drive: Towards multi-ability benchmarking of closed-loop end-to-end autonomous driving," *Advances in Neural Information Processing Systems*, 2024.
- [9] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Conference on robot learning*, 2017.
- [10] X. Weng, B. Ivanovic, Y. Wang, Y. Wang, and M. Pavone, "Paradrive: Parallelized architecture for real-time autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
- [11] 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.
- [12] 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.
- [13] 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*, vol. 45, no. 11, pp. 12 878–12 895, 2022.
- [14] W. Sun, X. Lin, Y. Shi, C. Zhang, H. Wu, and S. Zheng, "Sparsedriven: End-to-end autonomous driving via sparse scene representation," *2025 IEEE International Conference on Robotics and Automation*, 2025.
- [15] 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*, 2023.
- [16] X. Liu, Z. Zhong, Y. Guo, Y.-F. Liu, Z. Su, Q. Zhang, J. Wang, Y. Gao, Y. Zheng, Q. Lin *et al.*, "Reasonplan: Unified scene prediction and decision reasoning for closed-loop autonomous driving," *Conference on Robot Learning*, 2025.
- [17] P. Yang, Y. Zheng, Q. Zhang, K. Zhu, Z. Xing, Q. Lin, Y.-F. Liu, Z. Su, and D. Zhao, "Uncad: Towards safe end-to-end autonomous driving via online map uncertainty," *2025 IEEE International Conference on Robotics and Automation*, 2025.
- [18] Y. Zheng, P. Yang, Z. Xing, Q. Zhang, Y. Zheng, Y. Gao, P. Li, T. Zhang, Z. Xia, P. Jia *et al.*, "World4drive: End-to-end autonomous driving via intention-aware physical latent world model," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025.
- [19] P. Yang, B. Lu, Z. Xia, C. Han, Y. Gao, T. Zhang, K. Zhan, X. Lang, Y. Zheng, and Q. Zhang, "Worldrft: Latent world model planning with reinforcement fine-tuning for autonomous driving," *Proceedings of the AAAI conference on artificial intelligence*, 2026.
- [20] W. Peebles and S. Xie, "Scalable diffusion models with transformers," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023.
- [21] X. Zhai, A. Kolesnikov, N. Houlsby, and L. Beyer, "Scaling vision transformers," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022.
- [22] J.-B. Alayrac, J. Donahue, P. Luc, A. Miech, I. Barr, Y. Hasson, K. Lenc, A. Mensch, K. Millican, M. Reynolds *et al.*, "Flamingo: A visual language model for few-shot learning," *Advances in neural information processing systems*, 2022.
- [23] R. Dong, C. Han, Y. Peng, Z. Qi, Z. Ge, J. Yang, L. Zhao, J. Sun, H. Zhou, H. Wei *et al.*, "Dreamllm: Synergistic multimodal comprehension and creation," *arXiv preprint arXiv:2309.11499*, 2023.
- [24] J. Liu, H. Mao, Z. Chen, T. Zhao, N. Shah, and J. Tang, "Neural scaling laws on graphs," *CoRR*, 2024.
- [25] F. Lin, Y. Hu, P. Sheng, C. Wen, J. You, and Y. Gao, "Data scaling laws in imitation learning for robotic manipulation," *International Conference on Learning Representations*, 2025.
- [26] A. O'Neill, A. Rehman, A. Maddukuri, A. Gupta, A. Padalkar, A. Lee, A. Pooley, A. Gupta, A. Mandlkar, A. Jain *et al.*, "Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration 0," in *2024 IEEE International Conference on Robotics and Automation*, 2024.
- [27] H. Bharadhwaj, J. Vakil, M. Sharma, A. Gupta, S. Tulsiani, and V. Kumar, "Roboagent: Generalization and efficiency in robot manipulation via semantic augmentations and action chunking," in *2024 IEEE International Conference on Robotics and Automation*. IEEE, 2024.
- [28] H. R. Walke, K. Black, T. Z. Zhao, Q. Vuong, C. Zheng, P. Hansen-Estruch, A. W. He, V. Myers, M. J. Kim, M. Du *et al.*, "Bridgedata v2: A dataset for robot learning at scale," in *Conference on Robot Learning*, 2023.
- [29] A. Naumann, X. Gu, T. Dimlioglu, M. Bojarski, A. Degirmenci, A. Popov, D. Bislá, M. Pavone, U. Müller, and B. Ivanovic, "Data scaling laws for end-to-end autonomous driving," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025.
- [30] J. Philion and S. Fidler, "Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, 2020.
- [31] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018.
- [32] B. Liao, S. Chen, X. Wang, T. Cheng, Q. Zhang, W. Liu, and C. Huang, "Maptr: Structured modeling and learning for online vectorized hd map construction," in *International Conference on Learning Representations*, 2023.
- [33] A.-Q. Cao and R. De Charette, "Monoscene: Monocular 3d semantic scene completion," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [34] J. Ma, X. Chen, J. Huang, J. Xu, Z. Luo, J. Xu, W. Gu, R. Ai, and H. Wang, "Cam4docc: Benchmark for camera-only 4d occupancy forecasting in autonomous driving applications," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
- [35] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016.
- [36] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3d gaussian splatting for real-time radiance field rendering," *ACM Trans. Graph.*, vol. 42, no. 4, pp. 139–1, 2023.