

Generalizing Motion Planners with Mixture of Experts for Autonomous Driving

Qiao Sun^{1*}, Huimin Wang^{4*}, Jiahao Zhan², Fan Nie³, Xin Wen⁴, Leimeng Xu⁴,
Kun Zhan⁴, Peng Jia⁴, Xianpeng Lang⁴, Hang Zhao^{1,5}

Abstract—Large real-world driving datasets have sparked significant research into various aspects of data-driven motion planners for autonomous driving. These include data augmentation, model architecture, reward design, training strategies, and planner pipelines. These planners promise better generalizations on complicated and few-shot cases than previous methods. However, experiment results show that many of these approaches produce limited generalization abilities in planning performance due to overly complex designs or training paradigms. In this paper, we review and benchmark previous methods focusing on generalizations. The experimental results indicate that as models are appropriately scaled, many design elements become redundant. We introduce StateTransformer-2 (STR2), a scalable, decoder-only motion planner that uses a Vision Transformer (ViT) encoder and a mixture-of-experts (MoE) causal Transformer architecture. The MoE backbone addresses modality collapse and reward balancing by expert routing during training. Extensive experiments on the NuPlan dataset show that our method generalizes better than previous approaches across different test sets and closed-loop simulations. Furthermore, we assess its scalability on billions of real-world urban driving scenarios, demonstrating consistent accuracy improvements as both data and model size grow.

I. INTRODUCTION

Generalization bottlenecks the performance of autonomous driving motion planning for complex cases in the real world. These challenges arise from inconsistent objectives in complex environments. For example, human drivers might cross solid white lines while overtaking slow traffic ahead. Learning-based planners offer promising solutions by learning the complex mapping between observations and driving decisions from large datasets or simulations. A generalizable policy needs to react to similar scenarios and balance different objectives when scenarios get complicated. Considering the recent advances in the generalization of large language and vision models, scaling learning-based motion planners, including the training set and model sizes, could solve complicated, few-shot, and zero-shot driving problems. Additionally, MoE [1] Transformers architectures better learn and balance complex preferences.

Scaling for generalizations requires a large, diverse, and high-quality dataset for training and testing. We select the NuPlan [2] dataset for training, which offers a much larger and more comprehensive collection of diverse urban driving scenarios compared to previous datasets [3]–[5]. For testing, ensuring a sufficiently large test set is critical for

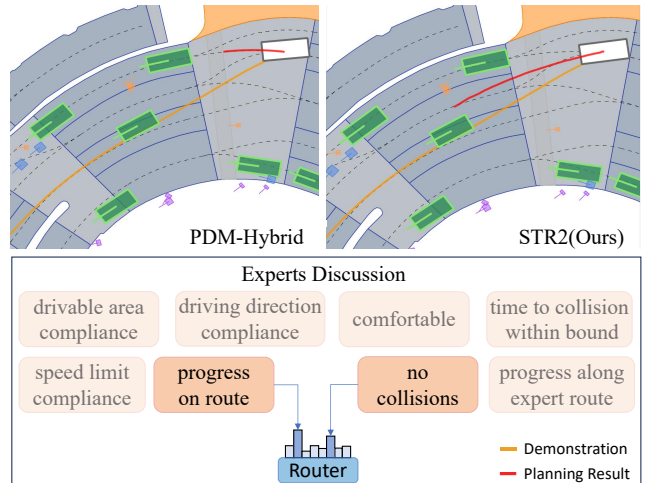


Fig. 1: The planning results, in red, from PDM-Hybrid and STR2 at the pickup area at the top, and an illustration of the MoE model learning and balancing different explicit rewards at the bottom. In this case, STR2 produces a better-nudging trajectory by balancing two conflicting rewards, making progress, and avoiding collisions.

testing generalities, especially for evaluating complicated cases. However, testing motion planners using closed-loop simulations is resource-intensive and time-consuming due to slow CPU model inferences. As a result, no open test server is available, which has led previous methods [6]–[8] to propose their own preferred subset to report results, complicating direct comparison. To address this issue, we accelerate the simulation process by GPU-parallel model inferences in batches for better efficiency with larger test sets. We extract two larger 4k test sets from the validation and the test dataset and provide a distinctively more comprehensive performance benchmark by thoroughly testing previous methods with their official checkpoints against various subsets. On the other hand, testing generality with complicated and few-shot cases presents significant challenges due to the difficulty of simulating realistic reactions during closed-loop simulations. We implement a wide spectrum of tests to benchmark the planning performance. First of all, closed-loop simulations challenge the models by drifting to different positions not visited in the training set. For example, models with low generalization abilities might not be able to correct themselves from the edge of the road because there are few or no samples at those dangerous positions in the training set for models to learn. Secondly, reactive simulations challenge

*Denotes equal contribution.

¹Shanghai Qi Zhi Institute ²Fudan University ³Stanford University
⁴LiAuto ⁵Tsinghua University
Corresponding at: hangzhao@mail.tsinghua.edu.cn

the models with the environment agents’ out-of-distribution reactions to potential conflicts. Finally, unseen new synthetic scenarios from InterPlan [9], like a crash site, challenge the zero-shot motion planning abilities of the models.

Our experiment results suggest that various previous methods greatly suffer from causality confusions, as mentioned in [7], producing over-smooth trajectories with modality collapse. Inspired by [10], we propose STR2 to model different driving rewards by different experts via a router at each layer by an MoE backbone, as illustrated in Fig. 1. On the other hand, inspired by multi-modality foundation models [11], we formulate the motion planning task as a general sequence modeling task and scale our decoder-only MoE backbone [1] up to 800 million parameters with a ViT [12] encoder and a two-layer MLP decoder. Unlike some previous methods, we do not add additional training paradigms like reinforcement learning (RL) [7], inverse reinforcement learning (IRL) [8] or contrastive learning [13]. In contrast, we train STR2 with a straightforward one-stage self-supervised learning without additional reward engineering for the testing metrics. We implement an autoregressive process for efficient training and flexible sampling. Inspired by [14], [15], we involve an additional proposal classification before generating key points and trajectories. The proposal classification further avoids modality collapse on the trajectory curvatures in regression on the continuous space during learning. As a result, comprehensive experiment results indicate a better performance on all metrics against all testing datasets than previous methods with a more general but more challenging raster representation of the environment. Additionally, our method produces much smaller performance drops when tested with reactive closed-loop simulations, as well as with unseen scenarios than other previous methods.

In summary, our contributions are:

- We propose a scalable MoE-based autoregressive model to learn different explicit rewards for motion planning and our method outperforms previous state-of-the-art methods by scaling with self-supervisions.
- We comprehensively benchmark and analyze the generality of previous methods across multiple test sets with multiple Closed-Loop metrics by speeding up the model inference during simulations.
- We unprecedentedly present scaling experiments with up to billions of diverse real-world urban driving scenarios.
- We release our codes for training, testing, and simulations on the NuPlan dataset for easy reproduction at <https://github.com/Tsinghua-MARS-Lab/StateTransformer>.

II. RELATED WORK

The non-data-driven planning systems are widely used in academia and industry [16] for their reliability in producing optimal solutions. However, as the testing environment gets more complicated, engineered cost functions are getting more difficult to cover the human-like balancing of simple objectives like comfort, safety, and efficiency. On the other

hand, reinforcement learning (RL) methods also guarantee an optimal solution with generalization abilities but suffer from significant sim-to-real gaps when deployed on large-scale AV fleets.

A. Imitation Learning based Planners

Imitation Learning (IL) methods learn driving policies from experienced drivers. These self-supervised learning models are easy to scale without additional data labeling and promise better performance as long as the size of the training data keeps growing. The challenge is that previous IL methods suffer from poor generalization abilities on motion planning tasks, often represented as covariate shift problems like accumulative error, due to complex environments and rewards. PDM-Hybrid [6] utilizes imitation loss to refine trajectories based on rule-based center lines. GameFormer [17] hierarchically models the relationships between scene elements by query-based cross-attention mechanisms. PlanTF [7] explores different augmentation and dropout strategies to mitigate compounding errors from a weak model. DTPP [18] employs a tree-structured policy planner and proposes a differentiable joint training framework for both ego-conditioned prediction and a cost model, which is trained by an inverse reinforcement learning (IRL) paradigm. PLUTO [13] provides additional hand-crafted reward functions, as heat maps which can be dated back to [19], [20]. Reward engineering is dangerous because it can be simply another way of overfitting the dataset. GUMP [21] uses Soft Actor-Critic (SAC) [22] with reinforcement learning (RL). Sampling and searching are also useful techniques widely used by IL planners [6], [8], [23], as well as self-supervised generative models on other tasks, like language modeling.

B. Scaling Laws

Previous studies [24], [25] provide adequate empirical results indicating a log-log relationship between the accuracy and size of the training dataset, as well as the accuracy and the model parameters on language modeling tasks. Considering these Transformer architectures are also good at modeling time series trajectories, it is intuitive to scale with them for better generalization abilities on the autonomous driving motion planning problem. Some early studies [21], [26] suggest similar scaling properties by training with models similar to the language models.

C. Simulation and Testings

NuPlan [2] is a large-scale dataset aiming to fulfill the learning-based planners’ needs of a huge training dataset. The dataset encompasses 1300 hours of recorded driving data collected from 4 urban centers, segmented into 75 scenario types using automated labeling tools. Although it is instinctive to benchmark different methods’ performance by trajectory accuracy, which is the offset between model outputs and the ground truth trajectory, some [6], [27] argue a significant gap between the open-loop prediction performance and the effectiveness in the real world. However, it is challenging to test the performance within a simulation.

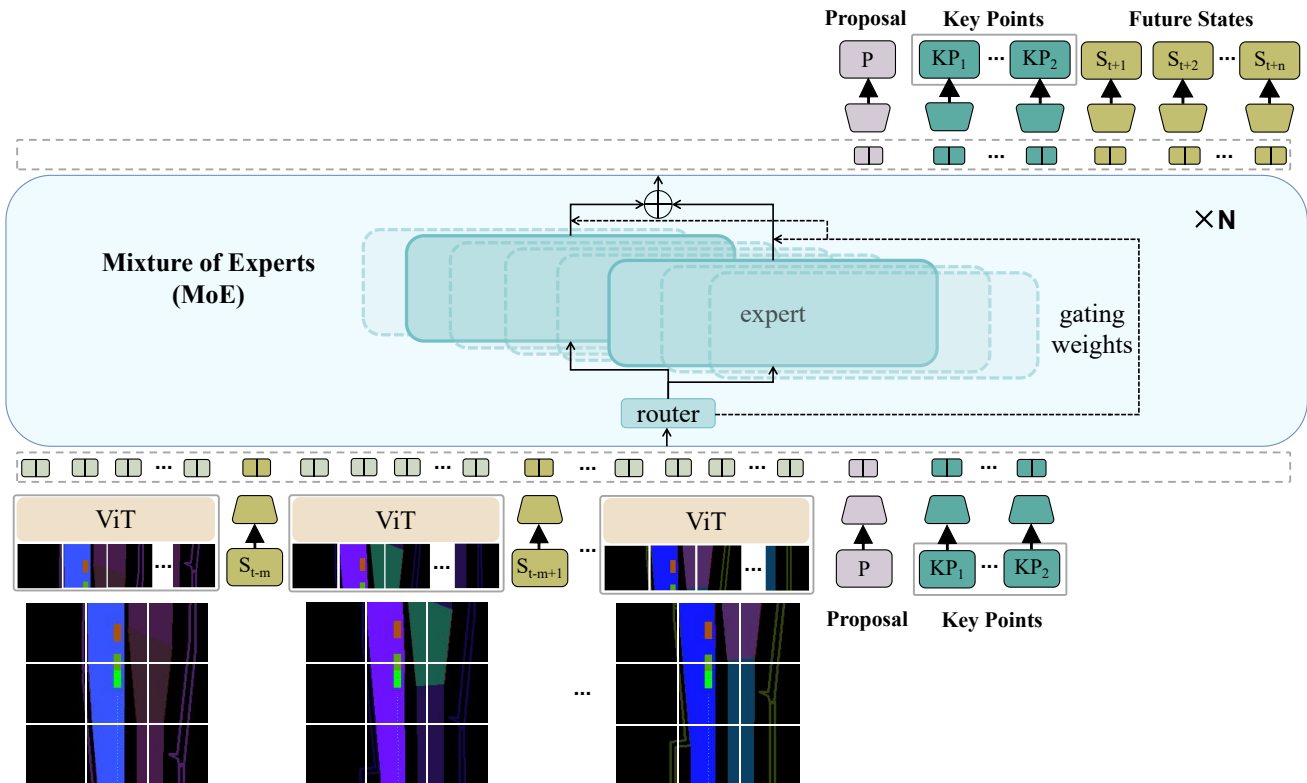


Fig. 2: An overview of the STR2-CPKS model which has a sequence of context, proposal, key points, and future states for the MoE backbone to model. For STR2-CKS, proposals are removed in the sequence for better efficiency. The context part has rasterized environment information encoded by scalable ViT encoders and past ego states.

Firstly, the controller might produce a tracking error [7] as a systematic bias. Furthermore, it is hard [9] to simulate the reaction of other road users responding to the models' output. Without a perfect solution, we cover as many tests and metrics as possible to comprehensively benchmark each method. Finally, rule-based planners and reward engineering cost maps can easily overfit the test set by tuning against failed cases in the test set.

III. SCALING FOR GENERALITY WITH MOE

Following previous learning-based motion planners, our model learns the mapping from observations of 2 past seconds to an optimal 8 seconds of future ego trajectory. To ensure scalability, we follow [26] to formulate this task as a sequence modeling problem to learn with an autoregressive model. In this section, we will discuss details about data preprocessing, rasterizations, model designs, training details, the dataset for training and testing, the metrics we choose to evaluate, and the LiAuto dataset for scaling.

A. Training and Testing Datasets

Learning-based motion planners promise better generalization performance than rule-based planners with enough data. Surprisingly, previous methods filter and only train with a subset, about 1 million scenarios, of all available training samples. Following [26], we remove scenario type filters leading to over 7 million NuPlan scenarios for training. For the test set, we sample 100 scenarios for each type

TABLE I: MoE Backbone Architecture

Parameter	MoE-100m	MoE-800m	MoE-1b
dim	320	512	1024
n_layers	16	32	16
hidden_dim	1280	2048	4096
n_heads	16	32	16
n_kv_heads	8	8	8
num_experts	8	8	8
top_k_experts	2	2	2

(14 different types in total provided for testing by the NuPlan devkit) leading to a total of 4790 scenarios for Val4k and 4639 scenarios for Test4k. We also test and compare all methods on the val14 (for better comparison) and testHard (testing generality only on hard cases) subsets. For the closed-loop simulations, reactive surrounding agents challenge the generality of learning-based methods more with out-of-distribution road users' trajectories in responses. Finally, we compare different methods with a more dynamic and complicated testset, InterPlan [9].

B. Data Preprocess and Rasterization

Preprocess and Caching. We make several optimizations to save disk usage: 1. we convert all data from float64 to float32; 2. we merge each crowd of pedestrians into one large shape; 3. we delete unused information like speed and accelerations of these road users. We sample scenarios with an interval of 1 second, producing 7 million scenarios for

TABLE II: Performance comparison with various state-of-the-arts with NuPlan Closed-Loop Metrics. NR is the overall score of the closed-loop simulation with non-reactive agents. R is the overall score of the closed-loop simulation with reactive agents controlled by a rule-based planner.

Methods	Val14 Set		TestHard Set		Val4k Set		Test4k Set		InterPlan Score \uparrow
	NR \uparrow	R \uparrow	NR \uparrow	R \uparrow	NR \uparrow	R \uparrow	NR \uparrow	R \uparrow	
Expert (log replay)	94	80	85.96	68.80	93.08	76.91	93.12	75.01	14.76
IDM [28]	75.59	77.32	56.16	62.26	72.52	76.03	72.45	74.10	-
GameFormer [29]	82.95	83.88	66.59	68.83	-	-	-	-	-
PDM-Hybrid [6]	92.84	92.12	65.07	75.18	90.31	90.62	90.65	90.13	41.61
PlanTF [7]	84.72	76.25	72.59	60.62	80.52	70.84	80.90	70.62	30.53
DTPP [8]	71.66	66.09	60.11	63.66	68.43	65.24	71.98	70.20	30.32
STR-16m [26]	45.06	49.69	27.59	36.13	42.27	45.16	39.66	39.62	-
STR2-CPKS-100m (Ours)	92.79	92.18	74.25	78.40	90.45	91.14	90.66	90.25	42.13
STR2-CKS-800m (Ours)	92.32	92.12	74.21	78.58	90.64	91.15	90.35	90.37	44.77
STR2-CPKS-800m (Ours)	93.91	92.51	77.54	82.02	91.41	91.53	92.14	91.38	46.03

TABLE III: Performance comparison on testHard Set with details of the closed-loop reactive simulations. Higher scores indicate better performance. Col. refers to No at-fault Collisions. Drivable refers to Drivable area compliance. Direction refers to Driving direction compliance. Making Prog. refers to Ego progress along the expert’s route ratio. TTC refers to Time to Collision (TTC) within bound. Speed Limit refers to Speed limit compliance. Prog. on Route refers to Ego progress along the expert’s route ratio.

Methods	Score	Col.	Drivable	Direction	Making Prog.	TTC	Speed Limit	Prog. on Route	Comfort
Expert (log replay)	68.80	77.02	95.96	98.16	100.00	69.85	94.12	98.48	99.26
GameFormer [29]	68.83	-	-	-	-	-	-	-	-
PlanTF [7]	60.62	90.07	94.85	97.98	80.51	85.66	97.97	65.22	92.28
PDM-Hybrid [6]	75.18	95.22	95.58	99.08	93.38	84.19	99.53	75.47	83.45
Hoplan [30]	75.06	89.33	94.85	96.13	97.05	80.51	95.28	85.02	98.52
GUMP hybrid [21]	77.77	94.36	98.98	98.95	94.41	87.46	97.51	77.08	79.84
STR2-CPKS-100m (Ours)	78.40	96.51	96.32	98.90	94.49	85.29	99.70	77.91	83.46
STR2-CKS-800m (Ours)	78.58	96.32	96.69	98.90	94.49	84.56	99.70	79.29	86.02
STR2-CPKS-800m (Ours)	82.02	97.98	96.69	99.08	94.12	87.87	99.27	78.86	95.59

training. Not like [6], we do not make route corrections on the training dataset. We filter total static scenarios, in which the ego vehicle keeps static for the whole 10 seconds (2 seconds in the past and 8 seconds in the future).

Rasterization. For rasterizations, we keep each type of road shape and each type of road user in a separate channel, represented as boolean values of occupancies. The total number of channels is 34. We rasterize the map and agents into two 224x224 images. One covers a smaller range for slow but delicate movements and one covers a larger range for fast movements. The fast OpenCV library proceeds the rasterization process during training and testing without the need for caching features ahead.

C. StateTransformer-2

In this section, we explain model designs of STR2 in detail. All encoded embeddings are formulated into one sequence for the decoder-only MoE backbone, shown in 2. Compared to STR [26], we eliminate the diffusion-based decoder for better efficiency, scalability, and overall performance.

ViT encoder. We employ a decoder-only ViT image encoder for better scalability and performance, which consists of a stacked 12 layers of Transformers. Rasterized images are sliced into 16 patches. We apply no attention dropouts on the ViT encoder. We select GeLU [31] as the activation function for the ViT encoder.

Mixture-of-Expert. Language modeling tasks also require models to learn and balance through complicated and often stochastically controversial rewards from expert data. Inspired by the generalization results of the MoE models on the language modeling tasks, we replace the GPT-2 backbone [33] with an MoE backbone for sequence modeling. The MoE backbone, inspired by [10], [34], is based on the Transformer architecture [35] with modification of changing the feed-forward blocks to Mixture-of-Expert layers. The MoE layers provide much better memory efficiency through specialized kernels and Expert Parallelism (EP). We also utilize the Flash Attention 2 [36] and the Data Parallelism (DP) for better training efficiency. The backbone architecture parameters for different sizes are summarized in Table I.

Autoregression and Proposal. We add a proposal embedding feature in the generation sequence for modality classifications with a Cross-Entropy loss, similar to the implementations in [15], [26] on the motion prediction task with the Waymo Open Motion Dataset (WOMD) [3]. Following [37], we extract 512 proposals with K-Means clustering with a mini-batch of 1000 from 0.7 million randomly selected dynamically feasible trajectories by their spatial-temporal distances. Each normalized trajectory includes 80 points (x, y, and yaw) for the future 8 seconds.

Sampling and Scoring: In NuPlan closed-loop simulations, the ego vehicle is controlled by a Linear-quadratic regulator (LQR) controller producing tracking errors to bring

TABLE IV: Performance comparison with various state-of-the-arts with NuPlan Open-Loop Metrics.

Methods	Val4k Set			Test4k Set			Val14 Set		
	OLS \uparrow	8sADE \downarrow	8sFDE \downarrow	OLS \uparrow	8sADE \downarrow	8sFDE \downarrow	OLS \uparrow	8sADE \downarrow	8sFDE \downarrow
PlanCNN [32]	-	-	-	-	-	-	64	2.468	5.936
PDM-Hybrid [6]	84.06	2.435	5.202	82.01	2.618	5.546	84	2.382	5.068
PlanTF [7]	88.59	1.774	3.892	87.30	1.855	4.042	89.18	1.697	3.714
DTPP [8]	65.15	4.196	9.231	64.18	4.117	9.181	67.33	4.088	8.846
STR-124m [26]	81.88	1.939	4.968	82.68	2.003	4.972	88.0	1.777	4.515
STR2-CKS-800m (Ours)	90.07	1.473	4.124	89.12	1.537	4.269	89.2	1.496	4.210

TABLE V: Scaling study on the impact of model parameters with STR2-CKS evaluated on NuPlan Open-Loop Metrics.

STR	Val14-OLS	Val14-8sADE	Val14-8sFDE	MR
STR2-100m	88.56	1.55	4.47	3.05%
STR2-800m	89.02	1.48	4.20	2.87%
STR2-1b	89.74	1.46	4.13	2.15%

a systematic bias from the model’s output and the final trajectory to compute the metrics. Inspired by [6], we apply similar sampling and scoring methods to enhance the model’s output. Specifically, we apply lateral offsets of [-1 m, 0 m, 1 m] to the model’s output trajectory with the route centerline, resulting in four lateral candidate paths. We set 5 different target speeds for diverse lateral sampling of [0.2, 0.4, 0.6, 0.8, 1.0] of the maximal speed limit. With the original model’s output, 21 trajectory candidates are sent to an LQR controller. The 21 tracking results are then scored as in [6].

D. Training and Testing Settings

We train both the ViT encoder and the MoE backbone from scratch. We train STR2 with a batch size of 16 on 8 H20 (96GB) Nvidia GPUs for 20 epochs for the STR2-CPKS-800m model. We train STR2 with a batch size of 64 on 8 3090 Nvidia GPUs for 20 epochs for the STR2-CPKS-100m model. We run Open-Loop and Closed-Loop simulations with a batch size of 50. The Val4k and Test4k simulations take about 5 hours to run closed-loop simulations on each. We use a Cosine-Restart learning rate scheduler, restarting a cosine LR schedule after each epoch, to avoid overfitting at a local minimum. All STR2 are trained with bfloat16 for training efficiency.

E. Scaling on the LiAuto Dataset

The LiAuto dataset is an industrial-level extra-large real-world driving dataset. This dataset includes a lane-level navigation map and tracking results from a sensor setting of 7 RGB cameras, 1 LiDAR, and 5 millimeter-wave radars (MMWR). The same dataset is used for training models which are later deployed on a fleet with a size of over 900,000 different models of vehicles. We select urban driving scenarios collected from the last 6 months without any human labeling. We filter the scenarios with the wrong navigation routes, not matched with actual future driving trajectories. Finally, we reformulate all driving logs into training and testing samples with a length of 10 seconds, 2 seconds in the past, and 8 seconds in the future. The final training dataset has over 1 billion training samples.

IV. RESULTS

In this section, we present, discuss, and analyze the experiment results of STR2 on the open NuPlan dataset and the large LiAuto dataset. Despite learning from more challenging raster representations, STR2 achieves comparable open-loop performance and beats other SOTA methods on the Closed-Loop metrics with better overall scores, suggesting better generalization in learning and balancing explicit driving rewards. Additionally, we present a comprehensive scaling analysis of training on the LiAuto dataset, implying strong scalability across extra-large datasets of up to 1 billion training samples.

A. Open-Loop Performance

The open-loop evaluation loss, or 8sADE can be considered a direct indicator of the learning abilities of different models. As shown in Table IV, STR2 outperforms the other methods on 8sADE accuracy and is comparable on other Open-loop metrics, suggesting strong fitting ability. Additionally, we also evaluate STR2 by training with different sizes, in response to the trainable model parameters. As shown in V, larger models tend to fit the dataset distribution better than smaller ones, indicating great scalability along model sizes.

B. Closed-Loop Performance

Closed-loop performance is the spotlight metric for benchmarking different motion planners. We test a wide spectrum of planners¹ with various test sets. To further challenge the generalization ability, we test these methods against the novel closed-loop benchmark, InterPlan [9]. InterPlan constructs out-of-distribution testing scenarios like dealing with a crush site in the middle of the road. As shown in Table II, STR2 outperforms the other methods on all 4 test sets, including the most popular Val14 set [6], the hard-cases TestHard set [7], and two larger scale Val4k and Test4k sets. Details of each metric are shown in Table III

Beyond numerical results, we uncover several insights through comparison. Firstly, in the reactive closed-loop simulations, other road users will always yield to the ego vehicle with each conflict. This means the overall score should always be higher on reactive simulations than on non-reactive simulations with the same test set due to fewer collisions. Surprisingly, we discover a performance drop of the PlanTF [7] and the DTPP [8] when tested with

¹No GameFormer [17] checkpoint available for re-evaluations.

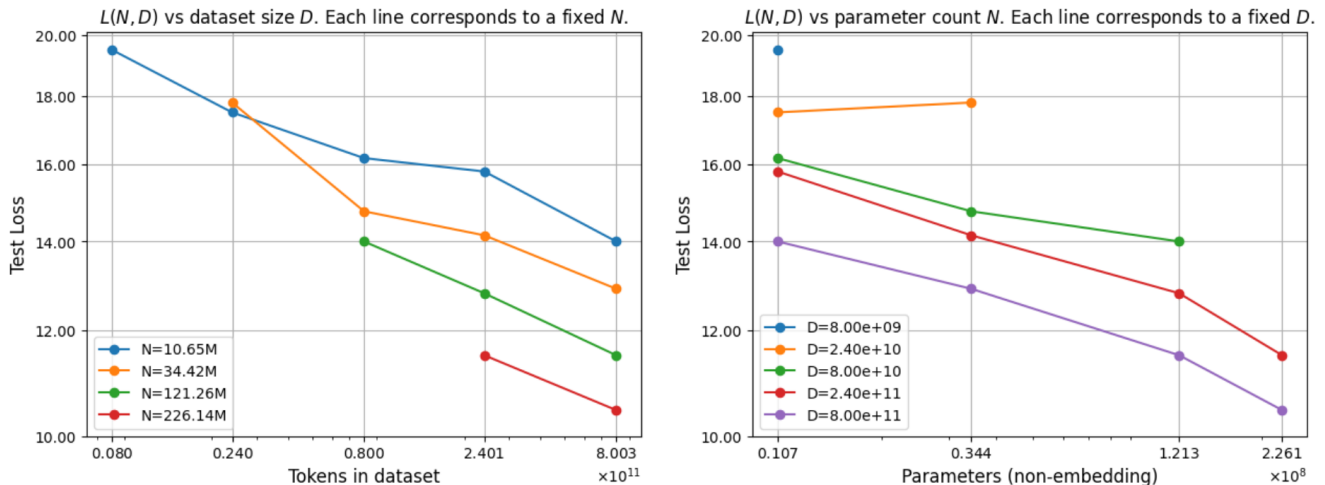


Fig. 3: Scaling results with the size of the training dataset, counted as the number of tokens D in the left and scaling results with model parameters N in the right. All axes are logarithmically scaled.

reactive surrounding agents. These drops indicate inferior generalizations of these planners than the others because they are more likely to be distracted by other road users’ out-of-distribution reactions due to a severe overfit to the original expert demonstrations, probably a result of causality confusion during training. Although they spot the problem of causality confusion on the ego past states and fix it by state dropouts, the same problem remains on the past trajectory of the other agents. This conclusion is double-confirmed in the InterPlan testings with similar performance drops. On the other hand, STR2 generalizes better than previous methods on reactive simulations. Additionally, we discover a significantly better generalization performance than other methods on the testHard and InterPlan test sets. The testHard set contains filtered in-distribution few-shot cases, like negotiating with other road users, and complicated cases, like balancing multiple conflicts. The InterPlan contains zero-shot cases, like driving through construction areas. STR2 suffers the least performance deterioration indicating the best generalization abilities by scaling the MoE backbone. Finally, STR2 significantly outperforms STR, the same model structure with a GPT2 backbone, with the closed-loop simulations.

C. Scaling to Billions on LiAuto Dataset

We study the scalability of our method by evaluating performance over the size of the dataset and model parameters. Results show that they follow a log-log relationship [38] with the evaluation/test loss, producing outstanding generalization performance up to billions of training samples. Following previous methods to test the scalability, we examine the converged test loss, which is the L2 loss for motion planning with STR2-CKS models. We test STR2 across 3 magnitudes on the size of the dataset (up to 1 billion) and over 2 magnitudes on the size of the models (up to 300M). As shown in Fig. 3, the left image illustrates the relationship between the test loss L and the dataset size D , denoted as $L(D)$, with varying model sizes N . The right image shows the relationship between the test loss L and the

number of parameters N , denoted as $L(N)$, with varying dataset sizes D . Specifically, we observe that $L(N)$ can be fitted with $L(N) = (N_c/N)^{\alpha_N}$ where N represents the number of trainable model parameters excluding encoding and positional embeddings and L denotes the L2 loss, and α_N is the power-law exponent characterizing the scaling behavior of the loss.

V. CONCLUSIONS

In this paper, we propose a scalable and strong model, STR2, for generalization on motion planning for autonomous driving. We demonstrated that previous methods often rely on overly complex designs that may hinder scalability when applied to large-scale datasets. Through extensive benchmarking on the NuPlan dataset and industrial-scale datasets, we showed that our decoder-only MoE architecture significantly improves generalization across diverse driving scenarios, including challenging out-of-distribution zero-shot cases. Our experiments confirm that MoE models when trained at scale, can achieve superior planning performance over more complex architectures and training paradigms by balancing multiple explicit rewards.

Limitation and future work. We leave comprehensive scaling analysis on the LiAuto dataset with larger models as future works due to large computation requirements. Testing the interaction-intensive scenarios with the NuPlan reactive simulations suffers from limited environment agents’ (vehicles only) behavior simulations. We leave more tests with more advanced simulators as future works. Another limitation of STR2 is the inference time. Larger models might generate slower. We leave speeding up the inference time on edge computing resources as future works.

ACKNOWLEDGMENT

We would like to thank Yujian Li for speeding up the NuPlan simulator, Qinghui Zhao for retesting previous methods, Tianyu Zhang for CPKS model training, Yue Wang for trajectory sampling, Weixing Cai, Yangang Zou, and Wentao Li for simulation visualization and analysis.

REFERENCES

- [1] A. Q. Jiang, A. Sablayrolles, A. Roux, A. Mensch, B. Savary, C. Bamford, D. S. Chaplot, D. d. l. Casas, E. B. Hanna, F. Bressand, *et al.*, “Mixtral of experts,” *arXiv preprint arXiv:2401.04088*, 2024.
- [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] S. Ettinger, S. Cheng, B. Caine, C. Liu, H. Zhao, S. Pradhan, Y. Chai, B. Sapp, C. R. Qi, Y. Zhou, *et al.*, “Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9710–9719.
- [4] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, “nuScenes: A multimodal dataset for autonomous driving,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 621–11 631.
- [5] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. K. Pontes, D. Ramanan, P. Carr, and J. Hays, “Argoverse 2: Next generation datasets for self-driving perception and forecasting,” in *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS Datasets and Benchmarks 2021)*, 2021.
- [6] D. Dauner, M. Hallgarten, A. Geiger, and K. Chitta, “Parting with misconceptions about learning-based vehicle motion planning,” in *Conference on Robot Learning*. PMLR, 2023, pp. 1268–1281.
- [7] J. Cheng, Y. Chen, X. Mei, B. Yang, B. Li, and M. Liu, “Rethinking imitation-based planners for autonomous driving,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 14 123–14 130.
- [8] Z. Huang, P. Karkus, B. Ivanovic, Y. Chen, M. Pavone, and C. Lv, “Dtp: Differentiable joint conditional prediction and cost evaluation for tree policy planning in autonomous driving,” *arXiv preprint arXiv:2310.05885*, 2023.
- [9] M. Hallgarten, J. Zapata, M. Stoll, K. Renz, and A. Zell, “Can vehicle motion planning generalize to realistic long-tail scenarios?” *arXiv preprint arXiv:2404.07569*, 2024.
- [10] W. Fedus, B. Zoph, and N. Shazeer, “Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity,” *Journal of Machine Learning Research*, vol. 23, no. 120, pp. 1–39, 2022.
- [11] A. Singh, R. Hu, V. Goswami, G. Couairon, W. Galuba, M. Rohrbach, and D. Kiela, “Flava: A foundational language and vision alignment model,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 15 638–15 650.
- [12] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, “An image is worth 16x16 words: Transformers for image recognition at scale,” *ICLR*, 2021.
- [13] J. Cheng, Y. Chen, and Q. Chen, “Pluto: Pushing the limit of imitation learning-based planning for autonomous driving,” *arXiv preprint arXiv:2404.14327*, 2024.
- [14] H. Zhao, J. Gao, T. Lan, C. Sun, B. Sapp, B. Varadarajan, Y. Shen, Y. Shen, Y. Chai, C. Schmid, *et al.*, “Tnt: Target-driven trajectory prediction,” in *Conference on Robot Learning*. PMLR, 2021, pp. 895–904.
- [15] S. Shi, L. Jiang, D. Dai, and B. Schiele, “Motion transformer with global intention localization and local movement refinement,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 6531–6543, 2022.
- [16] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli, “A survey of motion planning and control techniques for self-driving urban vehicles,” *IEEE Transactions on intelligent vehicles*, vol. 1, no. 1, pp. 33–55, 2016.
- [17] Z. Huang, H. Liu, and C. Lv, “Gameformer: Game-theoretic modeling and learning of transformer-based interactive prediction and planning for autonomous driving,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 3903–3913.
- [18] Z. Huang, P. Karkus, B. Ivanovic, Y. Chen, M. Pavone, and C. Lv, “Dtp: Differentiable joint conditional prediction and cost evaluation for tree policy planning in autonomous driving,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 6806–6812.
- [19] M. Bansal, A. Krizhevsky, and A. Ogale, “Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst,” *arXiv preprint arXiv:1812.03079*, 2018.
- [20] J. Zhou, R. Wang, X. Liu, Y. Jiang, S. Jiang, J. Tao, J. Miao, and S. Song, “Exploring imitation learning for autonomous driving with feedback synthesizer and differentiable rasterization,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 1450–1457.
- [21] Y. Hu, S. Chai, Z. Yang, J. Qian, K. Li, W. Shao, H. Zhang, W. Xu, and Q. Liu, “Solving motion planning tasks with a scalable generative model,” *arXiv preprint arXiv:2407.02797*, 2024.
- [22] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” in *International conference on machine learning*. PMLR, 2018, pp. 1861–1870.
- [23] A. Cui, S. Casas, A. Sadat, R. Liao, and R. Urtasun, “Lookout: Diverse multi-future prediction and planning for self-driving,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 16 107–16 116.
- [24] 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,” *CoRR*, vol. abs/2001.08361, 2020. [Online]. Available: <https://arxiv.org/abs/2001.08361>
- [25] J. Hoffmann, S. Borgeaud, A. Mensch, E. Buchatskaya, T. Cai, E. Rutherford, D. de Las Casas, L. A. Hendricks, J. Welbl, A. Clark, T. Hennigan, E. Noland, K. Millican, G. van den Driessche, B. Damoc, A. Guy, S. Osindero, K. Simonyan, E. Elsen, J. W. Rae, O. Vinyals, and L. Sifre, “Training compute-optimal large language models,” 2022. [Online]. Available: <https://arxiv.org/abs/2203.15556>
- [26] Q. Sun, S. Zhang, D. Ma, J. Shi, D. Li, S. Luo, Y. Wang, N. Xu, G. Cao, and H. Zhao, “Large trajectory models are scalable motion predictors and planners,” *arXiv preprint arXiv:2310.19620*, 2023.
- [27] Z. Li, Z. Yu, S. Lan, J. Li, J. Kautz, T. Lu, and J. M. Alvarez, “Is ego status all you need for open-loop end-to-end autonomous driving?” 2024. [Online]. Available: <https://arxiv.org/abs/2312.03031>
- [28] M. Treiber, A. Hennecke, and D. Helbing, “Congested traffic states in empirical observations and microscopic simulations,” *Physical review E*, vol. 62, no. 2, p. 1805, 2000.
- [29] Z. Huang, H. Liu, and C. Lv, “Gameformer: Game-theoretic modeling and learning of transformer-based interactive prediction and planning for autonomous driving,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023, pp. 3903–3913.
- [30] Y. Hu, K. Li, P. Liang, J. Qian, Z. Yang, H. Zhang, W. Shao, Z. Ding, W. Xu, and Q. Liu, “Imitation with spatial-temporal heatmap: 2nd place solution for nuPlan challenge,” *arXiv preprint arXiv:2306.15700*, 2023.
- [31] D. Hendrycks and K. Gimpel, “Gaussian error linear units (gelus),” *arXiv preprint arXiv:1606.08415*, 2016.
- [32] K. Renz, K. Chitta, O.-B. Mercea, A. S. Koepke, Z. Akata, and A. Geiger, “Plant: Explainable planning transformers via object-level representations,” in *6th Annual Conference on Robot Learning*, 2022.
- [33] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, *et al.*, “Language models are unsupervised multitask learners,” *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [34] W. Fedus, J. Dean, and B. Zoph, “A review of sparse expert models in deep learning,” *arXiv preprint arXiv:2209.01667*, 2022.
- [35] A. Vaswani, “Attention is all you need,” *Advances in Neural Information Processing Systems*, 2017.
- [36] T. Dao, “Flashattention-2: Faster attention with better parallelism and work partitioning,” *arXiv preprint arXiv:2307.08691*, 2023.
- [37] 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.
- [38] 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.