

# Generalized Trajectory Scoring for End-to-end Multimodal Planning

Zhenxin Li<sup>1,2</sup> Wenhao Yao<sup>2</sup> Zi Wang<sup>1</sup> Xinglong Sun<sup>1</sup> Joshua Chen<sup>1</sup>  
Nadine Chang<sup>1</sup> Maying Shen<sup>1</sup> Zuxuan Wu<sup>2</sup> Shiyi Lan<sup>1</sup> Jose M. Alvarez<sup>1</sup>  
<sup>1</sup>NVIDIA <sup>2</sup>Fudan University

## Abstract

*End-to-end multi-modal planning is a promising paradigm in autonomous driving, enabling decision-making with diverse trajectory candidates. A key component is a robust trajectory scorer capable of selecting the optimal trajectory from these candidates. While recent trajectory scorers focus on scoring either large sets of static trajectories or small sets of dynamically generated ones, both approaches face significant limitations in generalization. Static vocabularies provide effective coarse discretization but struggle to make fine-grained adaptation, while dynamic proposals offer detailed precision but fail to capture broader trajectory distributions. To overcome these challenges, we propose GTRS (Generalized Trajectory Scoring), a unified framework for end-to-end multi-modal planning that combines coarse and fine-grained trajectory evaluation. GTRS consists of three complementary innovations: (1) a diffusion-based trajectory generator that produces diverse fine-grained proposals; (2) a vocabulary generalization technique that trains a scorer on super-dense trajectory sets with dropout regularization, enabling its robust inference on smaller subsets; and (3) a sensor augmentation strategy that enhances out-of-domain generalization while incorporating refinement training for critical trajectory discrimination. As the winning solution of the Navsim v2 Challenge, GTRS demonstrates superior performance even with sub-optimal sensor inputs, approaching privileged methods that rely on ground-truth perception. Code will be available at <https://github.com/NVlabs/GTRS>.*

## 1. Introduction

End-to-end multi-modal planning has emerged as a powerful approach in autonomous driving. Unlike traditional unimodal planners that predict a single trajectory [5, 11, 12], multi-modal approaches [1, 3, 6, 15, 16, 18, 19] generate multiple candidates, enabling greater adaptability during inference. This adaptability supports a wide range of applications, including responding to language instructions [21–23, 28], accommodating different driving styles [15, 25], and navigating complex driving environments [16, 18, 24].

The typical problem of end-to-end multi-modal planning involves evaluating multiple trajectory proposals through scoring [1, 3, 16, 18, 19, 24] given raw sensor data, without access to ground-truth perception. The planner selects the trajectory with the highest likelihood as the decision.

Current trajectory scoring methods generally fall into two categories: (1) scoring a large static vocabulary [3, 16, 18, 24], and (2) scoring a small set of dynamically generated proposals [19, 30]. Both approaches face challenges in generalization. Fixed trajectory vocabularies [3, 16, 18] offer limited flexibility, as they cannot adapt to situations where dynamic proposals are needed. Meanwhile, methods that rely on a small number of dynamic proposals [19, 30] often fail to generalize to unseen trajectories, since the scorer is only exposed to a narrow subset during training. Ideally, a robust scorer should generalize across diverse trajectory distributions—whether static or dynamic—to handle the full complexity of real-world scenarios.

To address these limitations, we propose GTRS (Generalized Trajectory Scoring) for end-to-end multi-modal planning. GTRS is built on a key insight: an effective trajectory scorer must be trained on both coarse and fine-grained trajectory distributions to develop robust generalization capabilities. Our approach contains three complementary techniques, each leading to a dedicated sub-network as shown in Fig. 1:

1. **Diffusion-based Trajectory Generation (DP):** A diffusion policy (DP) [4] produces diverse trajectory candidates with BEV features as the condition. DP provides the fine-grained details crucial for safety-critical situations that coarse, fixed vocabularies cannot capture. (Sec. 2.1)
2. **Trajectory Vocabulary Generalization (GTRS-Dense):** We train on a super-dense vocabulary of trajectory samples (16,384 trajectories) covering a wide range of driving scenarios. To maximize generalization of fixed trajectory vocabularies, we propose a trajectory dropout training strategy. The idea is to deliberately create a mismatch between training and inference vocabularies—training the model to effectively generalize to unseen trajectory distributions during inference. (Sec. 2.2)
3. **Sensor Augmentation with Refinement (GTRS-Aug):** To handle unexpected trajectory events and data distri-

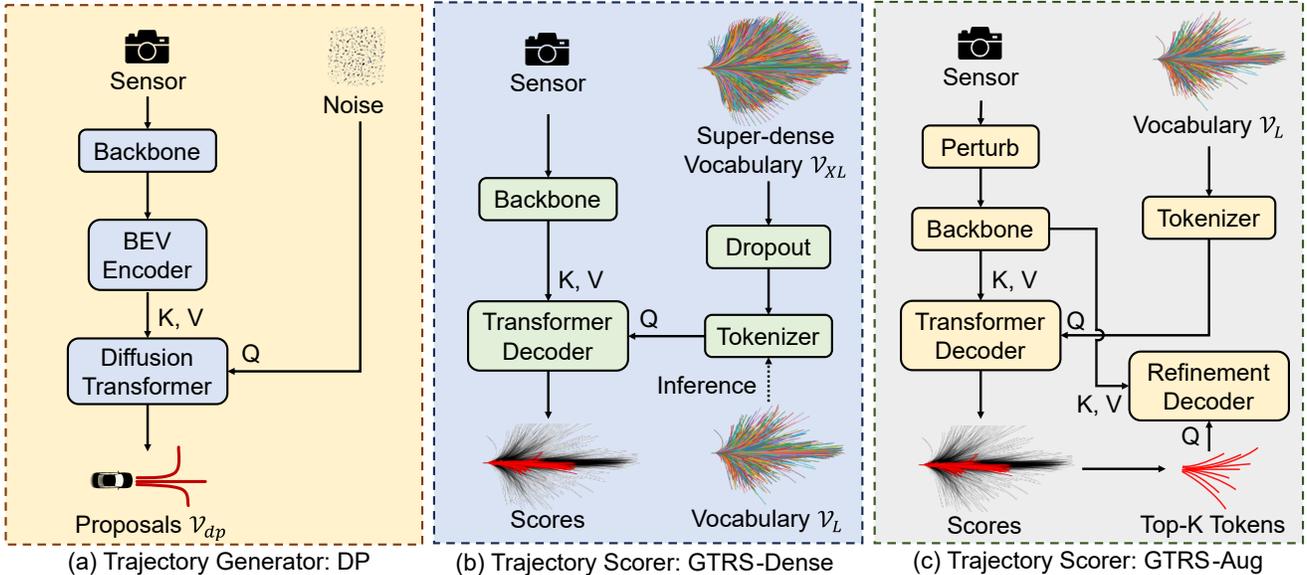


Figure 1. **The Three Pillars of GTRS.**

bution shifts in the form of viewpoint changes, we introduce a data augmentation strategy by applying rotation perturbations to sensor inputs, dramatically improving robustness to out-of-domain environments. Further, a refinement training mechanism enables the model to distinguish between subtly different trajectory options. (Sec. 2.3)

GTRS demonstrates strong trajectory scoring abilities even under sub-optimal sensor conditions, as evaluated on the Navhard benchmark. Our contributions are as follows:

1. We propose GTRS, a generalizable end-to-end multi-modal planning framework that combines diffusion-based trajectory generation with vocabulary scoring. With super-dense vocabularies, sensor augmentations, and refinement strategies, GTRS enables effective scoring across both dynamic and static candidates.
2. GTRS demonstrates superior planning performance and generalization to out-of-domain data on the Navhard benchmark [2]. With model ensembling, our sensor-based GTRS—the winning entry of the Navsim v2 Challenge—approaches the performance of the state-of-the-art planner PDM-Closed [6], which operates on ground-truth perception and is unaffected by degraded sensor inputs.

## 2. The Three Pillars of GTRS

### 2.1. Trajectory Generator

To obtain high-quality dynamic trajectory proposals during inference, diffusion models have become a popular practice in autonomous driving for its ability to generate multi-modal trajectory candidates [4, 13, 15, 19, 29–31, 33]. We adopt a Diffusion Policy-based trajectory generator following [4] to produce multiple trajectory proposals. This sub-network is composed of an image backbone to extract image fea-

tures, a BEV encoder where BEV queries attend to image features through a Transformer Encoder [26], and a Diffusion Transformer that generates  $N$  trajectory proposals  $\mathcal{V}_{dp}$  conditioned on the BEV features.

During training, we follow Transfuser [5] and include a BEV segmentation head to provide supervision on the BEV features. For the Diffusion Transformer, we apply first-order differentiation to the ground-truth trajectory waypoints to normalize the input and use the DDPM scheduling strategy [10] for denoising ground-truth trajectories.

### 2.2. Scorer with Vocabulary Generalization

While diffusion-based trajectory generators provide fine-grained trajectory proposals, they remain limited in their ability to capture the full breadth of possible driving scenarios. To achieve robust generalization, we introduce a novel vocabulary generalization technique that trains the model to effectively evaluate diverse trajectory distributions, even those not seen during training.

We propose the Generalized Vocabulary Scorer **GTRS-Dense**, which builds upon Hydra-MDP [18] but introduces critical innovations in trajectory evaluation. The architecture consists of an image backbone, a trajectory tokenizer that encodes candidates into feature representations, and a Transformer Decoder [26] that models complex interactions between trajectory and image tokens.

Our key innovation is twofold: First, we deliberately train on a super-dense trajectory vocabulary ( $\mathcal{V}_{XL}$  with 16,384 distinct trajectories) that comprehensively covers the trajectory space, while inferencing on a smaller vocabulary ( $\mathcal{V}_L$  with 8,192 trajectories)—forcing the model to develop generalizable representations. Second, we apply vocabulary dropout to  $\mathcal{V}_{XL}$  during training, randomly removing half of

| Method               | Img. Resolution | Backbone  | Training Vocab.              | Inference Vocab.                         | EPDMS <sub>1</sub> | EPDMS <sub>2</sub> | EPDMS       |
|----------------------|-----------------|-----------|------------------------------|--|--------------------|--------------------|-------------|
| DP w/o Scorer        | 512 × 2048      | V2-99     | -                            | $\mathcal{V}_{dp}$ (Random)              | 60.9               | 40.8               | 25.6        |
| GTRS-Dense           | 512 × 2048      | EVA-ViT-L | $\mathcal{V}_{XL}$           | $\mathcal{V}_{dp}$                       | 76.6               | 48.6               | 36.7        |
|                      | 512 × 2048      | EVA-ViT-L | $\mathcal{V}_{XL}$           | $\mathcal{V}_{XL}$                       | <b>78.1</b>        | 50.2               | 39.7        |
|                      | 512 × 2048      | EVA-ViT-L | $\mathcal{V}_{XL}$           | $\mathcal{V}_{dp} \cup \mathcal{V}_{XL}$ | 77.4               | 52.7               | 40.8        |
|                      | 512 × 2048      | EVA-ViT-L | $\mathcal{V}_{XL}$           | $\mathcal{V}_{dp} \cup \mathcal{V}_L$    | 71.8               | 57.3               | 42.0        |
|                      | 512 × 2048      | EVA-ViT-L | $\mathcal{V}_{XL}$ (Dropout) | $\mathcal{V}_{dp} \cup \mathcal{V}_L$    | 73.1               | <b>59.0</b>        | <b>43.4</b> |
| Baseline Scorer [18] | 256 × 1024      | ViT-L     | $\mathcal{V}_L$              | $\mathcal{V}_{dp} \cup \mathcal{V}_L$    | 73.4               | 54.1               | 40.6        |
| GTRS-Aug             | 256 × 1024      | ViT-L     | $\mathcal{V}_L$              | $\mathcal{V}_{dp} \cup \mathcal{V}_L$    | 75.0               | 56.9               | <b>43.4</b> |

Table 1. **Roadmap to Generalized Trajectory Scoring.** Random denotes that we randomly select a trajectory from  $\mathcal{V}_{dp}$  during inference. V2-99 [14] is pretrained from DD3D [20], EVA-ViT-L [9] is initialized from StreamPETR [27], and ViT-L [8] is from Depth Anything [32].

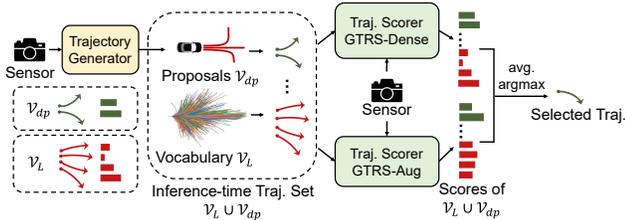


Figure 2. **The Inference Pipeline of GTRS.**

the trajectories in each batch. This serves multiple purposes: (1) it aligns the number of trajectory tokens during training and inference, (2) it creates intentional distribution shifts that improve robustness, and (3) it acts as an effective regularizer against overfitting to specific trajectory patterns.

This vocabulary generalization technique enables our model to effectively score both static trajectory vocabularies and dynamically generated proposals without requiring dedicated training on both types—a capability previous approaches have struggled to achieve.

### 2.3. Sensor-augmented Scorer with Refinement

To further enhance model robustness across diverse and out-of-domain environments, we develop a systematic sensor augmentation strategy, along with a refinement training mechanism. This approach focuses on two critical challenges: handling perceptual distribution shifts in sensor data and distinguishing between subtly different trajectory options in safety-critical scenarios.

First, we introduce structured sensor perturbations by applying controlled 2D horizontal view rotations to the input images. Rather than random augmentation, these perturbations specifically target the model’s ability to maintain consistent trajectory evaluation under varying viewing conditions. To maintain label consistency, we apply corresponding transformations to the ground-truths used for training.

Second, we develop a refinement training mechanism focused on fine-grained trajectory discrimination. As a training-only module, it incorporates an additional Transformer Decoder that progressively refines trajectory scores for the top-k most promising candidates, enabling the model to capture subtle differences between similar trajectories. The

refinement process is guided by a self-distillation framework where an exponential moving average (EMA) copy of the model provides soft supervision signals:

$$\tilde{y}_i^m = \hat{y}_i^m + \text{clip}(s_{i,\text{teacher}}^m - y_i^m, -\delta_m, \delta_m), \quad (1)$$

where  $\tilde{y}_i^m$  represents the refined target score,  $\hat{y}_i^m$  is the ground-truth score, and  $s_{i,\text{teacher}}^m$  is the teacher model’s prediction. The clipping parameter  $\delta_m$  ensures the refined targets remain within a reasonable range of the ground truth.

Together, these strategies enable **GTRS-Aug** to perform robustly in challenging out-of-domain settings without domain-specific adaptation.

## 3. Inference-time Integration

After training the sub-networks described above, we combine the trajectory generator and one of the trajectory scorers (*i.e.* GTRS-Dense, GTRS-Aug) during inference, as illustrated in Fig. 2. The dynamic proposals generated by the generator  $\mathcal{V}_{dp}$  are appended to the inference vocabulary  $\mathcal{V}_L$ , and the combined set  $\mathcal{V}_{dp} \cup \mathcal{V}_L$  is tokenized and scored by the scorer.

This sequential integration of dynamic proposals at inference time, rather than during training, is a deliberate design choice that leverages the strengths of both approaches. By training solely on a diverse static vocabulary, the scorer develops robust generalization abilities across a wide range of trajectory patterns. Then, at inference time, the diffusion-based generator provides fine-grained, context-aware trajectories specifically tailored to the current scene. Meanwhile, it avoids the computational overhead and potential instability of integrating diffusion sampling into the training loop, while still benefiting from the precision of dynamically generated trajectories during deployment.

## 4. Experiments

### 4.1. Dataset and metrics

The Navsim dataset [7] is designed for evaluating end-to-end driving systems, addressing prior limitations in benchmarking [17]. The Navsim v2 Challenge introduces a new split, Navhard, which features difficult real-world scenarios alongside their synthetic continuations generated using 3DGS.

| Method         | Img. Resolution   | Backbone  | Stage   | NC          | DAC         | DDC         | TLC         | EP          | TTC         | LK          | HC          | EC          | EPDMS       |
|----------------|-------------------|-----------|---------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| PDM-Closed [6] | GT Perception     | -         | Stage 1 | 94.4        | 98.8        | 100         | 99.5        | 100         | 93.5        | 99.3        | 87.7        | 36.0        | 51.3        |
|                |                   |           | Stage 2 | 88.1        | 90.6        | 96.3        | 98.5        | 100         | 83.1        | 73.7        | 91.5        | 25.4        |             |
| LTF [5]        | $256 \times 1024$ | ResNet34  | Stage 1 | 96.2        | 79.5        | 99.1        | 99.5        | <b>84.1</b> | 95.1        | 94.2        | 97.5        | <b>79.1</b> | 23.1        |
|                |                   |           | Stage 2 | 77.7        | 70.2        | 84.2        | 98.0        | <b>85.1</b> | 75.6        | 45.4        | 95.7        | <b>75.9</b> |             |
| GTRS-Dense     | $512 \times 2048$ | V2-99     | Stage 1 | 98.7        | 95.8        | 99.4        | 99.3        | 72.8        | 98.7        | 95.1        | 96.9        | 40.4        | 41.7        |
|                |                   |           | Stage 2 | 91.4        | 89.2        | 94.4        | 98.8        | 69.5        | 90.1        | 54.6        | 94.1        | 49.7        |             |
|                |                   | EVA-ViT-L | Stage 1 | 97.6        | 95.8        | 99.7        | <b>99.8</b> | 77.2        | 97.8        | 95.3        | 97.3        | 46.7        | 43.4        |
|                |                   |           | Stage 2 | 91.9        | 91.3        | 92.7        | 99.0        | 72.7        | 90.4        | 53.8        | 94.1        | 41.6        |             |
|                |                   | ViT-L     | Stage 1 | <b>98.9</b> | 98.2        | <b>99.8</b> | 99.6        | 73.9        | 98.9        | 95.3        | 97.3        | 40.0        | 45.3        |
|                |                   |           | Stage 2 | 91.5        | 90.8        | <b>94.7</b> | 98.5        | 70.8        | 90.1        | 55.4        | <b>97.2</b> | 54.2        |             |
| GTRS-Aug       | $512 \times 2048$ | V2-99     | Stage 1 | <b>98.9</b> | 95.1        | 99.2        | 99.6        | 76.1        | 99.1        | 94.7        | <b>97.6</b> | 54.2        | 42.1        |
|                |                   |           | Stage 2 | 87.9        | 88.8        | 89.6        | 98.8        | 80.3        | 86.0        | 53.5        | 97.1        | 56.1        |             |
|                |                   | EVA-ViT-L | Stage 1 | 98.7        | 98.0        | 99.1        | <b>99.8</b> | 75.9        | 98.7        | 94.7        | <b>97.6</b> | 49.8        | 44.7        |
|                |                   |           | Stage 2 | 89.5        | 89.6        | 92.9        | 98.5        | 78.9        | 86.4        | 55.3        | 96.5        | 52.7        |             |
|                | ViT-L             | Stage 1   | 97.8    | 97.3        | 98.9        | 99.3        | 77.1        | 98.2        | 95.8        | <b>97.6</b> | 50.2        | 43.4        |             |
|                |                   | Stage 2   | 90.3    | 88.9        | 90.8        | 98.9        | 81.1        | 87.4        | 54.2        | 95.1        | 48.3        |             |             |
| GTRS-E-Lite    | $512 \times 2048$ | EVA-ViT-L | Stage 1 | 98.2        | 98.9        | 99.6        | 99.6        | 75.8        | <b>98.4</b> | <b>96.9</b> | 97.3        | 53.3        | 46.6        |
|                |                   |           | Stage 2 | 90.6        | 91.7        | 93.6        | 98.5        | 77.1        | 89.0        | <b>56.1</b> | 96.3        | 47.8        |             |
| GTRS-E         | *                 | *         | Stage 1 | <b>98.9</b> | <b>99.3</b> | <b>99.8</b> | <b>99.8</b> | 75.2        | <b>98.4</b> | 96.0        | <b>97.6</b> | 51.6        | <b>49.4</b> |
|                |                   |           | Stage 2 | <b>92.3</b> | <b>93.3</b> | 94.6        | <b>99.2</b> | 73.1        | <b>91.2</b> | 53.9        | 96.7        | 56.8        |             |

Table 2. **Performance on the Navhard Benchmark.** Backbone settings follow Tab. 1. GTRS-E-Lite ensembles GTRS-Dense and GTRS-Aug with EVA-ViT-L. The challenge-winning entry **GTRS-E** ensembles all six models from GTRS-Dense and GTRS-Aug.

Nevertheless, we observe that the synthetic data exhibits sub-optimal quality, which often contains artifacts such as distortion and blurring and may impair the performance of sensor-based planners. The Navsim v2 challenge evaluates end-to-end models based on the extended PDM Score (EPDMS) [16], an extension of the PDM Score [7] by aggregating multiple rule-based metrics<sup>1</sup>. Finally, it uses a two-stage scoring pipeline to aggregate the metrics on real-world data (Stage 1) and synthetic data (Stage 2).

## 4.2. Implementation Details

We train all models on the Navtrain split with 24 NVIDIA A100 GPUs, while the Navhard split and other synthetic sensor data are not used for training. Training is conducted for 20 epochs with a total batch size of 528 by default, while the training lasts 50 epochs for the trajectory generator. The learning rate and weight decay are  $2 \times 10^{-4}$  and 0.0. We concatenate the frontal view with center-cropped front-left and front-right views to form the input image. For the trajectory generator DP, we formulate a similar input image from the back-view, back-right view, and back-left view for BEV construction. Finally, we use 100 denoising steps with the DDPM scheduler [10] and generate 100 proposals in  $\mathcal{V}_{dp}$ .

## 4.3. Roadmap to Generalized Trajectory Scoring

**Trajectory Vocabulary Generalization.** We evaluate GTRS-Dense with various inference vocabularies (Tab. 4.3). Notably, though it is trained only on the super-dense static vocabulary  $\mathcal{V}_{XL}$ , the scorer generalizes well to unseen dynamic proposals in  $\mathcal{V}_{dp}$  (EPDMS: 36.7), demonstrating strong zero-

shot generalization by outperforming the generator with random selection substantially (+ 11.1 EPDMS). When combining  $\mathcal{V}_{dp}$  with  $\mathcal{V}_{XL}$ , performance improves by +1.1 EPDMS over  $\mathcal{V}_{XL}$ , confirming the complementary benefit of dynamic proposals at inference. Interestingly,  $\mathcal{V}_{dp} \cup \mathcal{V}_L$  outperforms  $\mathcal{V}_{dp} \cup \mathcal{V}_{XL}$ , likely because the reduced vocabulary complexity leads to better generalization in out-of-domain synthetic data. Finally, applying dropout to  $\mathcal{V}_{XL}$  during training yields the best performance (EPDMS: 43.4), showing that vocabulary dropout enhances generalization significantly.

**Sensor Augmentation with Refinement.** Finally, we evaluate GTRS-Aug, which incorporates sensor augmentation and refinement training. This model achieves the same top-level performance (EPDMS: 43.4) as the best GTRS-Dense variant, surpassing the baseline scorer [18] by a large margin (+2.8 EPDMS). This confirms that augmentations and refinement training greatly improve trajectory scoring.

## 4.4. Main Results

As shown in Tab. 2, our GTRS variants achieve significant improvements over the LTF baseline [5]. By scaling the image backbone to ViT-L [8] and EVA-ViT-L [9], our best single model achieves 45.3 EPDMS on the Navhard Benchmark. Further, GTRS-E-Lite, an ensemble of GTRS-Dense with GTRS-Aug during scoring, achieves 46.6 EPDMS. Our challenge-winning entry GTRS-E, an ensemble of all six variants, reaches 49.4 EPDMS, approaching the performance of PDM-Closed [6]—a privileged planner that relies on ground-truth perception—despite ours using challenging synthetic sensor input. This demonstrates the exceptional generalization capabilities of our approach across both trajectory distributions and perceptual domains.

<sup>1</sup><https://github.com/autonomousvision/navsim/blob/main/docs/metrics.md>

## References

- [1] Sourav Biswas, Sergio Casas, Quinlan Sykora, Ben Agro, Abbas Sadat, and Raquel Urtasun. Quad: Query-based interpretable neural motion planning for autonomous driving. In *IEEE International Conference on Robotics and Automation*, pages 14236–14243. IEEE, 2024. 1
- [2] 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. 2
- [3] Shaoyu Chen, Bo Jiang, Hao Gao, Bencheng Liao, Qing Xu, Qian Zhang, Chang Huang, Wenyu Liu, and Xinggang Wang. Vad2: End-to-end vectorized autonomous driving via probabilistic planning. *arXiv preprint arXiv:2402.13243*, 2024. 1
- [4] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, page 02783649241273668, 2023. 1, 2
- [5] Kashyap Chitta, Aditya Prakash, Bernhard Jaeger, Zehao Yu, Katrin Renz, and Andreas Geiger. Transfuser: Imitation with transformer-based sensor fusion for autonomous driving. *IEEE TPAMI*, 2022. 1, 2, 4
- [6] Daniel Dauner, Marcel Hallgarten, Andreas Geiger, and Kashyap Chitta. Parting with misconceptions about learning-based vehicle motion planning. In *Conference on Robot Learning*, pages 1268–1281. PMLR, 2023. 1, 2, 4
- [7] Daniel Dauner, Marcel Hallgarten, Tianyu Li, Xinshuo Weng, Zhiyu Huang, Zetong Yang, Hongyang Li, Igor Gilitschenski, Boris Ivanovic, Marco Pavone, Andreas Geiger, and Kashyap Chitta. Navsim: Data-driven non-reactive autonomous vehicle simulation and benchmarking. In *NeurIPS*, 2024. 3, 4
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3, 4
- [9] Yuxin Fang, Quan Sun, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. Eva-02: A visual representation for neon genesis. *arXiv preprint arXiv:2303.11331*, 2023. 3, 4
- [10] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 33:6840–6851, 2020. 2, 4
- [11] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhui Wang, et al. Planning-oriented autonomous driving. In *CVPR*, pages 17853–17862, 2023. 1
- [12] 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 *ICCV*, pages 8340–8350, 2023. 1
- [13] Chiyu Jiang, Andre Cornman, Cheolho Park, Benjamin Sapp, Yin Zhou, Dragomir Anguelov, et al. Motiondiffuser: Controllable multi-agent motion prediction using diffusion. In *CVPR*, pages 9644–9653, 2023. 2
- [14] Youngwan Lee, Joong-won Hwang, Sangrok Lee, Yuseok Bae, and Jongyoul Park. An energy and gpu-computation efficient backbone network for real-time object detection. In *CVPRW*, pages 0–0, 2019. 3
- [15] 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 preprint arXiv:2503.10434*, 2025. 1, 2
- [16] Kailin Li, Zhenxin Li, Shiyi Lan, Yuan Xie, Zhizhong Zhang, Jiayi Liu, Zuxuan Wu, Zhiding Yu, and Jose M Alvarez. Hydra-mdp++: Advancing end-to-end driving via expert-guided hydra-distillation. *arXiv preprint arXiv:2503.12820*, 2025. 1, 4
- [17] Zhiqi Li, Zhiding Yu, Shiyi Lan, Jiahao Li, Jan Kautz, Tong Lu, and Jose M Alvarez. Is ego status all you need for open-loop end-to-end autonomous driving? *arXiv preprint arXiv:2312.03031*, 2023. 3
- [18] 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. 1, 2, 3, 4
- [19] 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 *CVPR*, 2025. 1, 2
- [20] Dennis Park, Rares Ambrus, Vitor Guizilini, Jie Li, and Adrien Gaidon. Is pseudo-lidar needed for monocular 3d object detection? In *ICCV*, pages 3142–3152, 2021. 3
- [21] Katrin Renz, Long Chen, Ana-Maria Marcu, Jan Hünermann, Benoit Hanotte, Alice Karnsund, Jamie Shotton, Elahe Arani, and Oleg Sinavski. Carllava: Vision language models for camera-only closed-loop driving. *arXiv preprint arXiv:2406.10165*, 2024. 1
- [22] Katrin Renz, Long Chen, Elahe Arani, and Oleg Sinavski. Simlingo: Vision-only closed-loop autonomous driving with language-action alignment. *arXiv preprint arXiv:2503.09594*, 2025.
- [23] Chonghao Sima, Katrin Renz, Kashyap Chitta, Li Chen, Hanxue Zhang, Chengen Xie, Jens Beißwenger, Ping Luo, Andreas Geiger, and Hongyang Li. Drivelm: Driving with graph visual question answering. In *ECCV*, pages 256–274. Springer, 2024. 1
- [24] Chonghao Sima, Kashyap Chitta, Zhiding Yu, Shiyi Lan, Ping Luo, Andreas Geiger, Hongyang Li, and Jose M Alvarez. Centaur: Robust end-to-end autonomous driving with test-time training. *arXiv preprint arXiv:2503.11650*, 2025. 1
- [25] 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. *arXiv preprint arXiv:2503.08612*, 2025. 1
- [26] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017. 2

- [27] Shihao Wang, Yingfei Liu, Tiancai Wang, Ying Li, and Xiangyu Zhang. Exploring object-centric temporal modeling for efficient multi-view 3d object detection. In *ICCV*, pages 3621–3631, 2023. [3](#)
- [28] Shihao Wang, Zhiding Yu, Xiaohui Jiang, Shiyi Lan, Min Shi, Nadine Chang, Jan Kautz, Ying Li, and Jose M Alvarez. Omnidrive: A holistic llm-agent framework for autonomous driving with 3d perception, reasoning and planning. *arXiv preprint arXiv:2405.01533*, 2024. [1](#)
- [29] Theodor Westny, Björn Olofsson, and Erik Frisk. Diffusion-based environment-aware trajectory prediction. *arXiv preprint arXiv:2403.11643*, 2024. [2](#)
- [30] Zebin Xing, Xingyu Zhang, Yang Hu, Bo Jiang, Tong He, Qian Zhang, Xiaoxiao Long, and Wei Yin. Goalflow: Goal-driven flow matching for multimodal trajectories generation in end-to-end autonomous driving. In *CVPR*, 2025. [1](#)
- [31] Brian Yang, Huangyuan Su, Nikolaos Gkanatsios, Tsung-Wei Ke, Ayush Jain, Jeff Schneider, and Katerina Fragkiadaki. Diffusion-es: Gradient-free planning with diffusion for autonomous and instruction-guided driving. In *CVPR*, pages 15342–15353, 2024. [2](#)
- [32] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jishi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. *arXiv preprint arXiv:2401.10891*, 2024. [3](#)
- [33] 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](#)