

# Foresight in Motion: Reinforcing Trajectory Prediction with Reward Heuristics

Muleilan Pei<sup>1</sup> Shaoshuai Shi<sup>2\*</sup> Xuesong Chen<sup>2</sup> Xu Liu<sup>3</sup> Shaojie Shen<sup>1</sup>  
<sup>1</sup>HKUST <sup>2</sup>Voyager Research, Didi Chuxing <sup>3</sup>Zhuoyu Technology  
 {mpei, eeshaojie}@ust.hk, shaoshuaics@gmail.com

## Abstract

Motion forecasting for on-road traffic agents presents both a significant challenge and a critical necessity for ensuring safety in autonomous driving systems. In contrast to most existing data-driven approaches that directly predict future trajectories, we rethink this task from a planning perspective, advocating a “First Reasoning, Then Forecasting” strategy that explicitly incorporates behavior intentions as spatial guidance for trajectory prediction. To achieve this, we introduce an interpretable, reward-driven intention reasoner grounded in a novel query-centric Inverse Reinforcement Learning (IRL) scheme. Our method first encodes traffic agents and scene elements into a unified vectorized representation, then aggregates contextual features through a query-centric paradigm. This enables the derivation of a reward distribution, a compact yet informative representation of the target agent’s behavior within the given scene context via IRL. Guided by this reward heuristic, we perform policy rollouts to reason about multiple plausible intentions, providing valuable priors for subsequent trajectory generation. Finally, we develop a hierarchical DETR-like decoder integrated with bidirectional selective state space models to produce accurate future trajectories along with their associated probabilities. Extensive experiments on the large-scale Argoverse and nuScenes motion forecasting datasets demonstrate that our approach significantly enhances trajectory prediction confidence, achieving highly competitive performance relative to state-of-the-art methods.

## 1. Introduction

Trajectory prediction is a crucial component of advanced autonomous driving systems, bridging upstream perception with downstream planning modules. Accurately forecasting the movements of surrounding traffic agents requires reasoning about unknown intentions, given the inherent uncertainty and multimodal nature of driving behavior [7, 34].

Most existing data-driven motion prediction models [14,

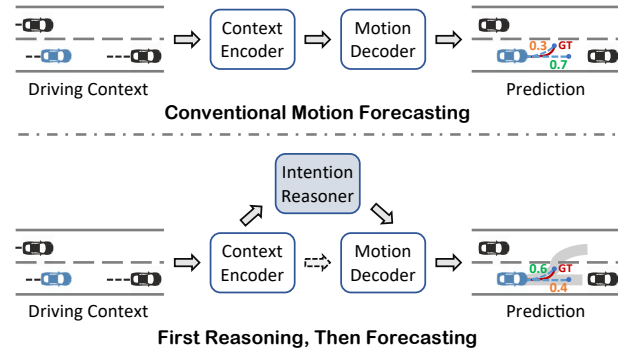


Figure 1. Comparison between the conventional motion forecasting pipeline and ours. Existing approaches mainly predict the trajectory as depicted in the upper framework, while our method follows a “First Reasoning, Then Forecasting” strategy, illustrated in the lower framework (with the dashed flow indicating an optional step). The blue vehicle represents the target agent to be predicted.

[19, 32, 46] utilize imitative approaches that either directly regress trajectories or classify endpoints based on data distribution matching from training datasets. However, these methods often lack adequate consideration of driving behavior, limiting interpretability and reliability. While numerous methods achieve strong performance on benchmark metrics, few explicitly reason about future intentions, creating a critical bottleneck in generating interpretable and robust multimodal predictions in real-world applications.

In contrast, human drivers typically maneuver vehicles in a hierarchical manner, making high-level decisions (e.g., lane changes or overtaking) before executing specific motion strategies. Consider the prediction module of the ego vehicle can be viewed as planning on behalf of other agents, under the assumption that road users behave rationally. Despite the intrinsic connection between trajectory prediction and planning, few studies have explored insights from the planning domain [33]. Motivated by these observations, we pose a key question: *Can trajectory prediction be approached from a planning perspective and enhanced with intention reasoning capabilities?*

To this end, we advocate a “First Reasoning, Then Forecasting” strategy, where behavioral intention reasoning pro-

\*Corresponding author.

vides critical prior guidance to facilitate accurate and confident multimodal motion forecasts. Consider an overtaking scenario as a motivating example: a model capable of explicitly reasoning both overtaking and lane-keeping intents in advance can generate more reliable predictions compared to direct forecasting without reasoning, as shown in Fig. 1. Besides, incorporating longer-term intention reasoning can further enhance prediction confidence (see Tab. 2).

Nevertheless, due to the inherent complexity of driving scenarios, reasoning about future intentions remains challenging when relying solely on handcrafted rules or predefined planners [34]. One promising alternative involves utilizing Large Reasoning Models (LRMs), such as OpenAI-o1 [17], to enable intention reasoning within trajectory predictors. However, their substantial computational demands render them impractical for onboard driving systems. Fortunately, recent advances in LRMs have demonstrated the remarkable reasoning capabilities of Reinforcement Learning (RL) techniques in domains such as mathematics and coding [15], prompting an intriguing question: *Can RL-based paradigms be leveraged to reason about an agent’s future intentions in trajectory prediction?*

In this regard, we explore the feasibility of applying RL paradigms to model the behavior reasoning of agents within autonomous driving scenarios. We formulate the task as a Markov Decision Process (MDP) and define the target agent’s intentions accordingly. To balance performance and computational efficiency, we construct a grid-level graph to represent the scene layout, where the intention is defined as a sequence of decisions over a discrete grid world, analogous to a *plan* in the traditional context of RL [6, 35]. The intention sequence is referred to as the Grid-based Reasoning Traversal (GRT) in our paper. However, a fundamental challenge of adopting RL for trajectory prediction lies in modeling the reward, as agent intentions remain unknown.

To overcome this challenge, we propose a reward-driven intention reasoner grounded in Maximum Entropy Inverse Reinforcement Learning (MaxEnt IRL) [48]. This framework first learns the agent-specific reward distribution from demonstrations and associated driving contexts via IRL. The learned reward serves as a compact representation [29], capturing observed behaviors and underlying intentions of the agent. Utilizing these inferred rewards as heuristics, we then perform policy rollouts to sample multiple plausible GRTs and extract their corresponding intention-informed features, providing prior guidance for trajectory prediction and thereby improving prediction accuracy and confidence.

Additionally, to further enhance feature extraction from scene contexts, we introduce a novel Query-centric IRL framework, termed QIRL, which integrates IRL with a query-based encoding mechanism. QIRL efficiently and flexibly aggregates vectorized scene context features into spatial grid-like tokens, facilitating structured reasoning.

With this dense grid representation, we augment our model with an auxiliary Occupancy Grid Map (OGM) prediction head, which enables dense predictions of future spatial-temporal occupancies for each agent in the scene. This auxiliary task effectively enhances the feature fusion process by capturing future interactions between agents, improving the overall prediction performance (see Tab. 7).

Finally, to fully exploit features provided by the intention reasoner, we develop a hierarchical DETR-like trajectory decoder. An anchor-free trajectory token initially generates proposals conditioned on GRT-derived features, which then act as initialized anchors for final trajectory decoding. Given the inherent sequential nature of trajectory states and recent advancements in selective state-space models (Mamba) for long-term, structured dynamic modeling [13, 47], we incorporate a bidirectional variant, Bi-Mamba, to effectively capture the sequential dependencies of trajectory states. This enhancement significantly improves both prediction accuracy and confidence (see Tab. 6).

In summary, our main contributions are as follows:

- (1) We introduce a “First Reasoning, Then Forecasting” strategy, rethinking trajectory prediction tasks from a planning perspective.
- (2) We propose a novel reward-driven intention reasoner for motion forecasting, with a QIRL module that integrates the MaxEnt IRL paradigm and vectorized context representation in a query-centric framework.
- (3) We develop a hierarchical DETR-like decoder with bidirectional selective state-space models (Bi-Mamba) to improve prediction accuracy and confidence.
- (4) Our approach significantly boosts the prediction confidence and achieves highly competitive performance on the Argoverse and nuScenes motion forecasting benchmarks, outperforming other state-of-the-art models.

## 2. Related Work

**Trajectory Prediction** for autonomous driving has been a subject of investigation for decades. Early efforts in this domain primarily relied on handcrafted rule-based or physics-based methods, which struggled to handle complex scenarios and lacked the capability for long-term predictions [20]. Recent approaches have transitioned to learning-based frameworks that leverage deep neural networks to encode agents’ historical motion profiles while integrating the topological and semantic information [26] of High-Definition (HD) maps. These maps are represented in either rasterized [3, 5] or vectorized [11, 18] formats. Rasterized representations typically use Bird’s-Eye-View (BEV) images as input, while vectorized representations rely on agent and map polylines as input. Convolutional Neural Networks (CNNs) and Graph Neural Networks (GNNs) are widely employed as feature extractors for these formats, playing a crucial role in encoding scene context. More recently, Transformer-

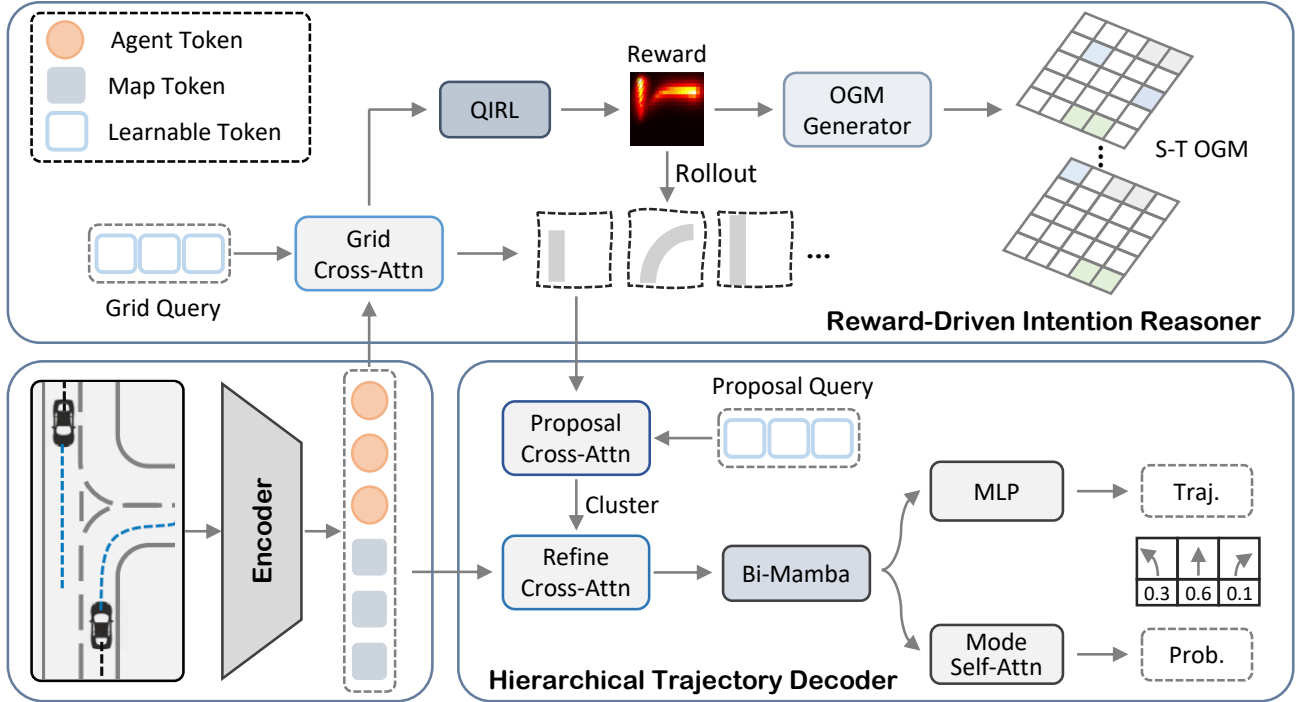


Figure 2. Overview of the proposed framework. Our approach integrates a reward-driven intention reasoner into an encoder-decoder motion forecasting structure. Scene context features are initially extracted using simple encoders and subsequently aggregated into grid queries. The proposed QIRL module generates reward distributions that reason about multiple plausible intention sequences as priors. These sampled tokens are then utilized to generate trajectory proposals, which, after clustering and refinement, yield multimodal future trajectories with associated confidences through a hierarchical DETR-like trajectory decoder enhanced by Bi-Mamba. Additionally, an auxiliary Spatial-Temporal OGM prediction head is incorporated to further improve feature fusion.

based architectures have garnered significant attention for feature extraction and fusion, owing to their ability to improve overall prediction performance [23, 32, 46]. Following this trend, our work adopts vectorized representations and leverages a query-centric Transformer encoder-decoder structure for feature aggregation and trajectory generation.

Despite these advancements, challenges persist in making trajectory prediction robust to out-of-distribution scenarios and generalizable to unseen environments. Our work addresses these gaps by rethinking trajectory prediction from a planning perspective, introducing a reward-driven intention reasoner to provide behavior guidance and contextually rich priors for advancing trajectory prediction.

**Reward** is a foundational concept in both planning and RL, serving as a guiding signal that shapes an agent’s behavior and decision-making process [24]. In planning, rewards are often designed to align with high-level objectives, such as reaching a goal while avoiding obstacles [25]. Typically, reward functions are either handcrafted or shaped through hierarchical frameworks, where high-level planners provide strategic guidance to lower-level controllers. Research on reward shaping has demonstrated that modifying the reward structure to emphasize specific behaviors or milestones can accelerate learning and improve policy robustness. In RL,

the reward function plays a central role, defining the agent’s objective and guiding it to perform actions that maximize cumulative rewards over time. The design of reward functions is crucial in both planning and RL; however, crafting effective rewards for complex tasks, such as autonomous driving, is exceptionally challenging. To address this challenge, Inverse RL (IRL) [22, 31] was introduced. IRL focuses on inferring a reward function from observed expert demonstrations, making it particularly valuable in scenarios where directly defining a reward function is infeasible. For example, MaxEnt IRL [48] has been widely applied to learn reward functions that capture the underlying intentions of expert behaviors [39], enabling agents to replicate nuanced, human-like decision-making in planning tasks [16].

Despite their utility, existing efficient IRL algorithms are typically tailored for structured and grid-like environments [10], limiting their flexibility in more complex domains. To overcome this limitation, we propose a novel query-centric framework that enhances the applicability and flexibility of MaxEnt IRL in our reward-driven intention reasoner. By leveraging this paradigm, our approach provides valuable reward heuristics that effectively reason about intentions of future behavior, providing informative priors for addressing the complexities inherent in motion forecasting tasks.

### 3. Methodology

#### 3.1. Problem Formulation

The goal of the standard trajectory prediction task is to forecast the future positions of a target agent over a time horizon  $T_f$ , given the driving context. We adopt vectorized representations for the scene input, including the historical observed states  $A \in \mathbb{R}^{N_a \times T_h \times C_a}$ , where  $N_a$  denotes the number of agents in the scene,  $T_h$  represents the number of past timestamps, and  $C_a$  captures motion profiles such as positions, velocities, headings, etc, along with HD map information  $M \in \mathbb{R}^{N_m \times N_s \times C_m}$ , where  $N_m$  and  $N_s$  correspond to the number of lane centerlines and lane segments, respectively, and  $C_m$  indicates the associated lane attributes.

Our approach employs a target-centric coordinate system in which all input elements are normalized to the current state of the target agent through translation and rotation operations. Given the inherent uncertainty of motion intentions, the predictor is tasked with providing  $K$  future trajectories  $Y \in \mathbb{R}^{K \times T_f \times 2}$ , along with the corresponding probabilities  $p \in \mathbb{R}^{K \times 1}$ .

#### 3.2. Framework Overview

As shown in Fig. 2, our motion forecasting approach adopts an encoder-decoder structure comprising a query-centric scene context encoder, a Mamba-enhanced hierarchical trajectory decoder, and a reward-driven intention reasoner. First, we represent the driving context in a vectorized format and utilize the agent and map encoder to extract scene features. These fused features are then aggregated into spatial grid tokens through cross-attention mechanisms. Subsequently, in the QIRL module, a grid-based MaxEnt IRL algorithm is leveraged to infer the reward distribution, thereby reasoning about multiple plausible intention sequences (i.e., GRTs) over the 2-D grid map through policy rollouts. Additionally, a dense prediction head for the Spatial-Temporal Occupancy Grid Map (S-T OGM) is incorporated to model future interactions among agents. Finally, we introduce a hierarchical DETR-like [2] trajectory decoder that generates trajectory proposals, which are further clustered and refined to produce multimodal future trajectories enhanced by the Bi-Mamba architecture.

#### 3.3. Query-Centric Context Encoding

Given the vectorized agent representations  $A$  and map representations  $M$ , we first tokenize them into separate feature sets. Specifically, we use an agent encoder, a simple 1-D CNN model [18], to obtain the agent features  $F_a \in \mathbb{R}^{N_a \times C}$ . For the map encoder, we adopt a PointNet-like network [11, 28] to extract static map features as  $F_m \in \mathbb{R}^{N_m \times C}$ . These resulting agent and map features are then concatenated to form context tokens  $F_c \in \mathbb{R}^{(N_a+N_m) \times C}$ , followed by a self-attention block to enhance feature fusion.

Since the reasoning process relies on a grid-level graph representation, we introduce learnable grid-shaped queries  $Q_G \in \mathbb{R}^{H \times W \times C}$  to integrate the scene features, where  $H$  and  $W$  define the spatial dimensions in the BEV plane. Each query  $Q_G^{s_i}$  at grid position  $s_i = (x_i, y_i)$  corresponds to a specific region in the real world with a resolution  $d$ . We then use the flattened grid queries with 2-D spatial learnable relative positional embeddings to aggregate the context tokens via a cross-attention mechanism.

#### 3.4. Reward-Driven Intention Reasoning

Given the grid tokens updated with context features, we first generate the reward distribution through our QIRL framework, which adapts the traditional grid-based MaxEnt IRL algorithm [48] in a query-centric paradigm. The MaxEnt IRL is typically defined as a finite MDP model, encompassing a state space, an action space, and a transition model. Its objective is to recover the reward distribution of the environment to yield a policy that mimics expert demonstrations by maximizing the log-likelihood of the demonstration data while adhering to the maximum entropy principle. Demonstrations consist of discrete state sequences, with rewards generally formulated as a combination of environmental features. The learning process involves an inner-loop forward RL process within each reward iteration until the loss  $\mathcal{L}_{IRL}$  converges. More technical details on the MaxEnt IRL can be found in Appx. 6.1 and [27, 40].

**QIRL.** In our QIRL framework, each grid  $s_i$  acts as a state, with its corresponding query  $Q_G^{s_i} \in \mathbb{R}^{1 \times C}$  representing contextual features. We aggregate features from grid tokens with a stack of  $1 \times 1$  CNN layers to establish a nonlinear mapping from driving context to the reward  $R \in \mathbb{R}^{H \times W \times 1}$ . Future trajectories are quantized to a resolution  $d$  to form expert demonstration states, which can also include paths for capturing long-term information, if available. Subsequently, the MaxEnt IRL algorithm is applied to derive the converged reward distribution alongside an optimal policy.

Rollouts are then performed based on the policy induced by the reward heuristic. We execute  $L$  rollouts over the grid map in parallel, producing multiple plausible GRTs as sequences of intentions,  $\Upsilon \in \mathbb{R}^{L \times \mathcal{H} \times 2}$ , where  $\mathcal{H}$  denotes the planning horizon. To better capture the multimodal future distribution, we set  $L \gg K$ . The grid tokens are then extracted according to the sampled GRT: for each position  $(x_i, y_i)$  associated with grid cell state  $s_i$  in the sampled GRT, the corresponding grid token  $Q_G^{s_i}$  is sequentially selected over  $\mathcal{H}$  steps. These grid tokens form the reasoning tokens  $Q_G^\Upsilon \in \mathbb{R}^{L \times \mathcal{H} \times C}$ . The GRT positions  $\Upsilon$  and associated reasoning tokens  $Q_G^\Upsilon$  serve as informative behavior intention priors for guiding subsequent motion forecasting.

**Auxiliary S-T OGM Prediction Head.** Leveraging the grid-shaped dense representations, we introduce an auxiliary S-T OGM prediction head to model future interactions

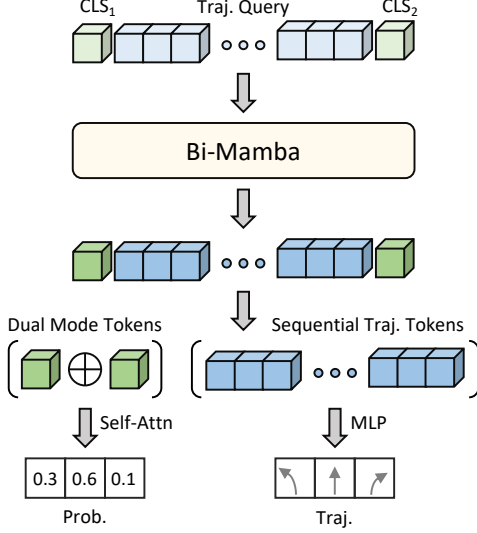


Figure 3. The process of Bi-Mamba-enhanced decoding.

among agents, thereby enhancing the fusion and aggregation of scene context features. We represent the occupancy map in binary form, where grid cells occupied in the BEV at future timestamps are set to 1, and unoccupied cells are set to 0. Our OGM generator takes the fused grid tokens  $Q_G$  and reward  $R$  as input and produces OGMs over  $T_f$  future timestamps using a U-Net-like architecture [30].

### 3.5. Mamba-Enhanced Trajectory Decoding

Given  $L$  plausible reasoning priors, we first use a DETR-like trajectory generator to produce  $L$  trajectories as proposals. We encode the GRT positions  $\Upsilon$  and reasoning tokens  $Q_G^r$  separately through simple MLP blocks, then concatenate and process them via an MLP-based feature fusion network to form the final reasoning tokens  $Q_r \in \mathbb{R}^{L \times H \times C}$ . Next, we introduce an anchor-free learnable trajectory proposal query  $Q_P \in \mathbb{R}^{L \times T_f \times C}$  to cross-attend to the prior features  $Q_r$  from the intention reasoner. This proposal query is then decoded into  $L$  trajectory proposals by a regression head with an MLP block. We apply the K-means algorithm to cluster these proposals into  $K$  multimodal trajectory proposals  $\bar{Y} \in \mathbb{R}^{K \times T_f \times 2}$ . Subsequently, we use an anchor-based trajectory refinement, as in many existing motion predictors [44], to further enhance trajectory query prediction performance. Each trajectory proposal, acting as an explicit anchor prior, is re-encoded into the trajectory query  $Q_T \in \mathbb{R}^{K \times T_f \times C}$ , which retrieves the original context features through a DETR-like architecture similar to that used for trajectory proposal generation. This hierarchical anchor-free proposal generation, together with the anchor-based refinement process, results in a trajectory query that integrates both reward-driven intentions and detailed scene context.

**Bi-Mamba Decoder.** Since the trajectory token  $Q_T$  contains significant sequential properties in both time and spa-

tial domains, we employ a selective state-space model to capture the coupled relationships within the series of trajectory queries, inspired by the recent success of the Mamba architecture in sequential modeling. Specifically, we adopt a Bi-Mamba model to process the trajectory token, leveraging its bidirectional scanning mechanism for more comprehensive information capture. In this Bi-Mamba-enhanced decoding process, we predict trajectory offsets  $\Delta Y \in \mathbb{R}^{K \times T_f \times 2}$  and the probability of each hypothesis  $p$ .

To better leverage the bidirectional capabilities of the Bi-Mamba structure, we design a learnable dual mode query  $Q_M \in \mathbb{R}^{K \times 2 \times C}$ , containing two classification (CLS) tokens. These tokens,  $CLS_1$  and  $CLS_2$ , are appended before and after the trajectory query  $Q_T$ , as illustrated in Fig. 3. The tokens aggregate backward and forward features, respectively, offering a more comprehensive fusion compared to a unidirectional Mamba with a single classification token, as validated in our ablation results (see Tab. 7). After the Bi-Mamba process, the two CLS tokens are combined through element-wise addition for feature fusion. A mode self-attention module then enables interaction among modes, further enhancing the multimodality of predictions.

Finally, the mode token undergoes classification with a softmax function to generate probabilities, and the sequential trajectory tokens are decoded with a regression head to produce trajectory offsets. More details are in Appx. 6.2.

The final predicted trajectory  $Y$  is obtained by summing the trajectory proposal  $\bar{Y}$  and its associated offset  $\Delta Y$ , as follows:

$$Y = \bar{Y} + \Delta Y. \quad (1)$$

### 3.6. Training Objectives

Our entire pipeline involves multiple training objectives. The reward-driven intention reasoner includes two subtask objectives: the QIRL and OGM generator. The QIRL objective employs the loss  $\mathcal{L}_{IRL}$  as discussed in Sec. 3.4, while the OGM generator, denoted as  $\mathcal{L}_{OGM}$ , uses a focal Binary Cross-Entropy (BCE) loss.

For the trajectory decoder, the training objectives include a regression loss  $\mathcal{L}_{REG}$  and a classification loss  $\mathcal{L}_{CLS}$ . To optimize trajectory regression, we apply Huber loss to both the trajectory proposals and the refined trajectories. Additionally, to address the mode collapse issue, we employ a winner-takes-all strategy, commonly used in similar works, where only the candidate with the minimum displacement error is selected for backpropagation. For mode classification, we adopt a max-margin loss following [18].

The overall loss  $\mathcal{L}$  integrates these components and can be optimized in an end-to-end fashion:

$$\mathcal{L} = \mathcal{L}_{IRL} + \alpha \mathcal{L}_{OGM} + \beta \mathcal{L}_{REG} + \gamma \mathcal{L}_{CLS}, \quad (2)$$

where  $\alpha$ ,  $\beta$ , and  $\gamma$  are hyperparameters to balance each training objective. More details can be found in Appx. 6.3.

Method	MR <sub>6</sub> ↓	minADE <sub>6</sub> ↓	minFDE <sub>6</sub> ↓	brier-minFDE <sub>6</sub> ↓	Brier score ↓
DenseTNT [14]	0.1258	0.8817	1.2815	1.9759	0.6944
HiVT [45]	0.1267	<b>0.7735</b>	<b>1.1693</b>	<u>1.8422</u>	0.6729
SceneTransformer [23]	<u>0.1255</u>	<u>0.8026</u>	1.2321	1.8868	0.6547
DSP [43]	0.1303	0.8194	1.2186	1.8584	<u>0.6398</u>
FiM (Ours)	<b>0.1250</b>	0.8296	<u>1.2048</u>	<b>1.8266</b>	<b>0.6218</b>
SIMPL [44]	0.1165	0.7693	1.1545	1.7469	0.5924
HPNet [36]	<b>0.1075</b>	<b>0.7478</b>	<b>1.0856</b>	<b>1.6768</b>	0.5912
Wayformer [21]	0.1186	<u>0.7676</u>	1.1616	1.7408	0.5792
MultiPath++ [37]	0.1324	0.7897	1.2144	1.7932	<u>0.5788</u>
FiM (Ours)	<u>0.1087</u>	0.7795	<u>1.1199</u>	<u>1.6931</u>	<b>0.5732</b>

Table 1. Performance comparison on the **Argoverse 1** motion forecasting leaderboard. The upper group lists results of single models, while the lower group lists those with ensemble methods. The best and the second-best results are in **bold** and underlined, respectively.

Method	minFDE <sub>6</sub> ↓	brier-minFDE <sub>6</sub> ↓	Brier score ↓
QCNet [46]	0.551	1.180	0.629
DeMo [42]	0.543	1.169	0.626
FiM w/ GRT-S	<u>0.529</u>	1.147	0.617
FiM w/ GRT-M	0.530	<u>1.134</u>	<u>0.604</u>
FiM w/ GRT-L	<b>0.528</b>	<b>1.131</b>	<b>0.603</b>

Table 2. Performance comparison on the customized **Argoverse 2** validation set. Results for QCNet and DeMo are produced using their official implementations.

## 4. Experiments and Results

### 4.1. Experimental Settings

**Datasets.** We use various large-scale public datasets, Argoverse 1 [4], Argoverse 2 [38], and nuScenes [1], for training and evaluating our approach. All datasets provide rich HD map information. The Argoverse 1 dataset includes over 1,000 hours of driving data from Miami and Pittsburgh, with approximately 206k, 39k, and 78k sequences for training, validation, and testing, respectively. Each sequence consists of 50 timestamps sampled at 10 Hz. The task is to forecast the next 30 positions given the past 20 observed states (i.e.,  $T_h = 20$ ,  $T_f = 30$ ). The Argoverse 2 dataset collects 250k scenarios from six cities, with about 200k and 25k sequences for training and validation, respectively. Each sequence comprises 110 timestamps sampled at 10 Hz, with  $T_h = 50$  for observations and  $T_f = 60$  for future predictions. The nuScenes dataset, collected in Boston and Singapore, presents dense traffic and challenging driving scenarios. It includes around 32k, 8k, and 9k sequences for training, validation, and testing, respectively, with each sequence containing 16 timestamps at a 2 Hz sampling frequency. Herein, the task is to predict 12 future positions given the previous 4 observed states (i.e.,  $T_h = 4$ ,  $T_f = 12$ ).

**Metrics.** We follow the standard metrics to evaluate prediction performance, including Miss Rate ( $MR_K$ ), minimum Average Displacement Error ( $\text{minADE}_K$ ), minimum Final Displacement Error ( $\text{minFDE}_K$ ), Brier minimum Final Dis-

Method	minADE <sub>5</sub> ↓	MR <sub>5</sub> ↓	minADE <sub>10</sub> ↓	MR <sub>10</sub> ↓
P2T [6]	1.45	0.64	1.16	0.46
THOMAS [12]	1.33	0.55	1.04	0.42
PGP [7]	1.27	0.52	0.94	0.34
DeMo [42]	1.22	0.43	0.89	0.34
MacFormer [8]	1.21	0.57	0.89	0.33
Goal-LBP [41]	1.02	<u>0.32</u>	0.93	<u>0.27</u>
UniTraj [9]	<u>0.96</u>	0.43	<u>0.84</u>	0.41
FiM (Ours)	<b>0.88</b>	<b>0.31</b>	<b>0.78</b>	<b>0.23</b>

Table 3. Performance comparison on the **nuScenes** prediction leaderboard.  $\text{minADE}_5$  is the official ranking metric.

placement Error ( $\text{brier-minFDE}_K$ ), and Brier score. Specifically,  $MR_K$  calculates the proportion of scenarios where none of the  $K$  predicted trajectories are within 2.0 meters of the Ground Truth (GT) endpoint.  $\text{minADE}_K$  measures the average  $\ell_2$  distance between the closest predicted trajectory and the GT, while  $\text{minFDE}_K$  focuses solely on endpoint error. Furthermore, the Brier score and  $\text{brier-minFDE}_K$ , introduced by Argoverse, incorporate prediction confidence into performance evaluation.  $\text{brier-minFDE}_K$  can be obtained by adding the Brier score  $(1.0 - p)^2$  to  $\text{minFDE}_K$ , where  $p$  is the probability of the best forecast. Note that the Brier score is particularly valuable in real-world applications, as accurate prediction confidence and uncertainty are crucial for safe and efficient decision-making and planning.

### 4.2. Comparison with State of the Art

We conduct a comprehensive comparison of our approach against state-of-the-art methods on the Argoverse 1, Argoverse 2, and nuScenes motion forecasting datasets. We refer to our method as FiM (Foresight in Motion) for brevity.

**Argoverse 1.** Quantitative results on the Argoverse 1 test split are presented in Tab. 1. We compare our FiM against several representative published methods evaluated on this challenging benchmark. According to the single-model results (upper group), FiM achieves highly competitive performance compared with strong baselines, including direct trajectory forecasting models such as HiVT [45] and Scene-

Method	minFDE <sub>6</sub> ↓	brier-minFDE <sub>6</sub> ↓	Brier score ↓
Vanilla	2.185	2.879	0.694
C.A.	1.490	2.132	0.642
Ours	<b>1.008</b>	<b>1.602</b>	<b>0.594</b>

Table 4. Ablation on the reward-driven reasoning strategy.

OGM	Refine	minFDE <sub>6</sub> ↓	brier-minFDE <sub>6</sub> ↓	Brier score ↓
✓		1.059	1.670	0.611
	✓	1.175	1.801	0.626
✓	✓	<b>1.008</b>	<b>1.602</b>	<b>0.594</b>

Table 5. Ablation on the OGM and refinement modules.

Transformer [23], as well as goal-based models like DSP [43] and DenseTNT [14]. FiM particularly excels in terms of the Brier score, brier-minFDE<sub>6</sub>, and MR<sub>6</sub>, highlighting its robust predictive capabilities.

We also apply model ensembling techniques to further enhance overall performance. The ensemble result (lower group) demonstrates a substantial performance improvement, indicating the significant potential and upper-bound capability of our proposed framework. Compared to other leading published methods such as HPNet [36] and Wayformer [21], FiM consistently achieves competitive performance across various evaluation metrics, notably excelling in the Brier score. This result emphasizes that our reasoning-enhanced predictor effectively generates predictions with greater reliability and confidence, as intended.

**Argoverse 2.** To further validate the effectiveness of our intention reasoning strategy, we introduce a customized evaluation benchmark built on top of the Argoverse 2 validation split. Specifically, the task requires forecasting the first 30 future positions, while during training, models have access to the subsequent 30 positions exclusively as auxiliary intention supervision. Notably, all models are constrained to training for trajectory generation using supervision strictly aligned with the first 30 future positions. This setup simulates practical applications where long-term paths are available for intention learning. Given that the proposed QIRL module is agnostic to the supervision format, whether trajectories or paths, we develop three model variants that incorporate different horizons of future supervision for the GRT training. These variants, denoted as GRT-S, GRT-M, and GRT-L, correspond to reasoning modules trained with 30, 45, and 60 future timestamps, respectively.

We compare our FiM against the two top-performing open-source models on the Argoverse 2 leaderboard, DeMo [42] and QCNet [46]. As shown in Tab. 2, all FiM variants surpass these two strong baselines, demonstrating substantial gains attributed to the intention reasoning module. Moreover, the results further indicate that longer-term intention supervision significantly enhances prediction confidence, facilitating more reliable trajectory forecasts.

MLP	Bi-Mamba	Self-Attn.	brier-minFDE <sub>6</sub> ↓	Brier score ↓
✓			1.720	0.622
✓	✓		1.649	0.605
✓		✓	1.682	0.617
✓	✓	✓	<b>1.602</b>	<b>0.594</b>

Table 6. Ablation on components of Mamba-enhanced decoder.

Mamba	minFDE <sub>6</sub> ↓	brier-minFDE <sub>6</sub> ↓	Brier score ↓
Unidirectional	1.034	1.636	0.603
Bidirectional	<b>1.008</b>	<b>1.602</b>	<b>0.594</b>

Table 7. Comparison between Uni-Mamba and Bi-Mamba.

# of layers	minFDE <sub>6</sub> ↓	brier-minFDE <sub>6</sub> ↓	Brier score ↓
2	1.027	1.635	0.609
4	<b>1.008</b>	<b>1.602</b>	<b>0.594</b>
8	1.024	1.628	0.609

Table 8. Effect of Mamba layer depth.

**nuScenes.** We also evaluate FiM on the nuScenes dataset, as reported in Tab. 3. Our model delivers top-tier performance on this prediction benchmark, outperforming all current entries on the leaderboard and further validating the robustness and advanced capabilities of our proposed framework in addressing complex motion forecasting challenges.

### 4.3. Ablation Study

We perform in-depth ablation studies on the Argoverse validation set to assess the effectiveness of key components in our approach, keeping all experimental settings consistent for a fair comparison.

**Effects of Reward Heuristics.** We first examined the effectiveness of the reward-driven intention reasoner by removing the reasoning branch from the pipeline. As shown in Tab. 4, the performance of the vanilla architecture dropped significantly compared to our full model, underscoring the critical contribution of the reasoning process to overall performance. Additionally, we explored the specific impact of the QIRL module by replacing it with a cross-attention block for feature extraction. The results in Tab. 4 show that our QIRL module outperformed this alternative by a large margin, demonstrating that QIRL can effectively gather essential intention priors and provide informative guidance that benefits subsequent motion forecasting.

**Effects of the OGM & Refinement.** We further evaluated the impact of the auxiliary spatial-temporal OGM module and refinement by ablating each separately, as shown in Tab. 5. Both modules contributed significantly to the final performance. In particular, the performance boost from OGM confirms that modeling future interactions enhances prediction quality, underscoring the importance of intention reasoning for improving trajectory predictions.

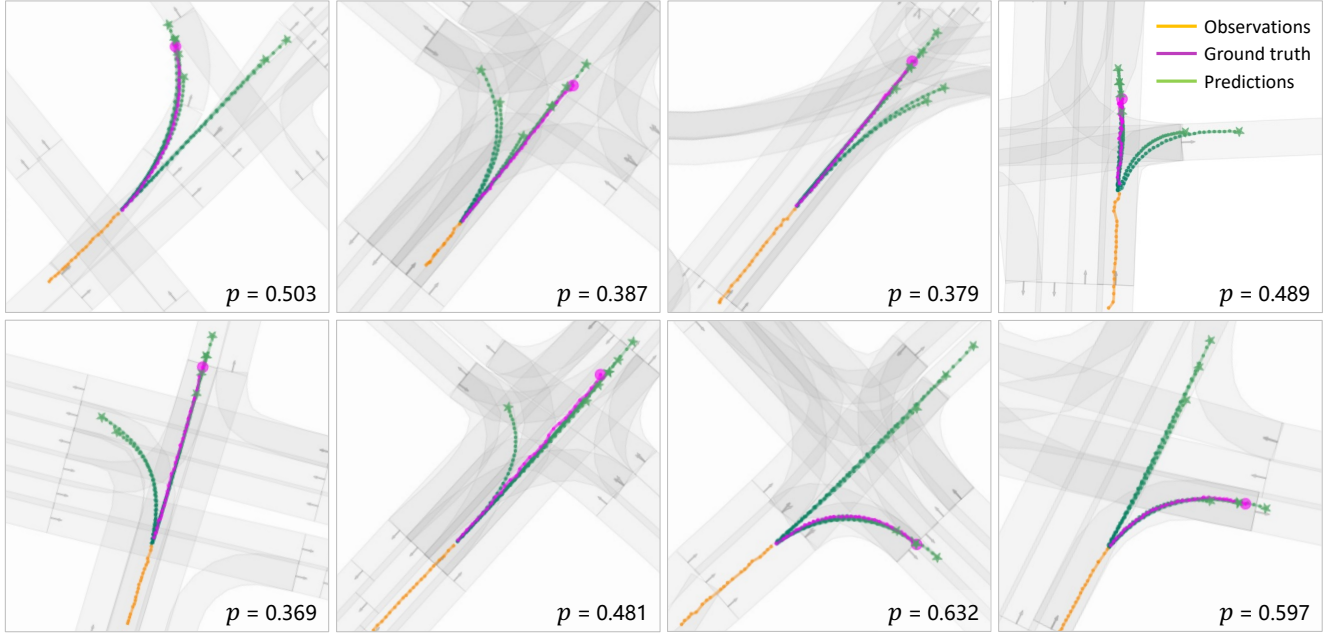


Figure 4. Qualitative results of our model on the Argoverse validation set. The historical trajectory, ground-truth future trajectory, and multimodal predictions are depicted in yellow, magenta, and green, respectively. The lower-right corner showcases the probability associated with the best forecast in terms of the endpoint. Other traffic participants are omitted to better highlight the predictions for the target agent.

**Effects of Components in Mamba-Based Decoder.** We conducted an ablation analysis of various decoder components to examine the benefits of the Mamba-like structure over conventional approaches. This analysis allowed us to determine whether the design offers meaningful feature extraction enhancements or constitutes over-engineering for trajectory decoding. Results in Tab. 6 highlight the advantages of this design. Compared to using MLPs as regression and classification heads, both the Bi-Mamba architecture and the self-attention mechanism between different modes significantly improved prediction performance and confidence. Additionally, we investigated the effect of the proposed dual mode tokens for classification, comparing it to a unidirectional Mamba model with a single mode token for aggregating trajectory query features. As shown in Tab. 7, the Bi-Mamba model demonstrated better performance, benefiting from its forward-backward scanning mechanism, which effectively fuses trajectory features into the two CLS tokens and validates the benefits of this design. We also examined the effects of different Mamba layer depths, as shown in Tab. 8. The results indicate that deeper layers can introduce unnecessary computational overhead and may also degrade performance due to overfitting, highlighting the importance of an optimal layer configuration for achieving strong performance.

#### 4.4. Qualitative Results

We present visualizations of our proposed approach across a variety of traffic scenarios from the Argoverse validation

set, as shown in Fig. 4. These qualitative results emphasize the strong capability of our model to generate accurate, feasible, and multimodal future trajectories that remain well-aligned with the scene layout under diverse conditions, including complex intersections and long-range forecasting scenarios. More qualitative results are provided in Appx. 7.

## 5. Conclusion

In this work, we reconceptualize trajectory prediction tasks from a planning perspective and advocate a “First Reasoning, Then Forecasting” strategy. We propose a novel and interpretable reward-driven intention reasoner, designed within a QIRL framework that combines the MaxEnt IRL paradigm with vectorized context representations through a query-centric pipeline, effectively providing informative intention priors for subsequent trajectory generation. In addition, we introduce a hierarchical DETR-like trajectory decoder integrated with a Bi-Mamba structure, which captures sequential trajectory dependencies and significantly enhances prediction accuracy and confidence. Experimental results demonstrate that our reasoning-enhanced predictor exhibits a strong capability for generating confident and reliable future trajectories that align well with the scene layout while achieving highly competitive performance against existing state-of-the-art models. Furthermore, our work underscores the critical role of intention reasoning in motion forecasting, substantiates the feasibility of RL paradigms in modeling driving behaviors, and establishes a promising baseline model for future research in trajectory prediction.

## Acknowledgment

This work was supported in part by the Hong Kong Ph.D. Fellowship Scheme, and in part by the HKUST-DJI Joint Innovation Laboratory.

## References

- [1] Holger Caesar, Varun Bankiti, Alex H Lang, et al. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 6
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 4
- [3] Yuning Chai, Benjamin Sapp, Mayank Bansal, and Dragomir Anguelov. Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. *arXiv preprint arXiv:1910.05449*, 2019. 2
- [4] Ming-Fang Chang, John Lambert, et al. Argoverse: 3d tracking and forecasting with rich maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8748–8757, 2019. 6
- [5] Henggang Cui, Vladan Radosavljevic, Fang-Chieh Chou, Tsung-Han Lin, Thi Nguyen, Tzu-Kuo Huang, Jeff Schneider, and Nemanja Djuric. Multimodal trajectory predictions for autonomous driving using deep convolutional networks. In *2019 international conference on robotics and automation (icra)*, pages 2090–2096. IEEE, 2019. 2
- [6] Nachiket Deo and Mohan M Trivedi. Trajectory forecasts in unknown environments conditioned on grid-based plans. *arXiv preprint arXiv:2001.00735*, 2020. 2, 6
- [7] Nachiket Deo, Eric Wolff, and Oscar Beijbom. Multimodal trajectory prediction conditioned on lane-graph traversals. In *Conference on Robot Learning*, pages 203–212, 2021. 1, 6
- [8] Chen Feng, Hangning Zhou, Huadong Lin, Zhigang Zhang, Ziyao Xu, Chi Zhang, Boyu Zhou, and Shaojie Shen. Macformer: Map-agent coupled transformer for real-time and robust trajectory prediction. *IEEE Robotics and Automation Letters*, 2023. 6
- [9] Lan Feng, Mohammadhossein Bahari, Kaouther Mersaoud Ben Amor, Éloi Zablocki, Matthieu Cord, and Alexandre Alahi. Unitraj: A unified framework for scalable vehicle trajectory prediction. *arXiv preprint arXiv:2403.15098*, 2024. 6
- [10] Tharindu Fernando, Simon Denman, Sridha Sridharan, and Clinton Fookes. Deep inverse reinforcement learning for behavior prediction in autonomous driving: Accurate forecasts of vehicle motion. *IEEE Signal Processing Magazine*, 38(1): 87–96, 2020. 3
- [11] Jiyang Gao, Chen Sun, Hang Zhao, Yi Shen, Dragomir Anguelov, Congcong Li, and Cordelia Schmid. Vectornet: Encoding hd maps and agent dynamics from vectorized representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11525–11533, 2020. 2, 4
- [12] Thomas Gilles, Stefano Sabatini, Dzmitry Tsishkou, Bogdan Stanciulescu, and Fabien Moutarde. Thomas: Trajectory heatmap output with learned multi-agent sampling. *arXiv preprint arXiv:2110.06607*, 2021. 6
- [13] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023. 2
- [14] Junru Gu, Chen Sun, and Hang Zhao. Densentn: End-to-end trajectory prediction from dense goal sets. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15303–15312, 2021. 1, 6, 7
- [15] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. 2
- [16] Zhiyu Huang, Haochen Liu, Jingda Wu, and Chen Lv. Conditional predictive behavior planning with inverse reinforcement learning for human-like autonomous driving. *IEEE Transactions on Intelligent Transportation Systems*, 2023. 3
- [17] Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv preprint arXiv:2412.16720*, 2024. 2
- [18] Ming Liang, Bin Yang, Rui Hu, Yun Chen, Renjie Liao, Song Feng, and Raquel Urtasun. Learning lane graph representations for motion forecasting. In *European Conference on Computer Vision*, pages 541–556, 2020. 2, 4, 5
- [19] Qiuqing Lu, Weiqiao Han, Jeffrey Ling, Minfa Wang, Haoyu Chen, Balakrishnan Varadarajan, and Paul Covington. Kemp: Keyframe-based hierarchical end-to-end deep model for long-term trajectory prediction. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 646–652. IEEE, 2022. 1
- [20] Sajjad Mozaffari, Omar Y Al-Jarrah, Mehrdad Dianati, Paul Jennings, and Alexandros Mouzakitis. Deep learning-based vehicle behavior prediction for autonomous driving applications: A review. *IEEE Transactions on Intelligent Transportation Systems*, 23(1):33–47, 2020. 2
- [21] Nigamaa Nayakanti, Rami Al-Rfou, Aurick Zhou, Kratharth Goel, et al. Wayformer: Motion forecasting via simple & efficient attention networks. In *2023 International Conference on Robotics and Automation (ICRA)*, pages 2980–2987. IEEE, 2023. 6, 7
- [22] Andrew Y Ng and Stuart Russell. Algorithms for inverse reinforcement learning. In *International conference on machine learning*, page 2, 2000. 3
- [23] Jiquan Ngiam, Benjamin Caine, Vijay Vasudevan, et al. Scene transformer: A unified architecture for predicting multiple agent trajectories. *arXiv preprint arXiv:2106.08417*, 2021. 3, 6, 7
- [24] Muleilan Pei, Dongping Wu, and Changhong Wang. Quadruped robot locomotion in unknown terrain using deep reinforcement learning. In *2020 3rd International Conference on Unmanned Systems (ICUS)*, pages 517–522. IEEE, 2020. 3
- [25] Muleilan Pei, Hao An, Bo Liu, and Changhong Wang. An improved dyna-q algorithm for mobile robot path planning

- in unknown dynamic environment. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 52(7):4415–4425, 2021. 3
- [26] Muleilan Pei, Jiayao Shan, Peiliang Li, Jieqi Shi, Jing Huo, Yang Gao, and Shaojie Shen. Sept: Standard-definition map enhanced scene perception and topology reasoning for autonomous driving. *IEEE Robotics and Automation Letters*, 10(7):7126–7133, 2025. 2
- [27] Muleilan Pei, Shaoshuai Shi, Lu Zhang, Peiliang Li, and Shaojie Shen. Goirl: Graph-oriented inverse reinforcement learning for multimodal trajectory prediction. *International Conference on Machine Learning (ICML)*, 2025. 4
- [28] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. 4
- [29] Daniela Ridel, Nachiket Deo, Denis Wolf, and Mohan Trivedi. Scene compliant trajectory forecast with agent-centric spatio-temporal grids. *IEEE Robotics and Automation Letters*, 5(2):2816–2823, 2020. 2
- [30] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 5
- [31] Stuart Russell. Learning agents for uncertain environments. In *Proceedings of the eleventh annual conference on Computational learning theory*, pages 101–103, 1998. 3
- [32] Shaoshuai Shi, Li Jiang, Dengxin Dai, and Bernt Schiele. Motion transformer with global intention localization and local movement refinement. *Advances in Neural Information Processing Systems*, 35:6531–6543, 2022. 1, 3
- [33] Haoran Song, Wenchao Ding, Yuxuan Chen, Shaojie Shen, Michael Yu Wang, and Qifeng Chen. Pip: Planning-informed trajectory prediction for autonomous driving. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pages 598–614. Springer, 2020. 1
- [34] Haoran Song, Di Luan, Wenchao Ding, Michael Y Wang, and Qifeng Chen. Learning to predict vehicle trajectories with model-based planning. In *Conference on Robot Learning*, pages 1035–1045, 2021. 1, 2
- [35] Richard S Sutton, Andrew G Barto, et al. *Reinforcement learning: An introduction*. MIT press Cambridge, 1998. 2
- [36] Xiaolong Tang, Meina Kan, Shiguang Shan, Zhilong Ji, Jinfeng Bai, and Xilin Chen. Hpnet: Dynamic trajectory forecasting with historical prediction attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15261–15270, 2024. 6, 7
- [37] Balakrishnan Varadarajan, Ahmed Hefny, Avikalp Srivastava, Khaled S Refaat, Nigamaa Nayakanti, et al. Multi-path++: Efficient information fusion and trajectory aggregation for behavior prediction. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 7814–7821. IEEE, 2022. 6
- [38] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, et al. Argoverse 2: Next generation datasets for self-driving perception and forecasting. *arXiv preprint arXiv:2301.00493*, 2023. 6
- [39] Zheng Wu, Liting Sun, Wei Zhan, Chenyu Yang, and Masayoshi Tomizuka. Efficient sampling-based maximum entropy inverse reinforcement learning with application to autonomous driving. *IEEE Robotics and Automation Letters*, 5(4):5355–5362, 2020. 3
- [40] Markus Wulfmeier, Dushyant Rao, Dominic Zeng Wang, Peter Ondruska, and Ingmar Posner. Large-scale cost function learning for path planning using deep inverse reinforcement learning. *The International Journal of Robotics Research*, 36(10):1073–1087, 2017. 4
- [41] Zhen Yao, Xin Li, Bo Lang, and Mooi Choo Chuah. Goal-lbp: Goal-based local behavior guided trajectory prediction for autonomous driving. *IEEE Transactions on Intelligent Transportation Systems*, 2023. 6
- [42] Bozhou Zhang, Nan Song, and Li Zhang. Decoupling motion forecasting into directional intentions and dynamic states. *arXiv preprint arXiv:2410.05982*, 2024. 6, 7
- [43] Lu Zhang, Peiliang Li, Jing Chen, and Shaojie Shen. Trajectory prediction with graph-based dual-scale context fusion. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11374–11381. IEEE, 2022. 6, 7
- [44] Lu Zhang, Peiliang Li, Sikang Liu, and Shaojie Shen. Simpl: A simple and efficient multi-agent motion prediction baseline for autonomous driving. *IEEE Robotics and Automation Letters*, 2024. 5, 6
- [45] Zikang Zhou, Luyao Ye, Jianping Wang, Kui Wu, and Kejie Lu. Hivt: Hierarchical vector transformer for multi-agent motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8823–8833, 2022. 6
- [46] Zikang Zhou, Jianping Wang, Yung-Hui Li, and Yu-Kai Huang. Query-centric trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17863–17873, 2023. 1, 3, 6, 7
- [47] Lianghui Zhu, Bencheng Liao, Qian Zhang, Xinlong Wang, Wenyu Liu, and Xinggang Wang. Vision mamba: Efficient visual representation learning with bidirectional state space model. *arXiv preprint arXiv:2401.09417*, 2024. 2
- [48] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, Anind K Dey, et al. Maximum entropy inverse reinforcement learning. In *AAAI*, pages 1433–1438. Chicago, IL, USA, 2008. 2, 3, 4