

HiMAP: History-aware Map-occupancy Prediction with Fallback

Yiming Xu¹, Yi Yang², Hao Cheng³, Monika Sester¹

Abstract—Accurate motion forecasting is critical for autonomous driving, yet most predictors rely on multi-object tracking (MOT) with identity association, assuming that objects are correctly and continuously tracked. When tracking fails due to, e.g., occlusion, identity switches, or missed detections, prediction quality degrades and safety risks increase. We present HiMAP, a tracking-free, trajectory prediction framework that remains reliable under MOT failures. HiMAP converts past detections into spatiotemporally invariant historical occupancy maps and introduces a historical query module that conditions on the current agent state to iteratively retrieve agent-specific history from unlabeled occupancy representations. The retrieved history is summarized by a temporal map embedding and, together with the final query and map context, drives a DETR-style decoder to produce multi-modal future trajectories. This design lifts identity reliance, supports streaming inference via reusable encodings, and serves as a robust fallback when tracking is unavailable. On Argoverse 2, HiMAP achieves performance comparable to tracking-based methods while operating without IDs, and it substantially outperforms strong baselines in the no-tracking setting, yielding relative gains of 11% in FDE, 12% in ADE, and a 4% reduction in MR over a fine-tuned QCNet. Beyond aggregate metrics, HiMAP delivers stable forecasts for all agents simultaneously without waiting for tracking to recover, highlighting its practical value for safety-critical autonomy. The code is available under: <https://github.com/XuYiming83/HiMAP>.

I. INTRODUCTION

In autonomous driving and robotic navigation, accurately predicting the future trajectories of surrounding agents is critical for understanding the environment and making safe decisions. A primary motivation for trajectory prediction is that reliable estimates of nearby agents’ future motions enable automated systems to plan safer paths. However, trajectory prediction is inherently unstable, as it is embedded within a complex pipeline of sequential tasks, including detection, segmentation, tracking, trajectory forecasting, and planning [1], [2]. Yet, integrating existing methods [3]–[5] into a single system often introduces instability across the pipeline.

Trajectory forecasting methods heavily depend on detection and tracking modules, assuming reliable tracks within a certain range to generate multi-modal predictions of future motions. In practice, however, tracking is prone to failures

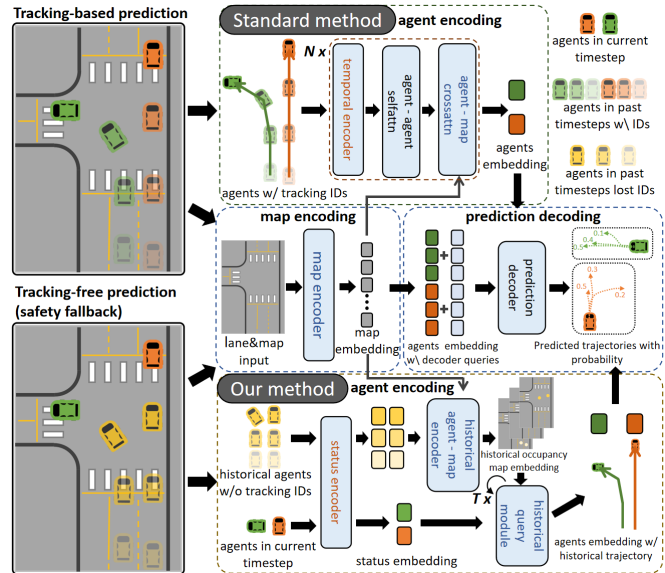


Fig. 1: Comparison between tracking-based prediction and our tracking-free fallback. Tracking-based methods rely on stable identity association, but fail when tracking breaks. Our tracking-free module provides a complementary safety mechanism by reconstructing history from historical occupancy maps, ensuring reliable prediction even under tracking failures. Our method matches the current agent state to historical occupancy maps, implicitly recovering its past states without explicit IDs for robust trajectory prediction.

such as flickering, identity drift, and ID switches, especially under occlusion or in dense multi-agent scenes [6]–[8]. These issues often lead to incomplete or erroneous histories, causing prediction accuracy and robustness to degrade.

In this work, we introduce **HiMAP**, a trajectory prediction module designed to mitigate failures caused by unstable or missing tracking. Our key idea is to track agents implicitly by leveraging historical detections alone—without requiring consistent tracking identities. Even if the tracking module collapses entirely, HiMAP can still produce reliable trajectory forecasting based solely on cached historical detections, enhancing the robustness and safety of autonomous systems.

Our method converts past time-based trajectories into space-based occupancy maps for tracking-free training. Specifically, we first encode each agent’s state and high-definition (HD) map information (including lane geometry). At each timestep, we integrate detected agent states into the map to build a sequence of historical occupancy maps. The observation module then uses the current agent state and the full history of occupancy maps to construct a query that iteratively decodes each past frame, implicitly

¹Yiming Xu and Monika Sester are with the Institute of Cartography and Geoinformatics, Leibniz University Hannover, Appelstr. 9a, 30167 Hannover, Germany {Yiming.Xu, Monika.Sester}@ikg.uni-hannover.de

²Yi Yang is with the Institute of Information Processing, Leibniz University Hannover, Schneiderberg 32, 30167 Hannover, Germany yangyi@tnt.uni-hannover.de

³Hao Cheng is with the Faculty of Geo-Information Science and Earth Observation, University of Twente, 7522 NH Enschede, The Netherlands. h.cheng-2@utwente.nl

reconstructing the agent’s historical trajectory. Finally, multi-modal future trajectories are predicted from the decoded history, the current agent state, and the map context, as shown in Fig. 1. The core principle is that *by projecting detections as occupancy information on per-frame maps, HiMAP can recover agents’ motion history even without explicit track IDs*. By aggregating these observations over time, the module stabilizes trajectory predictions and safeguards downstream decision-making in safety-critical systems.

Our main contributions are summarized as follows:

- We propose a trajectory prediction framework that remains reliable when tracking is unstable or unavailable. Without requiring tracking identities, our method reconstructs historical states from detections alone and predicts future trajectories with improved stability.
- We design a scene-level encoding of rich map context and a recurrent observation mechanism over historical occupancy maps. Parameters are shared across timesteps, and the module is aligned with standard agent–map encoders, enabling direct integration into existing forecasting models with minimal overhead.
- On the Argoverse 2 dataset, HiMAP achieves performance close to baselines that assume perfect tracking. In tracking failure scenarios, our method significantly improves robustness and substantially narrows the performance gap to tracking-based approaches. In particular, compared with tracking-based baselines in the tracking failure setting, HiMAP achieves relative improvements of 11% in FDE, 12% in ADE, and 4% in MR.

II. RELATED WORKS

A. Tracking in Motion Forecasting

Most state-of-the-art forecasting models rely on multi-object tracking (MOT) to provide identity-consistent histories for agents. Trackers such as AB3DMOT [9], CenterPoint with tracking [10], and recent detection-based MOT methods including ByteTrack [11], MOTRv2 [7], and BoT-SORT [12] are commonly adopted in autonomous driving benchmarks, and their outputs serve as the historical context for prediction [3], [13]–[15]. However, tracking remains brittle in crowded or occluded scenes, often producing fragmented trajectories or ID switches. In fact, most reported IDF1 [16] scores are below 72%, reflecting a substantial risk that can significantly degrade forecasting accuracy. Despite this limitation, nearly all forecasting methods still assume reliable tracking. To our knowledge, trajectory prediction in the complete absence of consistent IDs has not been explicitly studied. This gap is critical, as tracking failures are inevitable in practice and can persist long enough to jeopardize safety. Our work addresses this limitation by reconstructing historical trajectories directly from unlabeled detections, enabling reliable forecasting without stable tracking. Prior work has explored forecasting from detections without explicit track IDs, often relying on tracker intermediates such as affinity scores [17], [18]. In this work, we explicitly study trajectory prediction under tracking failures without requiring any affinity estimation or other tracker outputs.

B. Scene Context Representation

Scene context in motion forecasting refers to the complex traffic environment, including HD maps and multi-agent interactions, which autonomous vehicles must interpret to avoid collisions and conflicts [19], [20]. Early works rasterized the environment, storing maps and agent states as image-like tensors and applying convolutional networks for feature fusion [21], [22]. With the introduction of VectorNet [23], scene context was unified into a vectorized representation of points, polylines, and polygons for both maps and trajectories. This paradigm has since become dominant [3], [14], [15], [24]–[26], as vectorized representations provide broader receptive fields, finer map details, and lower computational cost. They further support encoding methods with permutation invariance and long-range dependency modeling, such as graph networks [27] and Transformers [28], enabling direct interaction among agents and maps.

Most forecasting methods normalize the scene to the target agent’s coordinates, aligning its heading with the y -axis (longitudinal direction) [15], [25], [26]. While effective, this approach requires re-normalizing at every timestep, which prevents reuse of encoded context. To address this limitation, QCNet [3] introduced spatiotemporally invariant encodings, ensuring that agent and map features remain stable across timesteps and locations. This property not only supports efficient streaming prediction, but also provides the foundation for exploring trajectory forecasting under missing or unstable tracking, which we build upon in this work.

C. Multimodal Trajectory Prediction

Agent behavior is inherently uncertain, leading to multi-modal distributions of possible futures in complex environments. Multi-modal decoders are therefore critical for trajectory prediction. Generative models such as GANs [29], CVAEs [30], and diffusion models [31], [32] have been widely explored. Another line of work employs winner-takes-all strategies, where multiple prediction heads specialize in different motion modes [15], [25], [26].

Recent advances increasingly leverage Transformer decoders [28], inspired by DETR [33], where each query directly decodes into a candidate trajectory, naturally supporting multi-modal forecasting [3], [14], [15]. In this work, we adopt this query-based decoding paradigm, which effectively combines scene context with past trajectories to produce diverse and accurate future predictions.

III. METHOD

A. Problem Formulation

We consider a scenario with A agents surrounding the autonomous vehicle at discrete timesteps $t \in \{1, \dots, T\}$. During online operation, the perception stack provides to the predictor the states of detected agents at fixed intervals. When tracking is available, a data association module links detections across time to form past trajectories. For the i -th agent at time t , we denote its state feature as $\mathbf{s}_i^t = \{p_{i,x}^t, p_{i,y}^t, v_{i,x}^t, v_{i,y}^t, h_i^t, a_i^t\}$, where p_i is the position, v_i is the velocity, h_i is the heading, and a_i is the agent category

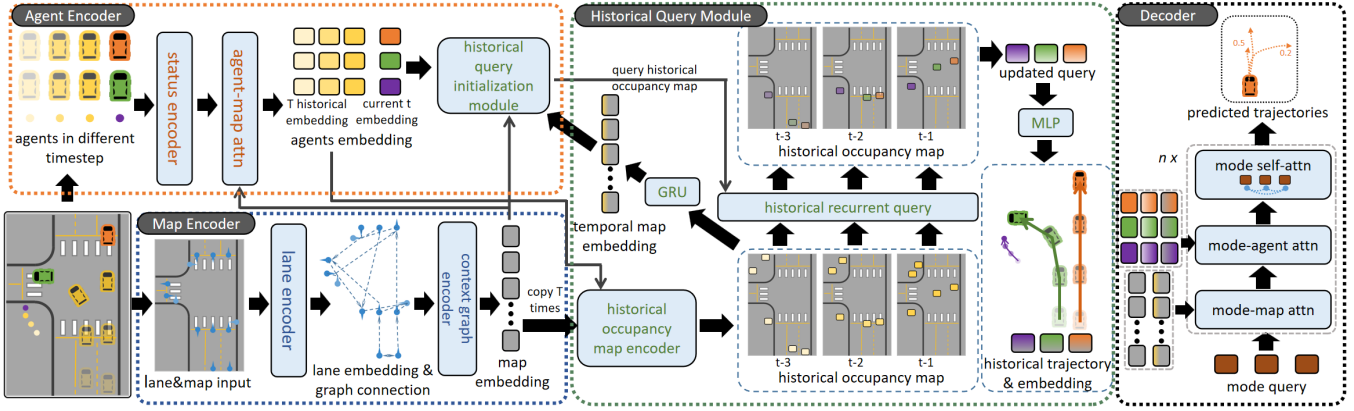


Fig. 2: Overview of our HiMAP pipeline. The framework consists of four main components: **Agent and Map Encoding**, which embeds agent states and HD map elements in spatiotemporally invariant local frames; **Historical Occupancy Map Encoder**, which aggregates per-frame agent-lane interactions into occupancy representations without relying on tracking IDs; **Historical Query Module**, which initializes a history-aware query from the current agent state and iteratively attends to past occupancy maps to reconstruct agent-specific trajectories; **Future Trajectory Decoder**, a DETR-style query decoder that generates multi-modal predictions from the reconstructed history, final query, and map context. This design provides a robust fallback when tracking fails, enabling reliable forecasting directly from historical detections.

(e.g., car, bus, pedestrian). The HD map is represented by M lane polygons, where the j -th polygon is $L_j = \{P_j, a_{L,j}\}$. Here $P_j = \{p_{j,k}\}$ is a set of sampled points with $p_{j,n} = \{p_{j,n,x}, p_{j,n,y}, a_{L,j,n}\}$ denoting position and lane point attributes, and $a_{L,j}$ encodes polygon-level attributes such as lane type or intersection membership. Given an observation of T steps for an agent, the prediction module aims to generate K possible future trajectories $\{Y_i^{(k)}\}_{k=1}^K$ along with their probabilities $\{\pi_i^{(k)}\}_{k=1}^K$.

In practice, however, tracking can be unstable or unavailable, leading to the loss of consistent agent identities over time. In such cases, the input within the observation window is no longer the identity-aware sequence $\{s_i^t\}$ but instead a set of unlabeled detections per frame, denoted as $\mathcal{D}^t = \{\tilde{s}_n^t\}_{n=1}^{N_t}$, where each \tilde{s}_n^t has the same structure as s_i^t but without a persistent agent ID, and there is no guaranteed correspondence across timesteps. The full observation history is then $\mathcal{D}_{1:T} = \{\mathcal{D}^t\}_{t=1}^T$. The prediction objective remains the same: given the current target agent and the HD map $\{L_j\}_{j=1}^M$, predict K plausible future trajectories with probabilities, now conditioned on the identity-agnostic history $\mathcal{D}_{1:T}$ rather than tracked sequences. Our method is explicitly designed to handle both regimes by leveraging historical detection sets without relying on tracking IDs.

B. Baseline Model

We adopt QCNet [3] as our baseline to investigate tracking failures. QCNet has two main advantages. First, it achieves excellent performance on the Argoverse 2 dataset, making it a strong backbone. Second, it employs spatiotemporally invariant encodings, ensuring consistent representations of agents and map segments across timesteps and locations.

Unlike agent-centric methods that re-normalize scene context to the target agent’s position and heading at each frame, QCNet avoids redundant re-encoding. Such normalization prevents reuse of past context and requires ID-based trajec-

tory storage, which is infeasible when tracking is unreliable. Because the map encoding space is invariant and agent encodings remain stable over time, historical states can be directly projected onto the map to construct spatiotemporally invariant occupancy information. Thus, QCNet supports efficient streaming prediction and joint forecasting of all agents under a unified, ID-agnostic representation, making it an ideal baseline for studying robustness to tracking failures.

C. Overall Architecture

Our prediction framework HiMAP consists of four modules shown in Fig. 2. The **Scene Encoding** module provides spatiotemporally invariant representations of agents and map elements. The **Historical Occupancy Map Encoder** aggregates per-frame agent-lane interactions into occupancy features without relying on tracking IDs. The **Historical Query Module** reconstructs agent-specific history by attending to occupancy maps conditioned on the current state. Finally, the **Future Trajectory Decoder** generates diverse multi-modal forecasts from the reconstructed history, the current state, and the map context. This modular design ensures robust prediction even under tracking failures.

1) *Scene Encoding*: We follow QCNet [3]: agents and lanes are embedded in local frames for invariance. Lane polylines are encoded via cross-attention, while agent motion states are encoded with Fourier features and semantic attributes. Unlike QCNet, we do not use agent IDs: we store the set of visible agent embeddings at each timestep, forming $E_A \in \mathbb{R}^{T \times A_t \times D}$, together with lane embeddings $E_L \in \mathbb{R}^{M \times D}$, where A_t denotes the number of visible agents at timestep t , M the number of map elements, T the observation length, and D the feature dimension. For streaming inference, only the current frame is recomputed.

2) *Historical Occupancy Map Encoder*: We introduce a historical occupancy map encoder that allows the model to incorporate scene information from each past frame and capture long-term agent interactions. Since tracking data is

unavailable, each frame contains an independent set of agent embeddings E_A^t . Leveraging the spatiotemporally invariant encodings, the connections between agents and lanes are represented as a graph, where each agent embedding is linked to nearby lane segments through relative positional encodings.

For each historical frame, the occupancy map representation E_{occ}^t is defined as

$$\begin{aligned} E_{\text{occ}}^t &= E_L + \hat{g}^t \odot \Phi(E_A^t \rightarrow E_L, d(p_A^t \rightarrow p_L)), \\ \hat{g}^t &= \sigma\left(\text{MLP}([E_A^t + d(p_A^t \rightarrow p_L)] \oplus E_L)\right). \end{aligned} \quad (1)$$

where E_L denotes lane embeddings, $d(p_A^t \rightarrow p_L)$ is the relative position encoding from agents to lanes, \oplus denotes concatenation, and $\Phi(\cdot)$ is a cross-attention operator defined on the agent–lane graph. The gating vector \hat{g}^t with Sigmoid activation σ adaptively modulates the influence of agents on lane occupancy, reducing the contribution of agents with limited relevance to lane dynamics. In this way, each lane’s historical occupancy encoding is derived from both its own features and the gated contributions of nearby agents, conditioned on their relative positions. This design ensures that only behaviorally relevant agents exert a strong influence on the historical occupancy representation, thereby improving robustness to noisy or weak interactions. Collecting all timesteps, we obtain the historical occupancy tensor of shape $[T, M, D]$, where T is the number of observed frames, M is the number of lane segments.

3) *Historical Query Module*: To incorporate past occupancy information, we first summarize the sequence of historical occupancy maps $[T, M, D]$ into a temporal map embedding $\tilde{E}_L \in \mathbb{R}^{M \times D}$ using a GRU. This embedding compactly represents long-term map–agent interactions across T frames.

At the current timestep t_c , the observed agent state $E_A^{t_c}$ is connected to its surrounding map within an r -meter radius. Since map coordinates are spatiotemporally invariant, the relative position encoding $d(p_L \rightarrow p_A^{t_c})$ is shared across the temporal map embedding \tilde{E}_L and all historical occupancy maps E_{occ}^t . The initial historical query is then formed by cross-attention between the current agent and the temporal map embedding:

$$A_{\text{query}}^{t_c} = E_A^{t_c} + \hat{g}^{t_c} \odot \Phi(\tilde{E}_L \rightarrow E_A^{t_c}, d(p_L \rightarrow p_A^{t_c})), \quad (2)$$

where $\hat{g}^{t_c} = \sigma(\text{MLP}([\tilde{E}_L + d(p_L \rightarrow p_A^{t_c})] \oplus E_A^{t_c}))$ is a gating vector, \oplus denotes concatenation, and $\Phi(\cdot)$ is cross-attention on the agent–map graph. This design allows the query to integrate the current agent state with coarse localization and historical dynamics summarized by \tilde{E}_L .

The historical query is then iteratively updated by attending to each occupancy map E_{occ}^t in reverse temporal order, analogous to a recurrent update:

$$A_{\text{query}}^{t-1} = A_{\text{query}}^t + \hat{g}^t \odot \Phi(E_{\text{occ}}^t \rightarrow A_{\text{query}}^t, d(p_L \rightarrow p_A^t)), \quad (3)$$

with $\hat{g}^t = \sigma(\text{MLP}([E_{\text{occ}}^t + d(p_L \rightarrow p_A^t)] \oplus A_{\text{query}}^t))$. Each updated query A_{query}^{t-1} is projected by an MLP into a local displacement vector h_p^{t-1} , aligned with the coordinate frame

of the current timestep. By iterating this process over T_h historical steps with shared parameters, we obtain the reconstructed historical trajectory in local coordinates, analogous to an RNN unrolling through time. Importantly, our method queries each historical occupancy map conditioned on the current agent state, thereby implicitly retrieving the agent-specific information from unlabeled occupancy representations and decoding a coherent historical trajectory, which together with the final historical query $A_{\text{query}}^{T_h}$, the agent’s current state, and the map context, forms the input to our multi-modal trajectory decoder.

4) *Future Trajectory Decoder*: Built on the reconstructed history h_p , the final historical query $A_{\text{query}}^{T_h}$, and the map context E_L together with the temporal map embedding \tilde{E}_L , we adopt a DETR-style decoder with K learnable mode queries $\{q_k\}_{k=1}^K$. Each query with decoded historical trajectory attends to the scenario context and decodes one candidate future trajectory, yielding K diverse hypotheses with associated probabilities:

$$\{\hat{Y}^{(k)}, \pi^{(k)}\}_{k=1}^K = \mathcal{D}(\{q_k\}_{k=1}^K, A_{\text{query}}^{T_h}, h_p, \tilde{E}_L, E_L), \quad (4)$$

where $\sum_{k=1}^K \pi^{(k)} = 1$. This query-based design provides a compact, end-to-end mechanism for multi-modal forecasting while naturally leveraging the reconstructed history and current agent state.

Unlike QCNNet, which combines anchor-free and anchor-based decoding in a two-stage manner, we employ only the anchor-free design. The anchor-based refinement offers limited benefit in our setting, as our decoder already incorporates fine-grained historical features; our experiments show that it reduces FDE by only 0.03 while substantially increasing computational cost. Hence, the anchor-free stage alone is sufficient and more efficient for our framework, striking a better balance between accuracy and efficiency.

D. Training

Following prior work [3], [25], we model the distribution of predicted trajectories as a mixture of Laplace components. The regression loss is defined as the negative log-likelihood of the ground-truth trajectory under the closest mode k_{best} :

$$\mathcal{L}_{\text{reg}} = \frac{1}{t_f} \sum_{t=1}^{t_f} -\log P(Y_t | \mu_t^{k_{\text{best}}}, b_t^{k_{\text{best}}}), \quad (5)$$

where μ_t and b_t are the location and scale parameters of the Laplace distribution decoded by the trajectory predictor, Y_t is the ground-truth position, and t_f is the prediction horizon.

To supervise the mode probabilities, we adopt a cross-entropy classification loss:

$$\mathcal{L}_{\text{cls}} = \sum_{k=1}^K -\pi_k \log(\hat{\pi}_k), \quad (6)$$

where $\hat{\pi}_k$ is the predicted probability of mode k , π_k is the one-hot ground-truth indicator, and $\sum_{k=1}^K \hat{\pi}_k = 1$.

In addition, we directly supervise the reconstruction of historical trajectories decoded by the historical query module

TABLE I: Results on the Argoverse 2 motion forecasting benchmark. All leaderboard methods assume perfect tracking, whereas HiMAP is evaluated under the no-tracking setting (historical detections only). HiMAP achieves competitive results and clearly outperforms QCNet baselines when tracking fails.

Method	w/o Tracking	minFDE ₁ ↓	minADE ₁ ↓	minFDE ₆ ↓	minADE ₆ ↓	MR ₆ ↓	b-minFDE ₆ ↓
FRM [34]		5.93	2.37	1.81	0.89	0.29	2.47
HDGT [35]		5.37	2.08	1.60	0.84	0.21	2.24
MTR [13]		4.39	1.74	1.44	0.73	0.15	1.98
HPTR [36]		4.61	1.84	1.43	0.73	0.19	2.03
HeteroGCN [14]		4.40	1.72	1.34	0.69	0.18	1.90
ProphNet [15]		4.74	1.80	1.33	0.68	0.18	1.88
QCNet [3]		4.30	1.69	1.29	0.65	0.16	1.91
SmartRefine [37]		4.17	1.65	1.23	0.63	0.15	1.86
DeMo [38]		3.74	1.49	1.17	0.61	0.13	1.84
QCNet w/o tracking (baseline)	✓	11.59	5.88	3.23	2.00	0.52	3.89
DeMo w/o tracking	✓	9.99	5.14	3.16	2.09	0.47	3.85
QCNet w/ fine-tuning (baseline)	✓	5.16	2.10	1.49	0.77	0.21	2.25
HiMAP (Ours)	✓	4.61	1.81	1.33	0.68	0.17	1.96

TABLE II: Results on the Argoverse 1 motion forecasting validation set.

Method	minFDE ₆ ↓	minADE ₆ ↓	MR ₆ ↓
QCNet [3]	0.86	0.63	0.08
QCNet w/o tracking	3.05	1.99	0.37
QCNet w/ fine-tuning	1.24	0.79	0.12
HiMAP (Ours)	0.94	0.66	0.09

using an ℓ_2 loss:

$$\mathcal{L}_{\text{his}} = \frac{1}{t_h} \sum_{t=1}^{t_h} \|\hat{h}_{p_t}^{\text{his}} - h_{p_t}^{\text{his}}\|_2^2, \quad (7)$$

where \hat{Y}_t^{his} and Y_t^{his} denote the predicted and ground-truth historical positions, respectively, and t_h is the number of historical steps.

The final training objective is a weighted sum of the three terms:

$$\mathcal{L} = \mathcal{L}_{\text{reg}} + \alpha \mathcal{L}_{\text{cls}} + \beta \mathcal{L}_{\text{his}}, \quad (8)$$

where α and β are balancing coefficients.

IV. EXPERIMENTS

A. Experimental Setup

1) *Dataset*: We evaluate our model on the Argoverse 2 motion forecasting dataset [39], which contains 25,000 real-world driving scenarios split into training, validation, and test sets with 199,908, 24,988, and 24,984 samples, respectively. Each sequence spans 11 seconds at 10 Hz, providing sufficiently long temporal horizons for both history and prediction. The large scale and extended duration of this dataset make it well-suited for assessing the temporal dependency and stability of our method.

2) *Metrics*: We adopt standard evaluation metrics [24]: minADE_K , minFDE_K , b-minFDE_K , and MR_K . minADE_K measures the average displacement error between the ground-truth trajectory and the closest of K predicted trajectories across all timesteps. minFDE_K evaluates the final displacement error at the prediction horizon. b-minFDE_K further incorporates the confidence score of each predicted trajectory. MR_K denotes the miss rate, defined as the proportion of cases where the final displacement error exceeds 2 meters.

3) *Implementation Details*: We train our model on 8 NVIDIA RTX 3090 GPUs using the AdamW optimizer [40]. The training schedule consists of 64 epochs with a batch size of 16 and an initial learning rate of 1×10^{-4} . We adopt a warm-up phase of 10,000 steps, followed by a cosine annealing scheduler [41] to decay the learning rate. The number of heads in all multi-head attention layers is set to 8. No ensemble methods or data augmentation are applied in our experiments.

B. Comparison on Argoverse 2 Benchmark

We evaluate our proposed **HiMAP** on the Argoverse 2 motion forecasting benchmark [39]. For fair comparison, ensemble methods are excluded. All leaderboard entries assume perfectly accurate tracking, whereas our method is designed as a fallback solution under tracking failures, relying only on historical detections without identity information. As shown in Table I, despite this stricter setting, HiMAP still outperforms most existing forecasting approaches on the leaderboard.

To further validate the practicality of our method, we investigate the performance of QCNet [3] under tracking failures. Using the publicly released checkpoint and masking out tracking IDs, QCNet’s performance drops sharply due to distributional shift: minFDE_6 increases from 1.29 to 3.23, minADE_6 from 0.65 to 2.00, and MR_6 from 16% to 52%, rendering the predictions nearly unusable for autonomous driving. We then fine-tune the QCNet decoder under the no-tracking setting by randomly masking past agent trajectories, which partially recovers performance but still yields inferior results ($\text{minFDE}_6 = 1.49$, $\text{minADE}_6 = 0.77$, $\text{MR}_6 = 21\%$).

By contrast, HiMAP achieves $\text{minFDE}_6 = 1.33$, $\text{minADE}_6 = 0.68$, and $\text{MR}_6 = 17\%$, corresponding to relative improvements of 11% in FDE, 12% in ADE, and 4% in MR over the fine-tuned QCNet baseline. These gains highlight the robustness of our approach in handling tracking failures, effectively bridging the gap between tracking-based forecasting methods and realistic perception scenarios. Overall, HiMAP demonstrates both strong predictive performance and robustness to tracking failures, highlighting its potential as a reliable module for safety-critical autonomous driving.

TABLE III: Ablation studies on Historical Query Module components. Experimental results are based on the Argoverse 2 validation set.

Index	Historical Query Module				Metrics				
	recurrent query	hi. occ. map	hi. query init.	hi. temp. map	reconstruct hi. ADE \downarrow	minFDE \downarrow	minADE \downarrow	MR \downarrow	b-minFDE \downarrow
1					—	1.46	0.85	0.21	2.06
2	✓				0.27	1.40	0.83	0.19	2.01
3	✓	✓			0.21	1.38	0.79	0.18	2.00
4	✓	✓	✓		0.14	1.37	0.79	0.17	1.99
5	✓	✓	✓	✓	0.12	1.33	0.76	0.17	1.96

TABLE IV: Ablation studies on decoder inputs. Results are reported on the Argoverse 2 validation set.

Index	Decoder				Metrics			
	current agent status & map	hi. temporal map	updated query	recurrent	minFDE \downarrow	minADE \downarrow	MR \downarrow	b-minFDE \downarrow
1	✓			✓	1.46	0.85	0.21	2.06
2	✓	✓		✓	1.39	0.79	0.18	2.01
3	✓	✓	✓		1.41	0.83	0.19	2.03
4	✓	✓	✓	✓	1.33	0.76	0.17	1.96

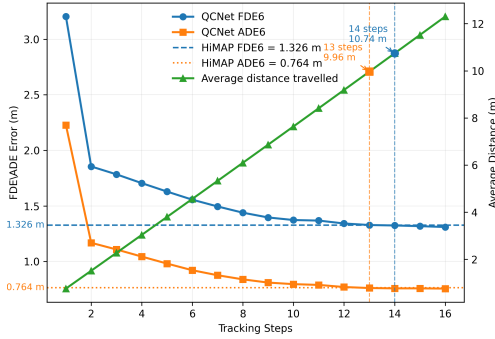


Fig. 3: Comparison between HiMAP and QCNet on the Argoverse 2 validation set under different numbers of available tracking steps. HiMAP maintains fixed performance without requiring tracking, shown as horizontal dashed lines. QCNet gradually improves as more tracked history is available, surpassing HiMAP after 13–14 steps. The green curve indicates the average distance traveled per timestep.

C. Results on Argoverse 1 Dataset

We further evaluate our proposed method on the Argoverse 1 validation set for comparison with our baseline in Table II. Following the settings in QCNet [3], we reproduce its results on Argoverse 1 and benchmark our tracking-free variant under the same conditions. The results show that our method significantly improves performance in the absence of tracking and achieves accuracy close to tracking-based models, thereby ensuring safer and more reliable trajectory forecasting for autonomous driving.

D. Impact of Tracking Failures

Figure 3 illustrates the necessity of our study and its implications for autonomous driving safety. The plot compares the performance of HiMAP with QCNet on the Argoverse 2 validation set under varying numbers of available tracking steps.

HiMAP achieves fixed performance on this benchmark ($\text{minFDE}_6 = 1.33\text{ m}$ and $\text{minADE}_6 = 0.76\text{ m}$), shown as the horizontal dashed lines. In contrast, the curves show how QCNet’s errors decrease as more historical tracking steps become available. The green curve represents the average

distance traveled by vehicles per timestep in the validation set, which increases approximately linearly.

The figure highlights a critical safety gap. When tracking fails, QCNet requires about 13 steps (1.3 s) for ADE and 14 steps (1.4 s) for FDE to surpass HiMAP’s fixed accuracy. During this recovery period, vehicles travel on average 10.7 m and fast-moving agents can cover over 30 m, before QCNet regains its predictive reliability. Such delays pose substantial risks for real-time decision-making in safety-critical systems. By contrast, HiMAP immediately provides stable predictions without relying on tracked histories, improving short-term safety. Within the first second, HiMAP reduces FDE by 20% and ADE by 23% compared to QCNet without tracking on average, ensuring reliable trajectory forecasts precisely when they are most critical. This underscores the importance of designing tracking-agnostic forecasting modules. This robustness entails a trade-off between accuracy and efficiency, as each additional history-reconstruction step increases average latency by 1.6%.

E. Ablation Study

We conduct ablation studies to analyze the contribution of each component in the proposed Historical Query Module and Decoder, with results reported in Table III and Table IV.

Table III evaluates the Historical Query Module. The first row removes the entire module, which is equivalent to fine-tuned QCNet. In this case, prediction relies only on the current agent state and map, failing to reconstruct historical trajectories and yielding poor results. The second row allows the current agent state to directly query past agent states without occupancy maps. This design becomes infeasible, as the computational complexity at each timestep scales exponentially with the number of agents, and the query vector is easily entangled with unrelated agents across timesteps, leading to degraded reconstruction. Moreover, when the observed history of an agent is shorter than the required reconstruction horizon, this approach introduces substantial noise. In contrast, using occupancy maps enables the model to infer unobserved timesteps from map structure, yielding more plausible historical trajectories. The third row introduces historical occupancy maps, which substantially

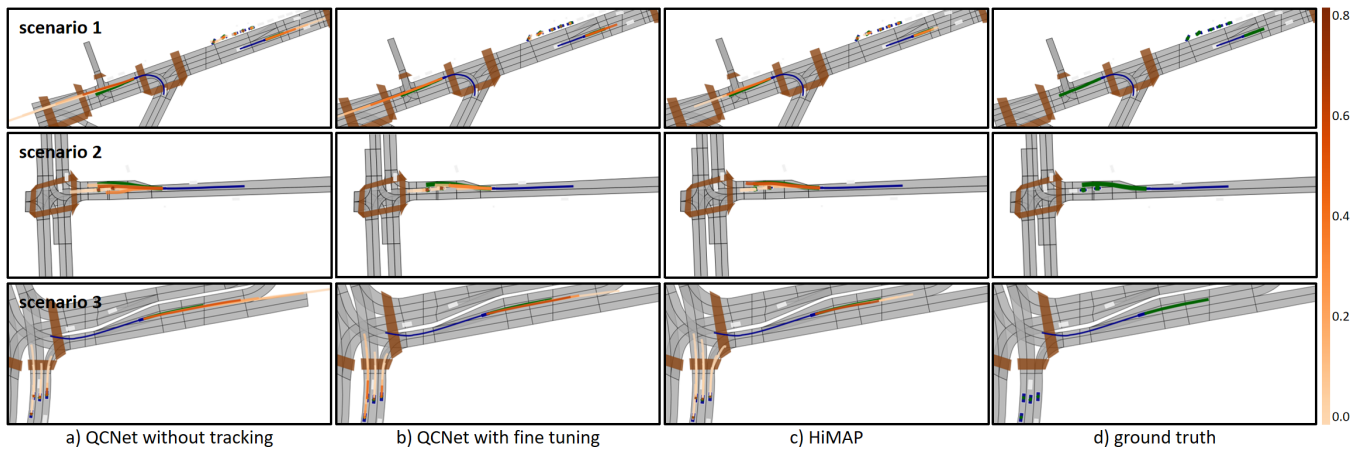


Fig. 4: Qualitative comparison of trajectory prediction results on the Argoverse 2 validation set. The blue box denotes the target agent. Blue lines show the agent’s past trajectory (not used during prediction), green lines represent the ground-truth future trajectory, and orange lines indicate predicted trajectories, with darker colors corresponding to higher predicted probabilities, and the colorbar is on the right.

TABLE V: Ablation studies on reconstructed historical steps T_h . Results are reported on the Argoverse 2 validation set.

reconstructed historical steps T_h	minFDE ₆ ↓	minADE ₆ ↓
10	1.37	0.78
20	1.34	0.76
30	1.33	0.76
40	1.35	0.77
50	1.36	0.77

improve performance by providing a natural filtering mechanism and topological constraints, reducing complexity and improving retrieval accuracy. The fourth row adds historical query initialization, fusing the current agent state with map features before querying to refine trajectory recovery. Finally, incorporating a temporal map via GRU in the fifth row achieves the best results, as it models temporal dynamics in occupancy maps and queries, allowing the network to capture both evolving scene dynamics and agent-specific history.

Table IV reports the ablation on decoder inputs. The first row uses only the current state and map, again equivalent to QCNet, and performs poorly in the no-tracking setting. The second row adds the historical temporal map, yielding large gains as decoded history is injected into DETR-style prediction queries (Section III). The third row introduces the updated query, where agents interact with historical map states, enriching the temporal context. The fourth row adds recurrence [3], [37], improving long-horizon accuracy. The complete design in the fourth row achieves the best performance, showing that updated queries and recurrent decoding are both crucial for modeling agent-specific dynamics.

F. Historical Timesteps Study

We conduct experiments on the number of reconstructed historical timesteps to determine the most suitable setting. Table V shows the impact of different choices. Using 30 timesteps achieves the best performance. Fewer timesteps may not provide sufficient historical context for accurate forecasting, while too many introduce noise and instability. In the Argoverse 2 dataset, many agents have fewer than 30 valid past observations, reflecting the practical challenge

that real-world vehicles often lack long historical trajectories, which can further increase prediction noise.

G. Qualitative Results

Figure 4 presents qualitative results comparing the baselines with our approach. Our method serves as a complementary fallback when tracking is unavailable, effectively reconstructing historical trajectories to ensure safer prediction. This reconstruction enables more accurate adjustment of agent speed and improved awareness of surrounding agents. In contrast, the original QCNet without tracking relies only on the current velocity, producing six map-constrained but overly simplistic predictions. After fine-tuning, QCNet can partially adapt its speed to scene context, yet its perception of motion dynamics and interactions remains limited. By incorporating reconstructed historical features, our approach provides robust and reliable forecasting even under tracking failure, thereby strengthening safety guarantees for autonomous driving.

V. CONCLUSION

In this work, we introduced HiMAP, a trajectory prediction framework designed to operate robustly under tracking failures. Unlike existing approaches that rely on stable tracking identities, HiMAP reconstructs historical trajectories from unlabeled detections by leveraging spatiotemporally invariant occupancy maps and a historical query mechanism. Our design enables accurate and reliable forecasting even when tracking modules fail. Experiments on Argoverse 2 show that HiMAP attains performance close to state-of-the-art tracking-based methods and substantially outperforms baselines in the no-tracking setting. These results highlight the importance of forecasting under imperfect perception and demonstrate HiMAP as a practical step toward safer and more reliable autonomy.

REFERENCES

- [1] Y. Hu, J. Yang, L. Chen, K. Li, C. Sima, X. Zhu, S. Chai, S. Du, T. Lin, W. Wang *et al.*, "Planning-oriented autonomous driving," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 17 853–17 862.
- [2] T. Ye, W. Jing, C. Hu, S. Huang, L. Gao, F. Li, J. Wang, K. Guo, W. Xiao, W. Mao *et al.*, "Fusionad: Multi-modality fusion for prediction and planning tasks of autonomous driving," *arXiv preprint arXiv:2308.01006*, 2023.
- [3] Z. Zhou, J. Wang, Y.-H. Li, and Y.-K. Huang, "Query-centric trajectory prediction," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 17 863–17 873.
- [4] Z. Li, W. Wang, H. Li, E. Xie, C. Sima, T. Lu, Q. Yu, and J. Dai, "Bev-former: learning bird's-eye-view representation from lidar-camera via spatiotemporal transformers," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [5] S. Wang, Y. Liu, T. Wang, Y. Li, and X. Zhang, "Exploring object-centric temporal modeling for efficient multi-view 3d object detection," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 3621–3631.
- [6] M. Feng and J. Su, "Rgbt tracking: A comprehensive review," *Information Fusion*, vol. 110, p. 102492, 2024.
- [7] Y. Zhang, T. Wang, and X. Zhang, "Motrv2: Bootstrapping end-to-end multi-object tracking by pretrained object detectors," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 22 056–22 065.
- [8] Z. Liu, X. Wang, C. Wang, W. Liu, and X. Bai, "Sparsetrack: Multi-object tracking by performing scene decomposition based on pseudo-depth," *IEEE Transactions on Circuits and Systems for Video Technology*, 2025.
- [9] X. Weng, J. Wang, D. Held, and K. Kitani, "Ab3dmot: A baseline for 3d multi-object tracking and new evaluation metrics," *arXiv preprint arXiv:2008.08063*, 2020.
- [10] T. Yin, X. Zhou, and P. Krahenbuhl, "Center-based 3d object detection and tracking," in *CVPR*, 2021.
- [11] Y. Zhang, P. Sun, Y. Jiang, D. Yu, Z. Yuan, P. Luo, W. Liu, and X. Wang, "Bytetrack: Multi-object tracking by associating every detection box," in *ECCV*, 2022.
- [12] N. Aharon, R. Orfaig, and B.-Z. Bobrovsky, "Bot-sort: Robust associations multi-pedestrian tracking," *arXiv preprint arXiv:2206.14651*, 2022.
- [13] 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.
- [14] X. Gao, X. Jia, Y. Li, and H. Xiong, "Dynamic scenario representation learning for motion forecasting with heterogeneous graph convolutional recurrent networks," *IEEE Robotics and Automation Letters*, vol. 8, no. 5, pp. 2946–2953, 2023.
- [15] X. Wang, T. Su, F. Da, and X. Yang, "Prophet: Efficient agent-centric motion forecasting with anchor-informed proposals," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 21 995–22 003.
- [16] J.-C. Chen, P.-S. Chang, and Y.-M. Huang, "A study on the features for multi-target dual-camera tracking and re-identification in a comparatively small environment," *Electronics*, vol. 14, no. 10, p. 1984, 2025.
- [17] X. Weng, B. Ivanovic, K. Kitani, and M. Pavone, "Whose track is it anyway? improving robustness to tracking errors with affinity-based trajectory prediction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 6573–6582.
- [18] P. Zhang, L. Bai, Y. Wang, J. Fang, J. Xue, N. Zheng, and W. Ouyang, "Towards trajectory forecasting from detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 10, pp. 12 550–12 561, 2023.
- [19] M. Liu, H. Cheng, and M. Y. Yang, "Tracing the influence of predecessors on trajectory prediction," in *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshop*, 2023, pp. 3253–3263.
- [20] S. Pellegrini, A. Ess, K. Schindler, and L. Van Gool, "You'll never walk alone: Modeling social behavior for multi-target tracking," in *2009 IEEE 12th international conference on computer vision*. IEEE, 2009, pp. 261–268.
- [21] Y. Chai, B. Sapp, M. Bansal, and D. Anguelov, "Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction," *arXiv preprint arXiv:1910.05449*, 2019.
- [22] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *European Conference on Computer Vision*. Springer, 2020, pp. 683–700.
- [23] J. Gao, C. Sun, H. Zhao, Y. Shen, D. Anguelov, C. Li, and C. Schmid, "Vectornet: Encoding hd maps and agent dynamics from vectorized representation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 525–11 533.
- [24] J. Gu, C. Sun, and H. Zhao, "Densett: End-to-end trajectory prediction from dense goal sets," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 15 303–15 312.
- [25] Z. Zhou, L. Ye, J. Wang, K. Wu, and K. Lu, "Hivt: Hierarchical vector transformer for multi-agent motion prediction," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 8823–8833.
- [26] M. Liu, H. Cheng, L. Chen, H. Broszio, J. Li, R. Zhao, M. Sester, and M. Y. Yang, "Laformer: Trajectory prediction for autonomous driving with lane-aware scene constraints," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshop*, 2024, pp. 2039–2049.
- [27] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, "The graph neural network model," *IEEE transactions on neural networks*, vol. 20, no. 1, pp. 61–80, 2008.
- [28] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [29] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social gan: Socially acceptable trajectories with generative adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2255–2264.
- [30] P. Xu, J.-B. Hayet, and I. Karamouzas, "Socialvae: Human trajectory prediction using timewise latents," in *European Conference on Computer Vision*. Springer, 2022, pp. 511–528.
- [31] Y. Xu, H. Cheng, and M. Sester, "Controllable diverse sampling for diffusion based motion behavior forecasting," in *2024 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2024, pp. 2397–2404.
- [32] Q. Yan, B. Zhang, Y. Zhang, D. Yang, J. White, D. Chen, J. Liu, L. Liu, B. Zhuang, S. Shi *et al.*, "Trajflow: Multi-modal motion prediction via flow matching," *arXiv preprint arXiv:2506.08541*, 2025.
- [33] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-end object detection with transformers," in *European conference on computer vision*. Springer, 2020, pp. 213–229.
- [34] D. Park, H. Ryu, Y. Yang, J. Cho, J. Kim, and K.-J. Yoon, "Leveraging future relationship reasoning for vehicle trajectory prediction," *arXiv preprint arXiv:2305.14715*, 2023.
- [35] X. Jia, P. Wu, L. Chen, Y. Liu, H. Li, and J. Yan, "Hdgt: Heterogeneous driving graph transformer for multi-agent trajectory prediction via scene encoding," *IEEE transactions on pattern analysis and machine intelligence*, vol. 45, no. 11, pp. 13 860–13 875, 2023.
- [36] Z. Zhang, A. Liniger, C. Sakaridis, F. Yu, and L. V. Gool, "Real-time motion prediction via heterogeneous polyline transformer with relative pose encoding," *Advances in Neural Information Processing Systems*, vol. 36, pp. 57 481–57 499, 2023.
- [37] Y. Zhou, H. Shao, L. Wang, S. L. Waslander, H. Li, and Y. Liu, "Smartrefine: A scenario-adaptive refinement framework for efficient motion prediction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 15 281–15 290.
- [38] B. Zhang, N. Song, and L. Zhang, "Decoupling motion forecasting into directional intentions and dynamic states," *Advances in Neural Information Processing Systems*, vol. 37, pp. 106 582–106 606, 2024.
- [39] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. K. Pontes *et al.*, "Argoverse 2: Next generation datasets for self-driving perception and forecasting," *arXiv preprint arXiv:2301.00493*, 2023.
- [40] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.
- [41] L. Ilya and H. Frank, "Sgdr: Stochastic gradient descent with warm restarts," *arXiv preprint arXiv:1608.03983*, 2016.