

MASAR: Motion–Appearance Synergy Refinement for Joint Detection and Trajectory Forecasting

Mohammed Amine Bencheikh Lehocine¹, Julian Schmidt¹, Frank Moosmann¹,
 Dikshant Gupta¹, and Fabian Flohr²

Abstract— Classical autonomous driving systems connect perception and prediction modules via hand-crafted bounding-box interfaces, limiting information flow and propagating errors to downstream tasks. Recent research aims to develop end-to-end models that jointly address perception and prediction; however, they often fail to fully exploit the synergy between appearance and motion cues, relying mainly on short-term visual features. We follow the idea of “looking backward to look forward”, and propose MASAR, a novel fully differentiable framework for joint 3D detection and trajectory forecasting compatible with any transformer-based 3D detector. MASAR employs an object-centric spatio-temporal mechanism that jointly encodes appearance and motion features. By predicting past trajectories with guidance from appearance cues, MASAR captures long-term temporal dependencies that enhance future trajectory forecasting. Experiments conducted on the nuScenes dataset demonstrate MASAR’s effectiveness, showing improvements of over 20% in minADE and minFDE while maintaining robust detection performance. Code and models are available at <https://github.com/aminmed/MASAR>.

I. INTRODUCTION

A fundamental requirement for safe and reliable autonomous driving is the ability to both perceive the surrounding scene and anticipate the future behaviors and interactions of nearby agents [1]. Vision-based systems have gained increasing attention due to their lower deployment cost, scalability, and rich semantic information. However, camera setups remain highly sensitive to localization errors caused by depth ambiguity and occlusions, which degrade detection performance and, consequently, future prediction.

To mitigate these issues, prior works [2], [3], [4], [5], [6] leverage historical frames to enhance the robustness of learned scene representations. These methods primarily augment appearance features (high level visual features describing how objects look) but do not model long-term past object dynamics. Bird’s-eye-view-based (BEV-based) approaches, for instance, compensate for ego-motion when aligning and aggregating past frames [7], [2], [6], [8] but ignore individual object motion. Perspective-based (sparse query-based) models, on the other hand, either focus only on modeling short-term inter-frame object motion [4], [9], [10] or rely on constant-velocity assumptions for temporal

This work is a result of the joint research project STADT:up (19A22006O). The project is supported by the German Federal Ministry for Economic Affairs and Energy (BMWE), based on a decision of the German Bundestag. The author is solely responsible for the content of this publication.

¹Mercedes-Benz AG, Germany.

²Munich University of Applied Sciences, Germany.

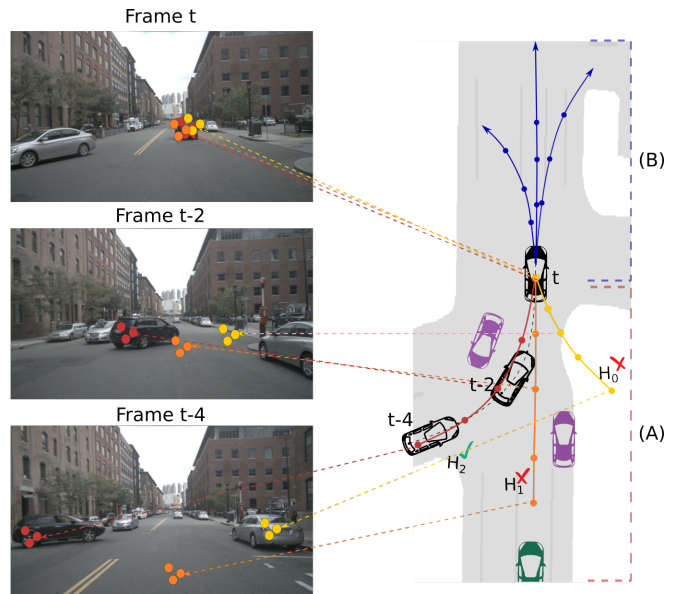


Fig. 1: Core idea of our method. *Left*: multi-frame input images, *Right*: (A) for each object query (i.e. hypothesis), taking the black car as example, we iteratively refine past trajectory hypotheses (yellow, orange, red), aggregate visual features along them, and perform appearance-guided scoring. H_2 wins because it hits more visual features corresponding to the object. (B) Based on the selected past trajectory and aggregated features, multiple modes of future trajectories (blue) are predicted.

features sampling [3], [5] thereby often overlooking long-term motion patterns that can potentially improve both detection and forecasting.

In modular autonomous driving pipelines, multi-object tracking is used to construct past trajectories for future prediction. Camera-based tracking is inherently noisy [11], [1], prior works [12], [9] reported that explicitly including noisy tracked trajectories leads to degradation in forecasting performance. Recent end-to-end approaches instead propagate tracking queries as input for forecasting [12], [13], [14], [9], but these queries primarily encode re-identification features and do not completely replace the accurate past trajectories.

To overcome these limitations, we bypass noise-prone tracking by directly predicting smooth past trajectories and introduce a new approach that refines past motion while modeling temporal object features. Our framework consists

of two key components: the Appearance-guided Past Motion Refinement (APR), which predicts multiple candidate past trajectories and selects the most compatible one using appearance cues, and the Past-conditioned Forecasting Decoder (PFD), which leverages these refined trajectories to improve future prediction (Figure 1). Unlike prior works [9], [13], [10] that rely solely on tracking queries to encode object dynamics, PFD explicitly incorporates past trajectories alongside object queries. Notably, MASAR outperforms previous works on end-to-end forecasting without tracking or any map information.

In summary, our contributions are:

- We propose MASAR, a tracking-free, map-free framework for joint 3D detection and trajectory forecasting, compatible with any transformer-based detector. To demonstrate its adaptability, we integrate it into two leading architectures: BEVFormer [2] and SparseBEV [3]¹, achieving consistent improvements.
- MASAR sets a new state-of-the-art performance for end-to-end forecasting on the nuScenes dataset [15], reducing minADE and minFDE by over 20% without relying on map information.
- Through extensive ablations, we show that conditioning on refined past trajectories provides up to 6% improvement in minFDE and 7% reduction in miss rate, highlighting the critical impact of our past-motion modeling approach.

II. RELATED WORKS

A. Camera-based 3D Object Detection

Early approaches address multi-view 3D detection by first applying standard 2D methods and then transform the 2D boxes to 3D by regressing additional 3D attributes [16]. LSS [17] introduces learning holistic BEV representations via view transformation, either by predicting per-pixel depth distributions to lift 2D features into 3D or BEV space [18], [7], [19], or by projecting BEV pillars onto the 2D image plane for features’ sampling [2], [20]. DETR3D [21] leverages sparse object queries with 3D reference points to sample image features. PETR [22] and SpatialDETR [23] further incorporate 3D geometric priors into image features, while StreamPETR [4] extends this design to a streaming setting.

B. Spatio-temporal Modeling

Spatio-temporal modeling is crucial for achieving high performance in camera-based 3D detection. Depending on the intermediate scene representation, two main directions can be distinguished:

Scene-level Modeling: Often adopted by BEV-based models. Methods such as BEVDet4D [7], FIERY [19], and BEVerse [24] wrap and fuse temporal BEVs using ego-motion compensation and convolutions, while BEVFormer [2] employs deformable attention to recurrently attend to previous BEVs. VideoBEV [6] uses a hybrid of

parallel and recurrent fusion but still relies on wrapping and convolutions. These methods, however, can lose important temporal information due to restricted BEV range and induce feature distortion [20] because they do not account for individual object motion.

Object-level Modeling: For perspective-based models, SparseBEV [3] and Sparse4D [5] use a constant velocity model to generate past trajectories for sampling features from previous perspective frames. Inspired by the tracking-by-attention mechanism [25], StreamPETR [4] and Sparse4Dv3 [26] adopt query propagation for temporal modeling. Additionally, StreamPETR [4] and [9], [10] incorporate inter-frame ego and object motion using a dedicated motion layer normalization.

A key limitation of prior works is that spatio-temporal modeling focuses on aggregating visual information from past frames without explicitly building up an understanding of objects’ past motion. Our APR introduces a novel approach that jointly predicts and refines both long-term past motion and appearance features.

C. Conventional Trajectory Forecasting

Given HD maps and past agent trajectories, the goal is to predict multiple future trajectories for surrounding agents. Early works rely on rasterized scene representations [27], [28], later [29] introduces vectorized representation to better handle heterogeneous inputs. Recent transformer-based methods [30], [31], [32] achieve state-of-the-art performance and solve challenging scenarios across multiple datasets. However, the assumption of access to curated past trajectories is often unrealistic in real-world settings [1], thereby limiting the applicability of these methods in practical autonomous driving scenarios.

D. Joint Detection and Forecasting

Detection and forecasting directly from sensor inputs has recently gained significant attention. Early LIDAR-based approaches employ convolutional networks to jointly detect, track, and forecast motion [33], [34]. More recent methods, such as FutureDet [35], and DeTra [36] address the joint detection and forecasting problem as a unified, single task. In camera-based systems, ViP3D [12] leverages HD maps with tracking queries to capture object dynamics and semantics for forecasting, while follow-up works [13], [10], [14] extend this framework to planning with additional online map construction. More recently, Foresight [9] builds on StreamPETR [4] and uses forward and backward query propagation to enhance both detection and forecasting performance. In contrast, MASAR eliminates the need for tracking by constructing smooth past trajectories using appearance guidance and removes the reliance on map information by leveraging joint optimization of both detection and forecasting tasks. We show through experiments substantial improvements over these SOTA models.

III. METHOD

In the following sections, we describe our method for BEV-based models. Subsection III-E then details how the

¹Despite its name, SparseBEV does not construct a BEV representation; its decoder directly interacts with perspective multi-scale image features.

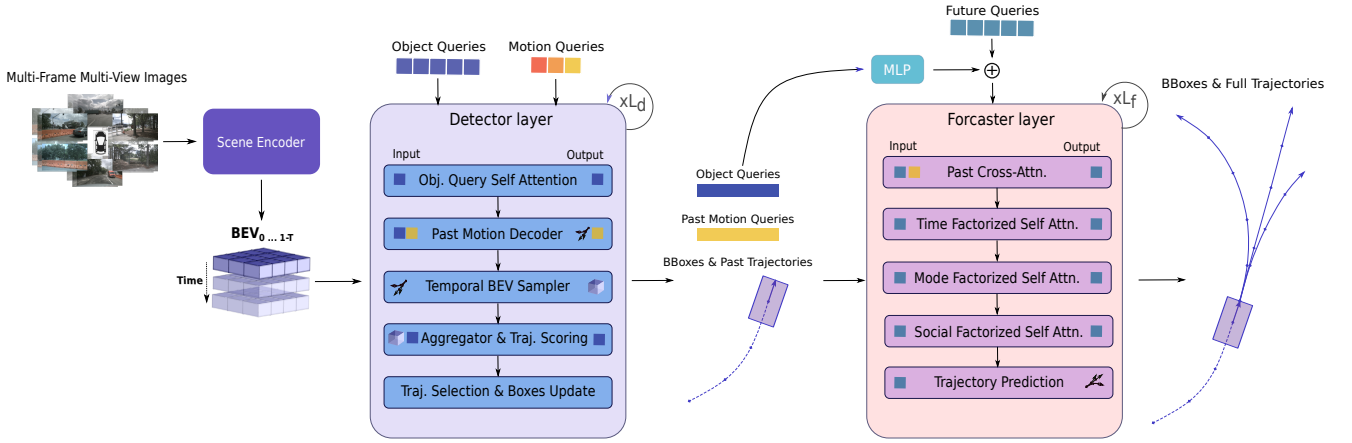


Fig. 2: MASAR architecture: The scene encoder encodes all multi-frame, multi-view images into BEVs. *Detector decoder*: L_d layers iteratively refine object detections together with their past trajectory estimates. *Forecasting decoder*: L_f layers iteratively forecast multiple trajectory modes for each detected object based on its estimated past trajectory and its visual features along that trajectory.

same approach can be adapted for perspective-based models.

A. Overall Architecture

MASAR extends any transformer-based 3D object detector (DETR-like) into a joint object detector and forecaster. The APR module estimates smooth past trajectories for detected objects without explicit tracking, while the forecasting decoder leverages both object features and past trajectories to predict multi-modal future trajectories, capturing the long-term behavior of surrounding objects. An overview of the framework is shown in Figure 2.

B. Scene Encoder

Given multi-frame, multi-view images, we use an image backbone [37] and a feature pyramid network (FPN) [38] to extract multi-scale multi-view feature maps \mathcal{F}_t for each time step t in parallel, with $-T_h < t \leq 0$, where T_h denotes the history length. For BEV-based detectors, the temporal multi-scale feature maps are fed to a BEV encoder to transform them into temporal BEV features \mathbf{B}_t :

$$\mathbf{B}_t = \text{BEVEncoder}(\mathcal{F}_t), \quad \mathbf{B}_t \in \mathbb{R}^{H_{\text{bev}} \times W_{\text{bev}} \times D} \quad (1)$$

where H_{bev} and W_{bev} denote the spatial dimensions of the BEV grid, and D is the feature dimension.

In our experiments, The BEVEncoder is a modification of the encoder proposed in [2], composed of self-attention and spatial cross-attention layers. We remove the original auto-regressive temporal modeling and replace it with our object-centric spatio-temporal mechanism APR.

C. Detector Decoder: Appearance-guided Past Motion Refinement (APR)

The detector is a transformer-based decoder [2], [3], [21] composed of L_d refinement layers as shown in Figure 2. It iteratively estimates how objects have moved over the past few seconds and leverages these trajectories to sample and aggregate corresponding features from previous frames [5],

[3]. We formulate past trajectory generation as a candidate selection process: for each object, multiple trajectories are generated, and the most plausible one is selected based on a compatibility score computed from aggregated appearance features, which indicates how closely a trajectory matches the object’s true past motion, see Figure 1.

Algorithm 1 illustrates the iterative refinement process. With N being the number of object queries, M_h the number of past trajectory hypotheses, and D the feature dimensionality, we denote object queries as $\mathbf{q}_\ell^{\text{obj}} \in \mathbb{R}^{N \times D}$, motion queries as $\mathbf{q}_\ell^{\text{mo}} \in \mathbb{R}^{N \times D}$, multi-hypothesis motion queries as $\mathbf{Q}_\ell^{\text{mo}} \in \mathbb{R}^{N \times M_h \times D}$, past trajectory proposals as $\mathbf{P}_\ell \in \mathbb{R}^{N \times T_h \times 2}$, and ego-motion-based transformation matrices $\{\mathbf{T}_{0 \rightarrow t}\}_{t=-1}^{1-T_h}$ for mapping positions from the current ego frame to past ego frames.

Initialization: We initialize object queries $\mathbf{q}_0^{\text{obj}}$ following standard DETR-based 3D detection practices [21], [3]. Trajectory proposals \mathbf{P}_0 are initialized using a constant velocity model. Motion queries \mathbf{q}_0^{mo} are derived from object queries through an MLP_{init} (line 1).

Past Motion Decoder (lines 3-4): We extend $\mathbf{q}_\ell^{\text{mo}}$ with a set of learnable embeddings $\mathbf{E}_{\text{past}} \in \mathbb{R}^{1 \times M_h \times D}$ to form trajectory hypothesis queries $\mathbf{Q}_\ell^{\text{mo}}$. \mathbf{E}_{past} acts as latent-space offsets, enabling the model to explore diverse variations of motion patterns efficiently. We then predict multiple candidate past trajectories per object using $\mathbf{Q}_\ell^{\text{mo}}$ and trajectory proposals \mathbf{P}_ℓ . The MotionDecoder is similar to a PFD layer (Section III-D) but employs only two factorized attention blocks: one over modes and one over time. Its output $\mathbf{H}_\ell \in \mathbb{R}^{N \times M_h \times T_h \times 4}$ encodes the means and scales of an isotropic Laplacian distributions for each time step and hypothesis, all expressed in the current $t = 0$ ego-frame.

Temporal BEV Sampler (lines 5-6): We use multi-scale deformable attention [39] to sample features from temporal BEVs along predicted past trajectory hypotheses. For each object and attention head k , a linear layer predicts N_{off} sampling offsets $\Delta(x_i^k, y_i^k)$, scaled by the object bounding

Algorithm 1 Appearance-guided Past Motion Refinement

Require: Temporal BEV \mathbf{B} , initial object queries $\mathbf{q}_0^{\text{obj}}$, motion emb. \mathbf{E}_{past} , trajectory proposals \mathbf{P}_0 , and transforms $\{\mathbf{T}_{0 \rightarrow t}\}_{t=-1}^{1-T_h}$.

Ensure: Object queries \mathbf{q}^{obj} , motion queries \mathbf{q}^{mo} , and past trajectory proposals \mathbf{P}

- 1: $\mathbf{q}_0^{\text{mo}} \leftarrow \text{MLP}_{\text{init}}(\mathbf{q}_0^{\text{obj}})$
- 2: **for** $\ell = 0$ to $L_d - 1$ **do**
- 3: $\mathbf{Q}_\ell^{\text{mo}} \leftarrow \mathbf{q}_\ell^{\text{mo}} + \mathbf{E}_{\text{past}}$
- 4: $\mathbf{H}_\ell \leftarrow \text{MotionDecoder}(\mathbf{Q}_\ell^{\text{mo}}, \mathbf{P}_\ell)$
- 5: $\hat{\mathbf{H}}_\ell \leftarrow \text{Concat}\{\mathbf{T}_{0 \rightarrow t} \cdot \mathbf{H}_\ell[:, :, t]\}_{t=-1}^{1-T_h}$
- 6: $\mathbf{F}_\ell \leftarrow \text{BEVSampler}(\mathbf{B}, \hat{\mathbf{H}}_\ell)$
- 7: $\mathbf{F}_\ell^{\text{traj}} \leftarrow \text{Aggregate}(\mathbf{F}_\ell, \mathbf{q}_\ell^{\text{obj}})$
- 8: $\mathbf{s}_\ell \leftarrow \text{MLP}_{\text{score}}(\mathbf{F}_\ell^{\text{traj}})$
- 9: $\mathbf{q}_{\ell+1}^{\text{obj}} \leftarrow \mathbf{q}_\ell^{\text{obj}} + \sum_m \text{Softmax}(\mathbf{s}_\ell)_m \cdot \mathbf{F}_\ell^{\text{traj}}[:, m]$
- 10: $m^* \leftarrow \arg \max_m \mathbf{s}_\ell$
- 11: $\mathbf{q}_{\ell+1}^{\text{mo}} \leftarrow \mathbf{Q}_\ell^{\text{mo}}[:, m^*] + \text{MLP}_{\text{update}}(\mathbf{q}_{\ell+1}^{\text{obj}})$
- 12: $\mathbf{P}_{\ell+1} \leftarrow \mathbf{H}_\ell[:, m^*, :, : 2]$
- 13: **end for**
- 14: **return** $\mathbf{q}^{\text{obj}}, \mathbf{q}^{\text{mo}}, \mathbf{P}$

box dimensions (w, l) similar to [5], [3]. Another linear layer predicts attention weights w_i^k from object queries $\mathbf{q}_\ell^{\text{obj}}$. Sampled features for head k are computed as:

$$\mathbf{f}_t^k = \sum_{i=1}^{N_{\text{off}}} w_i^k \text{Bilinear}(\mathbf{B}_t, \hat{\mathbf{H}}_\ell(t) + \Delta(x_i^k, y_i^k)), \quad (2)$$

where $\hat{\mathbf{H}}_\ell$ denotes transformed trajectory hypotheses \mathbf{H}_ℓ to each previous frame ego coordinate system. Sampled features $\mathbf{F}_\ell \in \mathbb{R}^{N \times M_h \times T_h \times D}$ are obtained by concatenating \mathbf{f}_t^k along the feature dimension and stacking across time.

Features Aggregator (lines 7-9): To aggregate the sampled features \mathbf{F}_ℓ of each object from temporal BEVs, we adopt adaptive mixing following [3], [40], that we found provides both superior performance and greater flexibility compared to conventional temporal fusion strategies such as simple summation or recursive aggregation like in [5]. Subsequently, an $\text{MLP}_{\text{score}}$ predicts a compatibility score $\mathbf{s}_\ell^{\text{traj}}$ based on the aggregated trajectory appearance features $\mathbf{F}_\ell^{\text{traj}}$. To update the object queries, we compute a weighted sum of the candidates' features $\mathbf{F}_\ell^{\text{traj}}$, using the softmax-normalized scores as weights.

Iterative Refinement (lines 10-12): The trajectory proposals $\mathbf{P}_{\ell+1}$ and motion queries $\mathbf{q}_{\ell+1}^{\text{mo}}$ for the next decoder layer $\ell + 1$ are obtained by selecting the highest-scoring hypothesis.

D. Past-conditioned Forecasting Decoder (PFD)

PFD is a transformer decoder where each block is composed of three factorized attention layers [32], [36], [30] with a past-conditioning layer using cross-attention as illustrated in Figure 2. PFD refines a set of future queries \mathbf{Q}^{fut} , used to predict multi-modal future trajectories in a tracking-free and map-free setting. PFD is positioned hierarchically on top of a

transformer-based object decoder forming a unified pipeline. The joint architecture is fully differentiable and trained end-to-end, enabling seamless optimization across both detection and forecasting tasks.

In contrast to prior works [9], [36], [41], [13], PFD does not interact with any scene-level representation (e.g. BEV features or multi-view feature maps). We posit that object queries \mathbf{q}^{obj} encode essential object semantics, while motion features \mathbf{q}^{mo} and explicit past trajectories \mathbf{P} provide sufficient information to capture object's dynamics.

Query Initialization: We employ two sets of learnable parameters: $\mathbf{e}_t^{\text{fut}} \in \mathbb{R}^d$, $1 \leq t < T_f$ for time steps embeddings and $\mathbf{e}_m^{\text{mode}} \in \mathbb{R}^d$, $1 \leq m < M_f$ for different future modes M_f [36], [32]. These embeddings are expanded to construct a future query volume $\mathbf{Q}^{\text{fut}} \in \mathbb{R}^{N \times M_f \times T_f \times D}$.

To make the future queries aware of the semantic features of each object, we incorporate information from the object queries \mathbf{q}^{obj} using an MLP:

$$\mathbf{Q}^{\text{fut}} \leftarrow \text{LayerNorm}(\mathbf{Q}^{\text{fut}} + \text{MLP}(\mathbf{q}^{\text{obj}})) \quad (3)$$

Attention Layers: We employ four types of attention layers in our forecaster: past cross-attention, factorized temporal attention, factorized mode attention, and factorized social attention.

For past-conditioning, the future queries \mathbf{Q}^{fut} attend to past motion queries \mathbf{q}^{mo} using cross-attention:

$$\begin{aligned} \tilde{\mathbf{Q}}^{\text{fut}} &\leftarrow \text{CrossAttn}(\mathbf{Q}^{\text{fut}}, \mathbf{q}^{\text{mo}} + \text{PE}(\mathbf{P})), \\ \mathbf{Q}^{\text{fut}} &\leftarrow \text{LayerNorm}(\mathbf{Q}^{\text{fut}} + \tilde{\mathbf{Q}}^{\text{fut}}), \end{aligned} \quad (4)$$

where $\text{PE}(\cdot)$ denotes sinusoidal positional encoding.

Factorized attention is widely adopted in motion prediction [14], [32], [30], as it is computationally efficient and captures rich dependencies across temporal, mode, and social dimensions in the query volume [36]. For each factorized attention block, we first permute and reshape the query volume \mathbf{Q}^{fut} so that the target factorized dimension (e.g., temporal T_f) is treated as the query dimension, while the remaining dimensions are flattened with the batch dimension [36]. For instance, in the case of factorized temporal attention:

$$\mathbf{Q}^{\text{fut}} \in \mathbb{R}^{N \times M_f \times T_f \times D} \mapsto \tilde{\mathbf{Q}}^{\text{fut}} \in \mathbb{R}^{(N \cdot M_f) \times T_f \times D} \quad (5)$$

Self attention is then applied along T_f , and the result is reshaped back to recover the original query volume:

$$\tilde{\mathbf{Q}}^{\text{fut}} \mapsto \mathbf{Q}^{\text{fut}} \in \mathbb{R}^{N \times M_f \times T_f \times D} \quad (6)$$

Each attention layer is followed by a residual connection and layer normalization, and then passed by a position-wise feed-forward network [42].

Multi-Modal Futures Regression and Scoring: At the end of each forecaster layer, we employ a regression head that for each object i and each future mode m , predicts the future location means $\mu_t^{i,m} \in \mathbb{R}^2$ and scales $\sigma_t^{i,m} \in \mathbb{R}^2$ of an isotropic Laplacian distribution [44]. We use another head to predict the scores associated with each future mode.

Method	Backbone	Map		3D Detection Metrics		Forecasting Metrics ($k = 6$)			
		Offline	Online	mAP \uparrow	NDS \uparrow	EPA \uparrow	minADE \downarrow	minFDE \downarrow	MR \downarrow
PnPNet [13]	R101	✓		–	–	0.222	1.15	1.95	0.226
PIP [43]	R50		✓	28.0	–	0.258	1.23	1.75	0.195
UniAD [13]	R101		✓	38.0	49.8	0.456	0.71	1.02	0.151
BEVFormer [2] + DeTra [36] †	R101			41.4	–	0.504	0.61	1.00	0.114
BEVFormer [2] + Ours	R101			43.0	52.9	0.519	0.55	0.82	0.101
Traditional [12]	R50	✓		–	–	0.209	2.06	3.02	0.277
ViP3D [12]	R50	✓		–	–	0.226	2.05	2.84	0.246
SparseDrive [14]	R50		✓	41.8	52.5	0.482	0.62	0.99	0.136
ForeSight [9]	R50	✓		46.6	56.0	0.499	0.70	–	–
SparseBEV [3] + Ours	R50			<u>43.5</u>	<u>53.8</u>	<u>0.492</u>	0.51	0.77	0.093
SparseDrive [14]	R101		✓	49.6	58.8	0.555	0.60	0.96	0.132
ForeSight [9]	R101	✓		50.2	58.9	0.549	0.68	0.93	0.102
SparseBEV [3] + Ours	R101			48.8	57.6	0.544	0.46	0.72	0.086

TABLE I: Comparison of BEV-based and perspective-based methods on nuScenes joint detection and forecasting tasks. † denotes DeTra’s refinement transformer integrated with BEVFormer (without map attention).

E. Adaptation to Perspective-based Models

For perspective-based models such as SparseBEV [3], we remove the BEV encoder and allow the detector decoder to directly interact with temporal multi-view features \mathcal{F}_t . We assume that object motion occurs only in the BEV plane [3], keeping the same z coordinate along the past trajectory. For sampling, we use the same mechanism introduced in [3]. PFD remains unchanged.

IV. EXPERIMENTS

In this section, we detail the dataset, evaluation metrics, and implementation specifics used to assess MASAR’s performance on joint detection and forecasting.

A. Datasets

To compare against most prior works on vision-based joint detection and forecasting, we train and evaluate our framework on nuScenes, a large-scale autonomous driving dataset [15]. It contains 1000 driving scenes, each lasting approximately 20 seconds. The dataset includes annotations for 3D detection and tracking at 2 Hz, which we use to construct ground-truth trajectories for end-to-end forecasting [13]. All results are reported on the validation split.

B. Evaluation Metrics

For 3D detection, we follow the official nuScenes 3D detection benchmark [15] and report main metrics, including mean Average Precision (mAP) and NuScenes detection score (NDS). For trajectory forecasting, we follow the evaluation protocol from [13], reporting metrics for $T_f = 12$ time steps, and over $M_f = 6$ predicted future modes, including minADE, minFDE, and Miss Rate. These metrics are computed on matched objects at 2.0 meters threshold and averaged similarly to TP metrics in detection. Additionally, we report the End-to-End Prediction Accuracy (EPA) [12], which accounts for both false positive detections and successful forecasts (i.e., true positive detections with an FDE less than 2 meters).

C. Implementation Details

As highlighted in Section III, our framework is flexible and can be integrated with both BEV-based and perspective-based object detectors. Accordingly, our experiments build upon two models: **BEVFormer** [2] a BEV-based 3D detector and **SparseBEV** [3] a perspective-based 3D detector. Unless otherwise specified, we follow the official training settings for each model, including image backbone, input resolution, number of object queries, optimizer, and detection loss \mathcal{L}_{det} .

The overall training objective for the joint detection, past refinement and forecasting is:

$$\mathcal{L} = \mathcal{L}_{\text{det}} + \lambda_p \cdot \mathcal{L}_{\text{past}} + \lambda_f \cdot \mathcal{L}_{\text{future}} \quad (7)$$

Our newly introduced trajectory-related losses $\mathcal{L}_{\text{past}}$ and $\mathcal{L}_{\text{future}}$ consist each of two components: (i) a regression loss, computed as the negative log-likelihood of an isotropic Laplacian distribution [44], and (ii) a scoring loss. For past refinement, the scoring loss is formulated as a binary cross-entropy loss, where targets are derived by scaling the negative average displacement error into $[0, 1]$ via a sigmoid function. For future prediction, the scoring loss is supervised using cross-entropy with a soft-target strategy [44]. We compute $\mathcal{L}_{\text{past}}$ and $\mathcal{L}_{\text{future}}$ only for assigned objects with center distance ≤ 1.0 m to the ground truth [36]. We adopt a teacher forcing with curriculum schedule training strategy: ground-truth past trajectories are used in place of the refined trajectories, and their ratio is gradually decreased over time.

The full model is trained end-to-end in a fully differentiable manner. We use $\lambda_p = 0.2$ and $\lambda_f = 0.1$ for weighting the past and future trajectory losses.

D. Joint Detection and Forecasting Performance

We evaluate our joint detection and forecasting framework against state-of-the-art models on the nuScenes validation split [15]. Since BEV-based methods typically fall behind perspective-based methods in 3D detection [8],

we group comparisons into two categories: BEV-based and perspective-based approaches. Within perspective-based models, we further group methods according to the image backbone used.

1) *BEV-based Models*: As shown in Table I, our BEVFormer [2] variant equipped with both APR and PFD achieves the best performance on end-to-end forecasting metrics, reducing minFDE by 17.6% and improving EPA by 12.5% compared to prior work [13]. Furthermore, our model significantly enhances BEVFormer [2] in 3D detection (3.3% improvement in mAP). Table II reports detailed forecasting metrics for cars and pedestrians. Our method improves forecasting for dynamic classes, achieving nearly 13% improvement in minADE for cars and 6% for pedestrians compared to the baseline (BEVFormer + DeTra [36]).

2) *Perspective-based Models*: Our integration with SparseBEV [3] surpasses all prior methods in motion prediction, achieving 20% lower minADE and minFDE and a 13% lower miss rate compared to [14], [9], all without using map information during training or inference. Although our model is not the top ranked model in 3D detection, it attains highly competitive EPA (within 1% of the best), demonstrating the effectiveness of our past-conditioned forecaster in modeling diverse future agent behaviors, which in turn leads to higher true positive joint detections and forecasts [12]. As shown in Table II, MASAR achieves the best performance for pedestrians, surpassing DualAD [10] by 8% in EPA and 4% in minADE, highlighting its ability to model complex agent behaviors. On an NVIDIA A100 GPU, our implementation introduces an additional 114.9 ms of inference time for the forecasting head and 60 ms for APR to the 58 ms baseline inference time of the default SparseBEV detector head.

Method	EPA \uparrow			minADE \downarrow		minFDE \downarrow	
	Car	Ped	Avg	Car	Ped	Car	Ped
UniAD [13]	0.46	0.35	0.405	0.71	0.78	1.02	1.05
BEVFormer [2] + DeTra	0.49	0.43	0.46	0.84	0.74	1.61	1.12
BEVFormer [2] + Ours	0.50	0.45	0.475	0.73	0.69	1.19	1.00
SparseDrive-S [14]	0.48	0.41	0.445	0.62	0.72	0.99	1.07
DualAD [10]	0.524	0.452	0.488	0.68	0.63	1.08	0.89
SparseBEV [3] + Ours	0.522	0.491	0.506	0.60	0.60	1.08	0.88

TABLE II: Comparison of motion prediction results on dynamic classes (cars and pedestrians). Results for ForeSight [9] are not reported, as per-class metrics are not publicly available.

E. Ablation Studies

We conduct several ablation experiments to assess the impact of each component on MASAR’s overall performance.

1) *Impact of History Length*: To examine the impact of longer historical context on joint detection and forecasting, we vary the number of past frames provided to BEVFormer-small [2]. As shown in Table III, increasing past frames to 8 (4 seconds) steadily improves detection and forecasting metrics (8.5% improvement on NDS). While it is shown that

BEVFormer’s recurrent temporal modeling saturates around 4 frames [2], [6], APR effectively leverages longer histories, yielding consistent gains.

Frames	NDS \uparrow	mAP \uparrow	EPA \uparrow	minADE \downarrow	minFDE \downarrow	MR \downarrow
3	47.8	37.0	–	–	–	–
4	50.7	40.0	0.491	0.592	0.868	0.108
6	51.4	40.8	0.498	0.587	0.862	0.105
8	51.9	41.5	0.500	0.583	0.856	0.104

TABLE III: Detection and forecasting performance with varying number of past frames on BEVFormer-small [2]. The gray row indicates the original BEVFormer-small performance taken from official BEVFormer Github repository.

2) *Past-conditioned Forecasting Decoder*: Table IV evaluates the contribution of each component in the forecasting decoder. We conduct this experiment using SparseBEV [3] with a R50 image backbone. Removing object queries from the future query initialization results in the loss of semantic and class information about the objects, while omitting the learnable time and mode weights limits the model’s ability to generate distinct queries for different future possibilities. Additionally, past conditioning significantly improves forecasting performance, confirming our first motivation of “looking backward to look forward”.

W	O	P	EPA \uparrow	minADE \downarrow	minFDE \downarrow	MR \downarrow
\times	\checkmark	\checkmark	0.432	1.064	1.786	0.194
\checkmark	\times	\checkmark	0.480	0.594	0.913	0.106
\checkmark	\checkmark	\times	0.489	0.539	0.820	0.100
\checkmark	\checkmark	\checkmark	0.492	0.511	0.770	0.093

TABLE IV: Ablation of forecasting decoder components. W: learnable weights for query initialization, O: object queries, P: past cross-attention.

3) *Past Motion Modeling*: We evaluate three past motion modeling variants using SparseBEV [3] with an R50 backbone: constant velocity, past motion prediction without appearance guidance (w/o AG), and with appearance guidance (w/ AG). As shown in Table V, past refinement with AG slightly reduces detection but improves forecasting, achieving a 6% gain in minFDE. This highlights the challenge of jointly performing detection and past predictions. While many objects in nuScenes [15] are static—making constant velocity sufficient for detection—appearance guidance better captures past object dynamics, enhancing forecasting performance while maintaining robust detection compared to not using guidance.

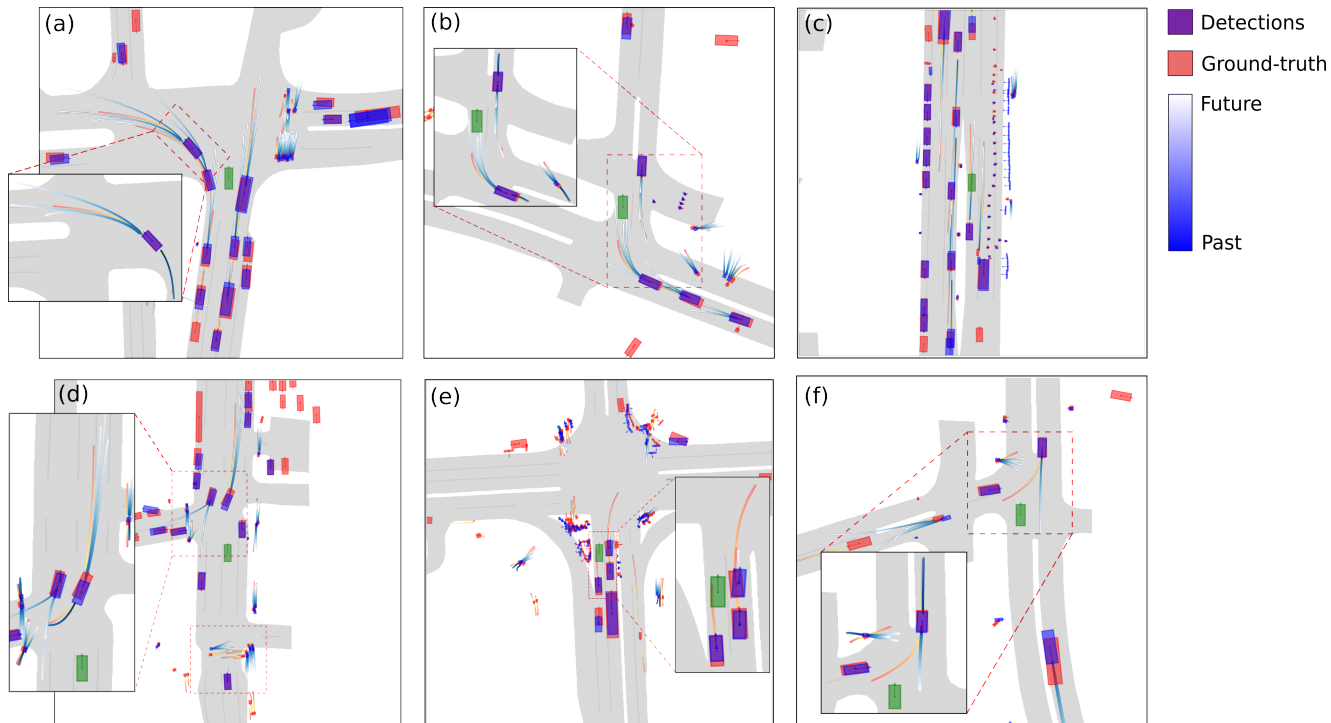


Fig. 3: Visualization of some nuScenes validation samples. Green: ego, purple: detections, red: ground-truth. Only future trajectories are plotted; past trajectories are shown in magnified views. The rendered maps are just for visualization and not input to the model. (a) and (c) show challenging crowded scenes, (b) and (d) diverse multi-modal futures, and (e) and (f) typical failure cases from missing context.

Past Motion	mAP \uparrow	EPA \uparrow	FDE $_{\text{past}}\downarrow$	minADE \downarrow	minFDE \downarrow	MR \downarrow
Const. Vel.	44.2	0.49	0.97	0.53	0.82	0.104
w/o AG	42.8	0.48	<u>0.84</u>	<u>0.52</u>	<u>0.78</u>	<u>0.095</u>
w/ AG	<u>43.5</u>	0.49	0.83	0.51	0.77	0.093

TABLE V: Ablation of past motion modeling strategies. FDE $_{\text{past}}$ denotes single hypothesis FDE computed for past motion on dynamic classes (cars and pedestrians).

F. Qualitative results

Figure 3 shows selected qualitative examples from nuScenes. Samples (a) and (c) highlight MASAR’s ability to generate diverse and consistent future trajectories for different agent types, including cars, pedestrians, motorcycles, and large vehicles. Samples (b) and (d) demonstrate MASAR’s capacity to capture plausible multi-modal futures; in (d), for instance, the model correctly predicts pedestrians’ intention to cross the street, although it fails to capture a change of direction for another group of pedestrians at the bottom of the figure. In contrast, (e) shows difficulty in handling behavior changes of stopped vehicles, while (f) illustrates a failure case where the model correctly estimates an oncoming car’s past trajectory but misses its future intention, suggesting the potential benefit of incorporating HD map context.

V. CONCLUSIONS

In this work, we introduced MASAR, a unified framework for joint 3D detection and trajectory forecasting in au-

tonomous driving. We designed a new object-centric spatio-temporal mechanism leveraging motion and appearance cues by refining past trajectories with appearance guidance. MASAR captures long-term temporal dependencies without relying on tracking or any map information. We integrated it with transformer-based detectors such as BEVFormer and SparseBEV, showing consistent improvements over prior end-to-end methods in trajectory forecasting. Extensive ablation studies demonstrate the effectiveness of each component in our framework, highlighting how past-conditioning substantially improves future forecasting.

REFERENCES

- [1] Y. Xu, L. Chambon, E. Zablocki, M. Chen, A. Alahi, M. Cord, and P. Pérez, “Towards motion forecasting with real-world perception inputs: Are end-to-end approaches competitive?” in *IEEE International Conference on Robotics and Automation*, 2024, pp. 18 428–18 435.
- [2] Z. Li, W. Wang, H. Li, E. Xie, C. Sima, T. Lu, Y. Qiao, and J. Dai, “Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers,” in *European Conference on Computer Vision*, 2022, pp. 1–18.
- [3] H. Liu, Y. Teng, T. Lu, H. Wang, and L. Wang, “Sparsebev: High-performance sparse 3d object detection from multi-camera videos,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 18 580–18 590.
- [4] 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.
- [5] X. Lin, T. Lin, Z. Pei, L. Huang, and Z. Su, “Sparse4d: Multi-view 3d object detection with sparse spatial-temporal fusion,” *arXiv preprint arXiv:2211.10581*, 2022.

- [6] C. Han, J. Yang, J. Sun, Z. Ge, R. Dong, H. Zhou, W. Mao, Y. Peng, and X. Zhang, "Exploring recurrent long-term temporal fusion for multi-view 3d perception," *IEEE Robotics and Automation Letters*, vol. 9, no. 7, pp. 6544–6551, 2024.
- [7] J. Huang and G. Huang, "Bevdet4d: Exploit temporal cues in multi-camera 3d object detection," *arXiv preprint arXiv:2203.17054*, 2022.
- [8] Z. Li, S. Lan, J. M. Alvarez, and Z. Wu, "Bevnext: Reviving dense bev frameworks for 3d object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 20 113–20 123.
- [9] S. Papais, L. Wang, B. Cheong, and S. L. Waslander, "Foresight: Multi-view streaming joint object detection and trajectory forecasting," *arXiv preprint arXiv:2508.07089*, 2025.
- [10] S. Doll, N. Hanselmann, L. Schneider, R. Schulz, M. Cordts, M. Enzweiler, and H. Lensch, "Dualad: Disentangling the dynamic and static world for end-to-end driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 14 728–14 737.
- [11] K.-C. Huang, M.-H. Yang, and Y.-H. Tsai, "Delving into motion-aware matching for monocular 3d object tracking," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 6909–6918.
- [12] J. Gu, C. Hu, T. Zhang, X. Chen, Y. Wang, Y. Wang, and H. Zhao, "Vip3d: End-to-end visual trajectory prediction via 3d agent queries," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 5496–5506.
- [13] 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.
- [14] W. Sun, X. Lin, Y. Shi, C. Zhang, H. Wu, and S. Zheng, "Sparsedrive: End-to-end autonomous driving via sparse scene representation," in *IEEE International Conference on Robotics and Automation*, 2025, pp. 8795–8801.
- [15] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 621–11 631.
- [16] F. Manhardt, W. Kehl, and A. Gaidon, "Roi-10d: Monocular lifting of 2d detection to 6d pose and metric shape," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2069–2078.
- [17] J. Philion and S. Fidler, "Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d," in *European Conference on Computer Vision*. Springer, 2020, pp. 194–210.
- [18] C. Reading, A. Harakeh, J. Chae, and S. L. Waslander, "Categorical depth distribution network for monocular 3d object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 8555–8564.
- [19] A. Hu, Z. Murez, N. Mohan, S. Dudas, J. Hawke, V. Badrinarayanan, R. Cipolla, and A. Kendall, "Fiery: Future instance prediction in bird's-eye view from surround monocular cameras," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15 273–15 282.
- [20] S. Fang, Z. Wang, Y. Zhong, J. Ge, and S. Chen, "Tbp-former: Learning temporal bird's-eye-view pyramid for joint perception and prediction in vision-centric autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 1368–1378.
- [21] Y. Wang, V. C. Guizilini, T. Zhang, Y. Wang, H. Zhao, and J. Solomon, "Detr3d: 3d object detection from multi-view images via 3d-to-2d queries," in *Conference on Robot Learning*. PMLR, 2022, pp. 180–191.
- [22] Y. Liu, T. Wang, X. Zhang, and J. Sun, "Pet: Position embedding transformation for multi-view 3d object detection," in *European Conference on Computer Vision*. Springer, 2022, pp. 531–548.
- [23] S. Doll, R. Schulz, L. Schneider, V. Benzin, M. Enzweiler, and H. P. Lensch, "Spatialdetr: Robust scalable transformer-based 3d object detection from multi-view camera images with global cross-sensor attention," in *European Conference on Computer Vision*. Springer, 2022, pp. 230–245.
- [24] Y. Zhang, Z. Zhu, W. Zheng, J. Huang, G. Huang, J. Zhou, and J. Lu, "Beverse: Unified perception and prediction in birds-eye-view for vision-centric autonomous driving," *arXiv preprint arXiv:2205.09743*, 2022.
- [25] T. Zhang, X. Chen, Y. Wang, Y. Wang, and H. Zhao, "Mutr3d: A multi-camera tracking framework via 3d-to-2d queries," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 4537–4546.
- [26] X. Lin, Z. Pei, T. Lin, L. Huang, and Z. Su, "Sparse4d v3: Advancing end-to-end 3d detection and tracking," *arXiv preprint arXiv:2311.11722*, 2023.
- [27] M. Bansal, A. Krizhevsky, and A. Ogale, "Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst," *arXiv preprint arXiv:1812.03079*, 2018.
- [28] Y. Chai, B. Sapp, M. Bansal, and D. Anguelov, "Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction," *arXiv preprint arXiv:1910.05449*, 2019.
- [29] 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.
- [30] J. Ngiam, B. Caine, V. Vasudevan, Z. Zhang, H.-T. L. Chiang, J. Ling, R. Roelofs, A. Bewley, C. Liu, A. Venugopal *et al.*, "Scene transformer: A unified architecture for predicting multiple agent trajectories," *arXiv preprint arXiv:2106.08417*, 2021.
- [31] 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.
- [32] 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.
- [33] W. Luo, B. Yang, and R. Urtasun, "Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3569–3577.
- [34] M. Liang, B. Yang, W. Zeng, Y. Chen, R. Hu, S. Casas, and R. Urtasun, "Pnpnet: End-to-end perception and prediction with tracking in the loop," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 553–11 562.
- [35] N. Peri, J. Luiten, M. Li, A. Ošep, L. Leal-Taixé, and D. Ramanan, "Forecasting from lidar via future object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 17 202–17 211.
- [36] S. Casas, B. Agro, J. Mao, T. Gilles, A. Cui, T. Li, and R. Urtasun, "Detra: A unified model for object detection and trajectory forecasting," in *European Conference on Computer Vision*. Springer, 2024, pp. 326–342.
- [37] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [38] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017, pp. 2117–2125.
- [39] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai, "Deformable detr: Deformable transformers for end-to-end object detection," *arXiv preprint arXiv:2010.04159*, 2020.
- [40] Z. Gao, L. Wang, B. Han, and S. Guo, "Adamixer: A fast-converging query-based object detector," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 5364–5373.
- [41] B. Lang, X. Li, and M. C. Chuah, "Bev-tp: End-to-end visual perception and trajectory prediction for autonomous driving," *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 11, pp. 18 537–18 546, 2024.
- [42] 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.
- [43] B. Jiang, S. Chen, X. Wang, B. Liao, T. Cheng, J. Chen, H. Zhou, Q. Zhang, W. Liu, and C. Huang, "Perceive, interact, predict: Learning dynamic and static clues for end-to-end motion prediction," *arXiv preprint arXiv:2212.02181*, 2022.
- [44] 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.