

GraphAD: Interaction Scene Graph for End-to-end Autonomous Driving

Yunpeng Zhang¹, Deheng Qian¹, Ding Li¹, Yifeng Pan¹, Yong Chen²,
Zhenbao Liang², Zhiyao Zhang², Yingzong Liu², Jianhui Mei², Maolei Fu²,
Yun Ye¹, Zhujiu Liang¹, Yi Shan¹ and Dalong Du^{1*}

¹PhiGent Robotics

²Geely Automobile Research Institute

yunpengzhang97@gmail.com, dalong.du@phigent.ai

Abstract

Modeling complicated interactions among the ego-vehicle, road agents, and map elements has been a crucial part for safety-critical autonomous driving. Previous work on end-to-end autonomous driving relies on the attention mechanism to handle heterogeneous interactions, which fails to capture geometric priors and is also computationally intensive. In this paper, we propose the Interaction Scene Graph (ISG) as a unified method to model the interactions among the ego-vehicle, road agents, and map elements. With the representation of the ISG, the driving agents aggregate essential information from the most influential elements, including the road agents with potential collisions and the map elements to follow. Since a mass of unnecessary interactions are omitted, the more efficient scene-graph-based framework is able to focus on indispensable connections and leads to better performance. We evaluate the proposed method for end-to-end autonomous driving on the nuScenes dataset. Compared with strong baselines, our method significantly outperforms in full-stack driving tasks.

1 Introduction

The conventional Autonomous Driving (AD) system is manually divided into multiple sequential modules, including perception [Huang *et al.*, 2021; Li *et al.*, 2022b; Liu *et al.*, 2022], prediction [Hu *et al.*, 2021a; Jia *et al.*, 2023a], planning [Hu *et al.*, 2021b], and control [Wu *et al.*, 2022]. However, the manual division prevents the system from being optimized jointly and globally, resulting in sub-optimal performance. To address this issue, end-to-end autonomous driving algorithms [Hu *et al.*, 2023; Jiang *et al.*, 2023; Zhang *et al.*, 2021] optimize different modules altogether, making the whole system differentiable. With the potential of reducing accumulated errors and achieving higher performance, end-to-end algorithms are drawing increasing attention [Chen *et al.*, 2023]. In end-to-end driving algorithms, both the prediction [Gu *et al.*, 2023] and the planning [Zeng *et al.*, 2019;

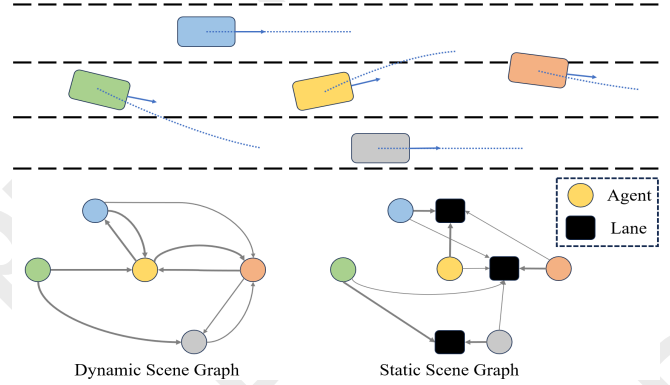


Figure 1: The Interaction Scene Graph is composed of the Dynamic Scene Graph (DSG) and the Static Scene Graph (SSG). In DSG, the traffic agents, represented by the round nodes, pay attention to the surrounding agents by the directed connections. In SSG, the traffic agents reason about their trajectories based on the connected lanes which are represented by the rectangular nodes.

Hu *et al.*, 2021b] modules share the same task of predicting future trajectories of agents (*i.e.*, road agents in the prediction task and the ego vehicle in the planning task). The future trajectories are affected by the interactions among the agents and surrounding environments. Hence, modeling the interactions plays a central role in conventional end-to-end algorithms, which is commonly concreted by the attention mechanism [Hu *et al.*, 2023; Jiang *et al.*, 2023]. However, the attention mechanism, mainly based on correlations of implicit features, lacks the prior knowledge of geometry about which driving elements are more important. As a result, the attention-based interactions inevitably waste their modeling capacities on the unimportant driving elements, while performing worse when being impaired by nuisance elements.

In this paper, we propose the Interaction Scene **Graph** for end-to-end **A**utonomous **D**riving (**GraphAD**) to enhance the interactions among driving elements. GraphAD encodes strong prior knowledge of the interactions into a graph model, the Interaction Scene Graph (ISG). The ISG is a directed graph model whose nodes represent key driving elements in the environment, including traffic agents and lanes. The driving elements are carefully selected for information aggregation, such that only important driving elements are represented. The directed edges in the graph represent the inter-

*Corresponding author.

actions between the nodes. Each node is linked to only a small number of other nodes, making the edges sparse. As a result, the ISG is a concise and efficient representation of the interactions.

Specifically, the ISG consists of two complementary parts, including the Dynamic Scene Graph (DSG) and the Static Scene Graphs (SSG), as shown in Fig. 1. The DSG focuses on the interactions among agents. Each node of DSG corresponds to an agent. Each edge has a weight, measuring the attention one agent pays to the other. The weights are employed to predict the future trajectories of the agents. Note that the interactions among agents depend on the future trajectories, *i.e.*, an agent would pay more attention to one another if their trajectories would collide in the future. Hence, the weights in DSG and the predicted future trajectories are interdependent. The predicted future trajectories can in turn refine the weights in the DSG. So we optimize the DSG and the predicted future trajectories iteratively. The SSG depicts the interaction between the agent and the surrounding map elements. Each agent is represented as a node in the SSG. In the meanwhile, the surrounding lanes in the map are also represented as nodes. The directed edges come from the agent and go to the lanes, modeling the attention the agents pay to the lanes. We apply graph neural networks [Kipf and Welling, 2016; Velivckovic *et al.*, 2017] on the DSG and SSG. The extracted features are utilized to predict the future trajectories of the ego and all other agents. By such means, we are able to adopt a unified method to accomplish both the prediction and the planning tasks.

We evaluate our method on the nuScenes dataset [Caesar *et al.*, 2020]. Extensive ablation studies are conducted to demonstrate the effectiveness of our design choices. We summarize our main contributions as follows: First, GraphAD is the first end-to-end autonomous algorithm which employs a graph model to describe the complex interactions in traffic scenes. The graph model allows us to introduce strong prior knowledge of the traffic scene into the algorithm effectively and efficiently. Second, we elaborately devise the ISG which concisely presents the heterogeneous interactions among ego vehicle, traffic agents, and map elements. In particular, the DSG is able to iteratively refine the prediction of future trajectories and describe subtle interactive games among agents. Finally, when compared with strong baselines [Hu *et al.*, 2023; Hu *et al.*, 2022; Jiang *et al.*, 2023], our method achieves state-of-the-art performance on multiple tasks.

2 Related Work

2.1 End-to-End Autonomous Driving

Instead of adopting a modular paradigm in the traditional AD framework, end-to-end methods, which aim to output future actions based on sensor inputs, have attracted considerable attention. When formulated in an end-to-end manner, the whole framework can be optimized towards the ultimate planning task with high computational efficiency [Chen *et al.*, 2023; Hu *et al.*, 2023]. Some pioneering approaches attempt to directly predict the planned trajectory while lacking explicit supervision of intermediate perception and prediction

tasks [Prakash *et al.*, 2021; Wu *et al.*, 2022; Wu *et al.*, 2023; Zhang *et al.*, 2021; Jia *et al.*, 2023b]. Considering the transparency and interpretability for safety, recent works [Casas *et al.*, 2021; Hu *et al.*, 2022; Chen and Krahenbuhl, 2022; Hu *et al.*, 2023; Ye *et al.*, 2023; Jiang *et al.*, 2023] introduce requisite preceding tasks in the end-to-end framework, thus unifying perception, prediction, and planning into a holistic model. For instance, UniAD [Hu *et al.*, 2023], which regards task-specific queries as a powerful tool for message passing throughout the AD pipeline, has achieved remarkable performance in both multi-object tracking, online mapping, motion forecasting, occupancy prediction, and planning. FusionAD [Ye *et al.*, 2023] extends the capacity of UniAD [Hu *et al.*, 2023] with multi-modal input. In the meantime, some researchers focus on the impact of different privileged inputs. VAD [Jiang *et al.*, 2023] contends that end-to-end AD can be performed in a fully vectorized manner with high efficiency, while OccNet [Tong *et al.*, 2023] attempts to perform the planning task based on the predicted occupancy.

Despite previous methods that have gained impressive performance, the interactions between traffic agents and the surrounding environment are not fully explored. In this work, we propose the Interaction Scene Graph to explicitly model the heterogeneous interactions between the dynamic and static driving elements.

2.2 Graph Neural Networks

Thanks to the success of Graph Neural Networks in graph data, GNNs [Kipf and Welling, 2016; Hamilton *et al.*, 2017; Velivckovic *et al.*, 2017] have been widely adopted in various fields, such as object detection [Wang and Solomon, 2021; Shi and Rajkumar, 2020; Xu *et al.*, 2019], skeleton-based action recognition [Yan *et al.*, 2018; Li *et al.*, 2023a], person re-identification [Yang *et al.*, 2020]. Also, GNN-related advances attract researchers in the autonomous driving community, several studies propose to leverage the ability of GNNs for scene perception and motion prediction. GNN3DMOT [Weng *et al.*, 2020] and PTP [Weng *et al.*, 2021] attempt to model inherit interactions among detected targets for 3D multi-object tracking. For online mapping, LaneGCN [Liang *et al.*, 2020] constructs a lane graph from an HD map, and TopoNet [Li *et al.*, 2023b] introduces relation modeling between lane and traffic elements with a learned scene knowledge graph. In multi-agent motion forecasting, both moving agents and map elements are designed as nodes in graph construction, and the introduction of the relationship among them would benefit trajectory prediction [Salzmann *et al.*, 2020; Mo *et al.*, 2022]. HDGT [Jia *et al.*, 2023a] devises a heterogeneous graph and explicitly models all semantics and relations in the scene. Different from prior works, GraphAD is the first to capture the interactions among dynamic agents and static map elements in the end-to-end AD framework. Also, GraphAD proposes to consider the potential movements of dynamic agents in graph construction by introducing the trajectory proposals.

3 Method

The overall framework of GraphAD is presented in Fig. 2. First, with multi-view video sequences, camera parameters,

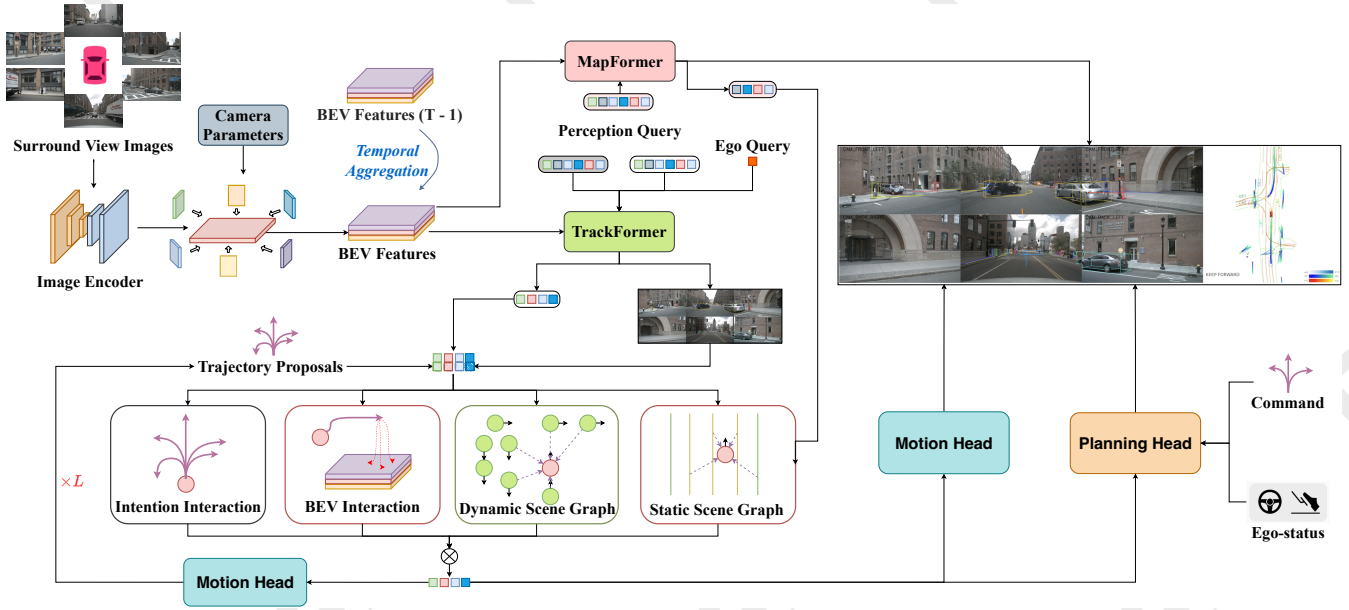


Figure 2: GraphAD features the graph-based interactions between the structured instances in the driving environment, including the dynamic traffic agents and the static map elements. GraphAD first constructs the spatiotemporal scene feature on the Bird-Eye-View as the unified representation for downstream tasks. Then, GraphAD extracts the structured instances by the TrackFormer and the MapFormer. Taking these instances as graph nodes, GraphAD proposes the Interaction Scene Graph to iteratively refine the features of dynamic nodes, by considering the inter-agent and agent-map interactions. Finally, the processed node features are utilized for motion prediction and end-to-end planning.

and ego-poses as input, the image features are extracted by the image encoder and then lifted to the Bird-Eye-View (BEV) features. The multi-frame BEV features are further aggregated to form the spatiotemporal scene representation. Second, GraphAD employs two transformer decoders, *i.e.* the TrackFormer and the MapFormer, to extract the structured representations for the dynamic and static driving elements. Third, the Interaction Scene Graph is explicitly constructed to model the interactions among the ego vehicle, dynamic elements, and static elements, by considering the potential movements. Finally, the graph-aggregated ego query feature, combined with ego status features and high-level driving command, is processed by the planning head to predict the ego-vehicle trajectory. We elaborate on the designs of these steps in the following sections.

3.1 Spatiotemporal Scene Representation

Image Encoder. The image encoder includes a backbone network for multi-scale feature extraction and a neck for fusing these features. Formally, with the multi-view images $I \in \mathbb{R}^{N \times 3 \times H_I \times W_I}$ as input, the image encoder creates the extracted visual features $\mathbf{F}_{2d} \in \mathbb{R}^{N \times C_I \times H'_I \times W'_I}$, where N is the number of camera views, C_I is the channel number, (H_I, W_I) and (H'_I, W'_I) are the input and downsampled image sizes. The output visual feature can contain the fundamental semantics and geometry of the surrounding environment.

Image-to-BEV Transformation. To build a unified scene representation for temporal aggregation and multi-task inference, we lift the multi-view image features to BEV representations with the Lift-Splat-Shoot paradigm [Pillion and

Fidler, 2020; Huang *et al.*, 2021; Li *et al.*, 2022a]. Specifically, the image features \mathbf{F}_{2d}^t at time t are processed to create the context features $\mathbf{F}_{con}^t \in \mathbb{R}^{N \times C \times H'_I \times W'_I}$ and categorical depth distributions $\mathbf{D}^t \in \mathbb{R}^{N \times D \times H'_I \times W'_I}$, where C is the channel number and D is the number of depth bins. The outer product $\mathbf{F}_{con}^t \otimes \mathbf{D}^t$ is then computed as lifted feature point cloud $\mathbf{P}^t \in \mathbb{R}^{NDH'_I W'_I \times C}$. Finally, the voxel-pooling is employed to process the feature points and generate the BEV feature $\mathbf{F}_{BEV}^t \in \mathbb{R}^{C \times H \times W}$ at time t .

Temporal Feature Aggregation. The multi-frame BEV features $\{\mathbf{F}_{BEV}^t\}_{t=t_{cur}-T+1}^{t_{cur}}$, where t_{cur} is the current time and $T \in \mathbb{N}_+$ is the number of frames, are first warped into the ego-centric coordinate system at the current time so that the ego-motion misalignment is removed. Afterward, the aligned multi-frame BEV features are concatenated along the channel dimension and further processed by a convolutional BEV encoder. The output spatiotemporal BEV feature $\mathbf{F}_{BEV} \in \mathbb{R}^{C_o \times H \times W}$ will serve as the unified spatiotemporal scene representation for downstream tasks.

3.2 Structured Element Learning

Based on the spatiotemporal scene features, the extraction of structured elements, including traffic agents and map elements, is important for safety-critical planning in autonomous driving. Therefore, GraphAD utilizes the TrackFormer and the MapFormer to predict these driving-related instances.

TrackFormer. With the spatiotemporal BEV representation, the TrackFormer aims to perform end-to-end 3D object detection and tracking. Following the design of [Hu *et al.*, 2023], we employ two groups of object queries and the trans-

former decoder to solve the problem. Specifically, one group of track queries, which corresponds to previously detected objects, is still required to predict the updated 3D bounding boxes of the same object identities. The other group of detection queries is responsible for the objects which are visible for the first time. For each timestamp, the positive queries, including the tracked and newborn, will serve as the track queries for the next timestamp. The transformer decoder layer includes the self-attention between all object queries and the deformable attention for attending to the spatiotemporal BEV features.

MapFormer. To better capture the geometric constraints of map elements, we follow recent practices [Liao *et al.*, 2022; Jiang *et al.*, 2023] to learn vectorized representations of local maps. Specifically, the MapFormer utilizes the instance-level and point-level queries to form the hierarchical map queries, which are processed by a similar transformer decoder as in TrackFormer. Finally, the output map queries are projected to the class scores and a series of BEV coordinates of potential map elements. To fully capture the map information, four kinds of elements are modeled, including the lane centerline, lane divider, road boundary, and pedestrian crossing.

3.3 Interaction Scene Graph

With the extracted driving instances in structured formats, including traffic agents and map elements, the key challenge lies in how the network can perceive heterogeneous interactions. These interactions, including the driving game between dynamic agents, or the simple centerline-following heuristics, are important for forecasting the future of surrounding environment and making driving decisions. To this end, we construct the Interaction Scene Graph to capture these heterogeneous interactions. As an iterative process, the Interaction Scene Graph functions in three steps. First, all dynamic and static elements are formulated as graph node representations, including explicit geometry and implicit features. Second, the Interaction Scene Graph is constructed with strong geometric priors. Third, the graph node features are updated based on the established graph edges, which are further processed to update the geometry. The detailed formulation is elaborated in the following paragraphs.

Graph Node Representation. The Interaction Scene Graph is constructed on the structured nodes of traffic agents and map elements. Each graph node is designed to include both the explicit geometry and the implicit features. Note that the ego-vehicle is treated as one of the traffic agents to participate in graph-based interactions.

Specifically, the graph nodes of traffic agents, *i.e.* dynamic graph nodes, are organized as one set $P^d = \{p_1^d, \dots, p_{N_d}^d\}$, where N_d is the number of dynamic graph nodes. Also, $p_i^d = (\mathbf{x}_i^d, \mathbf{f}_i^d)$ represents the node representation with its trajectory proposal $\mathbf{x}_i^d \in \mathbb{R}^{M_d \times 2}$ as BEV coordinates and its node feature $\mathbf{f}_i^d \in \mathbb{R}^{C_g}$ with C_g channels, where M_d is the time horizon of trajectory prediction. The trajectory proposals are the trajectory predictions from the previous layer. For the first layer, the clustering results from k-means are utilized instead. The implicit node features are computed as the combination of previous node features, queries from the Track-

Former, embeddings of trajectory proposals, and learnable intention embeddings, following [Hu *et al.*, 2023]. For the unified formulation, we treat different modalities of the same agent as different dynamic graph nodes.

Similarly, the graph nodes of map elements, *i.e.* static graph nodes, are organized as the other set $P^s = \{p_1^s, \dots, p_{N_s}^s\}$, where N_s is the number of static graph nodes and $p_i^s = (\mathbf{x}_i^s, \mathbf{f}_i^s)$ represents one map element by a series of BEV coordinates $\mathbf{x}_i^s \in \mathbb{R}^{M_s \times 2}$ with M_s points and its node feature $\mathbf{f}_i^s \in \mathbb{R}^{C_g}$ with C_g channels. The structured predictions from the MapFormer, including the BEV coordinates and the output query features, are directly used as the static graph nodes. Since the map elements in the driving scenes usually serve as constant environment constraints, their node features are not updated in the iterative layers.

Graph Connection Construction. To capture the heterogeneous interactions between all graph nodes, the Interaction Scene Graph consists of the Dynamic Scene Graph (DSG) and the Static Scene Graph (SSG). The Dynamic Scene Graph is formulated as $G^d = (P^d, E^d)$ by using the traffic agents as dynamic graph nodes, which intends to model the driving game between these agents. The Static Scene Graph is formulated as $G^s = (P^d, P^s, E^s)$ by incorporating both the dynamic and static graph nodes, which focuses on providing the appropriate map information for the dynamic agents. For both DSG and SSG, we follow the same high-level philosophy for computing the edge connections. Specifically, we compute the pair-wise distances between nodes and connect each node to its K nearest neighbors. Despite the straightforward formulation, the design choices of pair-wise distance functions are still underexplored.

Existing graph-based methods [Wang and Solomon, 2021; Jia *et al.*, 2023a] usually exploit the pair-wise distance in feature or coordinate spaces. However, the heterogeneous and evolutionary interactions in the constructed scene graph, with dynamic agents and map elements, cannot be well processed by existing approaches. To this end, we propose to utilize the geometric distances based on trajectory proposals to measure the correlations between graph nodes. On the Dynamic Scene Graph, the distance $\mathcal{H}^d(p_i^d, p_j^d)$ between two dynamic graph nodes is computed as the minimal distance between their trajectory proposals at each time, as in Eq. (1):

$$\mathcal{H}^d(p_i^d, p_j^d) = \min_{t=1}^{M_d} \|\mathbf{x}_i^d(t) - \mathbf{x}_j^d(t)\|_2, \quad (1)$$

where $\mathbf{x}_i^d(t)$ refers to the predicted future position at time t . On the Static Scene Graph, the distance $\mathcal{H}^s(p_i^d, p_j^s)$ between a dynamic node and a static node is computed as the minimal distance between the dynamic trajectory proposal and the static map coordinates, as in Eq. (2):

$$\mathcal{H}^s(p_i^d, p_j^s) = \min_{t=1}^{M_d} \left(\min_{k=1}^{M_s} \|\mathbf{x}_i^d(t) - \mathbf{x}_j^s(k)\|_2 \right), \quad (2)$$

where $\mathbf{x}_i^d(t)$ refers to the predicted future position at time t and $\mathbf{x}_j^s(k)$ refers to the k -th coordinate point of the predicted map element. When the pair-wise distances are computed, the nearest K graph nodes with minimal distances are selected as the graph neighbors.

Method	L2 (m) ↓				Collision (%) ↓			
	1s	2s	3s	Avg.	1s	2s	3s	Avg.
NMP [†] [Zeng <i>et al.</i> , 2019]	-	-	2.31	-	-	-	1.92	-
SA-NMP [†] [Zeng <i>et al.</i> , 2019]	-	-	2.05	-	-	-	1.59	-
FF [†] [Hu <i>et al.</i> , 2021b]	0.55	1.20	2.54	1.43	0.06	0.17	1.07	0.43
EO [†] [Khurana <i>et al.</i> , 2022]	0.67	1.36	2.78	1.60	0.04	0.09	0.88	0.33
ST-P3 [Hu <i>et al.</i> , 2022]	1.33	2.11	2.90	2.11	0.23	0.62	1.27	0.71
UniAD [Hu <i>et al.</i> , 2023]	0.48	0.96	1.65	1.03	0.05	0.17	0.71	0.31
VAD* [Jiang <i>et al.</i> , 2023]	0.54	1.15	1.98	1.22	0.00	0.33	1.07	0.47
GPT-Driver [Mao <i>et al.</i> , 2023a]	0.27	0.74	1.52	0.84	0.07	0.15	1.10	0.44
Agent-Driver [Mao <i>et al.</i> , 2023b]	0.22	0.65	1.34	0.74	0.02	0.13	0.48	0.21
GraphAD	0.32	0.61	1.10	0.68	0.03	0.07	0.25	0.12

Table 1: **Benchmark results for open-loop planning performance.** † denotes LiDAR-based methods. * represents the reproduced results with official checkpoints. GraphAD achieves the state-of-the-art planning performance.

Graph Feature Aggregation. Since the interaction connections have been established, the final part is to refine the node feature by aggregating the information from its connected neighbors. A simple yet effective approach is proposed for the feature aggregation in the Interaction Scene Graph. Specifically, the feature of each neighbor node is concatenated with the target node and then processed by a Multi-Layer Perceptron (MLP). Finally, the permutation-invariant max-pooling is used to aggregate the processed neighbor features into the target node. Also, the Dynamic Scene Graph and Static Scene Graph share the same approach for graph feature aggregation. At the end of each iteration layer, the updated features of dynamic agents are utilized to predict their multi-modal trajectories, including the probability score and the trajectory points for each modality. The predicted trajectory points are further used to update the geometric node features into the next iteration layer.

3.4 Planning Head

Planning Head Structure. The input information for the planning head includes the high-level driving command, the ego-status features, and the processed ego-query from the Interaction Scene Graph. The three groups of features are concatenated and processed by a simple MLP for the final planning predictions.

Ego-status Features. The ego-status information, which mainly includes the velocity, acceleration, and angular velocity, is important for the open-loop planning performance. Therefore, we use a small Multi-Layer Perceptron (MLP) to encode the ego-status information, along with the history trajectories of the ego-vehicle, into the ego-status features.

Occupancy-based Post-optimization. To further avoid the collision with other road agents and ensure the driving safety, we follow the implementation of UniAD [Hu *et al.*, 2023] to train an occupancy head, whose predictions can be utilized to post-optimize the predicted planning trajectories.

3.5 Training

Loss Functions. The loss functions include the depth estimation loss \mathcal{L}_{depth} , the TrackFormer loss \mathcal{L}_{track} , the MapFormer loss \mathcal{L}_{map} , the motion trajectory loss \mathcal{L}_{motion} , the

Method	minADE(m)↓	minFDE(m)↓	MR↓	EPA↑
Constant Pos.	5.80	10.27	0.347	-
Constant Vel.	2.13	4.01	0.318	-
PnPNet	1.15	1.95	0.226	0.222
ViP3D	2.05	2.84	0.246	0.226
UniAD	0.71	1.02	0.151	0.456
GraphAD	0.68	0.98	0.161	0.514

Table 2: **Benchmark results for motion-forecasting.**

DSG	SSG	minADE(m)↓	minFDE(m)↓	MR↓
✓		0.683	1.014	0.165
	✓	0.684	1.018	0.167
✓	✓	0.665	0.989	0.160
Attention	Attention	0.678	1.000	0.160

Table 3: The ablation studies for the Interaction Scene Graph.

occupancy loss \mathcal{L}_{occ} , and the planning loss \mathcal{L}_{plan} . GraphAD is end-to-end trained with the summation of multi-task losses:

$$\mathcal{L} = \mathcal{L}_{depth} + \mathcal{L}_{track} + \mathcal{L}_{map} + \mathcal{L}_{motion} + \mathcal{L}_{occ} + \mathcal{L}_{plan}. \quad (3)$$

Specifically, we use binary cross-entropy for \mathcal{L}_{depth} and follow existing methods [Hu *et al.*, 2023; Jiang *et al.*, 2023] for training other tasks.

Multi-stage Training. First, GraphAD is trained to jointly predict the 3D object detection and vectorized map elements. Second, we freeze the image backbone and train GraphAD for tracking, vectorized map, and graph-based motion prediction. Finally, we further add the tasks of occupancy prediction and planning for end-to-end training.

4 Experiments

Our experiments are conducted on the challenging nuScenes dataset [Caesar *et al.*, 2020], where 1000 complex driving scenes are included. For annotations, over 1.4M 3D bounding boxes of 23 categories are provided in total, and the keyframes are annotated at 2 Hz.

4.1 Implementation Details

For benchmark results, GraphAD adopts the input size of 640 and ResNet101-DCN [He *et al.*, 2016] as the image back-

Similarity	minADE(m)↓	minFDE(m)↓	MR↓
Feature Distance	0.673	0.993	0.160
Current Distance	0.677	0.999	0.160
Trajectory Distance	0.665	0.989	0.160

Table 4: The ablation studies for the choices of node similarity functions.

bone. The image neck generates feature maps with 512 channels and $16\times$ downsampling. For image-to-BEV transformation, GraphAD uses the method in BEVDepth [Li *et al.*, 2022a] to generate the BEV features with 80 channels. Four frames of BEV features are fused to create the spatiotemporal scene representation $\mathbf{F}_{BEV} \in \mathbb{R}^{256 \times 200 \times 200}$. The TrackFormer strictly follows the settings of UniAD [Hu *et al.*, 2023], while the MapFormer uses 100 map queries and a six-layer transformer decoder. The Interaction Scene Graph stacks three iterative layers for motion prediction with six modalities. The number of neighbours is set to 24 for the Dynamic Scene Graph and 8 for the Static Scene Graph. For ego-status features in the planning head, we follow the preprocessing of CAN-bus information from VAD [Jiang *et al.*, 2023]. For ablation studies, we adopt the input size of 256×704 and ResNet50 as the image backbone.

4.2 Metrics

We follow the same evaluation protocol of previous state-of-the-art method UniAD [Hu *et al.*, 2023]. For motion prediction, we employ End-to-end Prediction Accuracy (EPA), Average Displacement Error (ADE), Final Displacement Error (FDE), and Miss Rate (MR). In planning task, Displacement Error (DE, L2 distance) and Collision Rate (CR) are employed, where CR is considered as the main metric.

4.3 Benchmark Results

Planning Results. As shown in Tab. 1, GraphAD achieves the state-of-the-art performance for open-loop planning on the nuScenes validation set. When compared to the second best method, Agent-Driver [Mao *et al.*, 2023b], GraphAD achieves a 42.9% reduction of collision rate, which demonstrates the effectiveness of the proposed Interaction Scene Graph for aggregating information from related traffic agents and map elements.

Prediction Results. The benchmark results for motion prediction on the nuScenes validation set are summarized in Tab. 2. GraphAD achieves the best performance with 0.68 minADE and 0.514 EPA, significantly outperforming the previous best method UniAD [Hu *et al.*, 2023]. The improved performance validates the enhanced capacity of Interaction Scene Graph in modeling the map guidance and intention interaction from other driving instances.

4.4 Ablation Studies

To demonstrate the effectiveness of the proposed Interaction Scene Graph, we conduct extensive ablation studies on the nuScenes validation set.

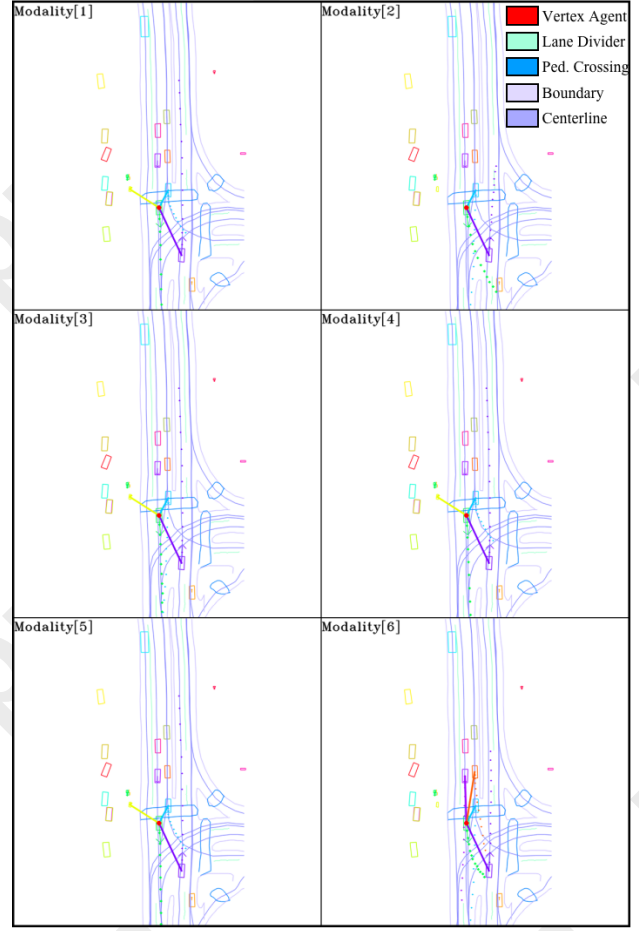


Figure 3: **The qualitative visualization of the Dynamic Scene Graph.** The agent of interest, marked by the red dot, has 6 different modalities of future trajectories. With each motion intention, this agent interacts with the most influential traffic agents, which are denoted by the connections. Faraway connections are omitted for clarity.

Effectiveness of Interaction Scene Graph. In Tab. 3, we ablate the influence of Dynamic Scene Graph (DSG) and Static Scene Graph (SSG) on the motion prediction of traffic agents. We can observe that both types of scene graphs make significant contribution to the performance boost. Since the DSG can model the driving game between dynamic agents and the SSG is able to provide explicit map constraints, both types of graph-based interactions can provide valuable and complementary information for the trajectory prediction. For comprehensive evaluation, we also implement an attention-based variant, where the inter-agent and agent-map interactions are entirely realized by the vanilla attention mechanism. However, we find the attention-based variant, without explicit geometric prior, fails to extract valid information and generates inferior performance.

Design choices of graph node distance. In Tab. 4, we analyze the influence of different methods for computing the distance between graph nodes. ‘‘Feature Distance’’ and ‘‘Current Distance’’ denote distances in the feature space and dis-



Figure 4: **The qualitative visualization of the planning trajectories.** The images from six cameras are shown on the left. The predicted trajectories of traffic agents and the planning result of the ego vehicle are shown on the right. The color intensities of these trajectories vary according to the probability p and the time t . The red arrows highlight the environments which most likely influence the ego vehicle planning.

Method	minADE(m)↓	minFDE(m)↓	MR↓
Attention	0.682	1.017	0.170
MLP + Avg-pooling	0.680	1.014	0.164
MLP + Max-pooling	0.665	0.989	0.160

Table 5: The ablation studies for the graph aggregation methods.

tances between current locations respectively, while “Trajectory Distance” is the distance between potential trajectories. Since the distance function directly determines which neighbour nodes will participate in the feature aggregation, its design choice is of vital importance. From the experimental results, we can find the proposed trajectory distance significantly outperforms the current distance because it explicitly considers the potential interactions in the future, which is crucial for accurate trajectory estimation. On the other hand, the geometric distance on trajectories also outperforms the feature distance. It is possibly because the graph nodes, including both traffic agents and map elements, with different sources and modalities have heterogeneous features.

Design choices of methods for graph feature aggregation. In Tab. 5, we compare different methods for aggregating the neighbour node features to update the vertex. As observed in the table, MLP-based aggregation methods performs better than attention-based methods. Furthermore, the max-pooling operation outperforms the avg-pooling method, reaching 0.665m minADE, 0.989m minFDE and 0.160 MR. Thus, we choose MLP with max-pooling as default setting.

Design choices of planning head. In Tab. 6, we explore the effects of different components for the planning task, where “Graph” refers to the proposed Interaction Scene Graph, “Ego-states” means the utilization of ego-vehicle status, and “Post-optim.” represents the optimization strategy with the predicted occupancy. The following effects can be observed: (1) The incorporation of ego-state features can bring a significant improvement on the planning performance, since the information, like velocity and acceleration, makes it much easier to recover the ego-trajectory. (2) Whether or not the ego-state features are utilized, the proposed method of Interaction Scene Graph consistently improves the planning performance. (3) The post-optimization with the predicted occupancy plays an important role in ensuring the driving safety,

Graph	Ego-states	Post-optim.	Planning	
			L2 (m) ↓	Col. (%) ↓
✓	✓	✓	1.39	1.13
			1.35	1.07
			0.65	0.63
✓	✓		0.64	0.47
✓	✓	✓	0.73	0.15
	✓	✓	0.74	0.22

Table 6: The ablation studies for designs in the planning head.

by avoiding the potential collisions with explicit adjustments. With all above components, GraphAD, with smaller input sizes and image backbone, achieves a remarkable collision rate of 0.15%.

4.5 Qualitative Results

To qualitatively evaluate our method for better understanding, we visualize both the intermediate interactions and the final results of GraphAD. As shown in Fig. 3, the agent of interest has 6 predicted future trajectories for different potential intentions. The dynamic scene graph for each trajectory automatically links the agent to other traffic agents nearby. With these explicit geometry priors, the agent can focus on the interactions with the important agents. From the cases in Fig. 4, GraphAD enables the ego vehicle to maneuver safely in complex situations like road junction and opposite meeting. These planning abilities results from the accurate motion prediction and necessary inter-agent interactions, based on the proposed graph designs.

5 Conclusion

In this paper, we propose a new end-to-end autonomous driving algorithm, GraphAD, which employs an elaborately designed graph to describe heterogeneous interactions in complex traffic scenes. The graph explicitly encodes key driving elements and their relations, allowing us to introduce strong prior knowledge into the algorithm. As a consequence, GraphAD achieves state-of-the-art performance in both the prediction and the planning tasks. The way using graphs to encode more complex interactions among diverse traffic instances, such as traffic lights and routing decisions, needs further exploration.

Acknowledgements

This research is supported by PhiGent Robotics, Beijing Nova Program (20240484562), and Geely Automobile Research Institute.

References

- [Caesar *et al.*, 2020] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *CVPR*, 2020.
- [Casas *et al.*, 2021] Sergio Casas, Abbas Sadat, and Raquel Urtasun. Mp3: A unified model to map, perceive, predict and plan. In *CVPR*, 2021.
- [Chen and Krahenbuhl, 2022] Dian Chen and Philipp Krahenbuhl. Learning from all vehicles. In *CVPR*, 2022.
- [Chen *et al.*, 2023] Li Chen, Penghao Wu, Kashyap Chitta, Bernhard Jaeger, Andreas Geiger, and Hongyang Li. End-to-end autonomous driving: Challenges and frontiers. *arXiv preprint arXiv:2306.16927*, 2023.
- [Gu *et al.*, 2023] Junru Gu, Chenxu Hu, Tianyuan Zhang, Xuanyao Chen, Yilun Wang, Yue Wang, and Hang Zhao. ViP3D: End-to-end visual trajectory prediction via 3d agent queries. In *CVPR*, 2023.
- [Hamilton *et al.*, 2017] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *NeurIPS*, 2017.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [Hu *et al.*, 2021a] Anthony Hu, Zak Murez, Nikhil Mohan, Sofia Dudas, Jeffrey Hawke, Vijay Badrinarayanan, Roberto Cipolla, and Alex Kendall. Fiery: future instance prediction in bird’s-eye view from surround monocular cameras. In *ICCV*, 2021.
- [Hu *et al.*, 2021b] Peiyun Hu, Aaron Huang, John Dolan, David Held, and Deva Ramanan. Safe local motion planning with self-supervised freespace forecasting. In *CVPR*, 2021.
- [Hu *et al.*, 2022] Shengchao Hu, Li Chen, Penghao Wu, Hongyang Li, Junchi Yan, and Dacheng Tao. St-p3: End-to-end vision-based autonomous driving via spatial-temporal feature learning. In *ECCV*, 2022.
- [Hu *et al.*, 2023] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhai Wang, et al. Planning-oriented autonomous driving. In *CVPR*, 2023.
- [Huang *et al.*, 2021] Junjie Huang, Guan Huang, Zheng Zhu, and Dalong Du. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*, 2021.
- [Jia *et al.*, 2023a] Xiaosong Jia, Penghao Wu, Li Chen, Yu Liu, Hongyang Li, and Junchi Yan. Hdgt: Heterogeneous driving graph transformer for multi-agent trajectory prediction via scene encoding. *TPAMI*, 2023.
- [Jia *et al.*, 2023b] Xiaosong Jia, Penghao Wu, Li Chen, Jiangwei Xie, Conghui He, Junchi Yan, and Hongyang Li. Think twice before driving: Towards scalable decoders for end-to-end autonomous driving. In *CVPR*, 2023.
- [Jiang *et al.*, 2023] Bo Jiang, Shaoyu Chen, Qing Xu, Bencheng Liao, Jiajie Chen, Helong Zhou, Qian Zhang, Wenyu Liu, Chang Huang, and Xinggang Wang. Vad: Vectorized scene representation for efficient autonomous driving. In *ICCV*, 2023.
- [Khurana *et al.*, 2022] Tarasha Khurana, Peiyun Hu, Achal Dave, Jason Ziglar, David Held, and Deva Ramanan. Differentiable raycasting for self-supervised occupancy forecasting. In *ECCV*, 2022.
- [Kipf and Welling, 2016] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [Li *et al.*, 2022a] Yinhao Li, Zheng Ge, Guanyi Yu, Jinrong Yang, Zengran Wang, Yukang Shi, Jianjian Sun, and Zeming Li. Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. *arXiv preprint arXiv:2206.10092*, 2022.
- [Li *et al.*, 2022b] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Qiao Yu, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. *arXiv preprint arXiv:2203.17270*, 2022.
- [Li *et al.*, 2023a] Ding Li, Yongqiang Tang, Zhizhong Zhang, and Wensheng Zhang. Cross-stream contrastive learning for self-supervised skeleton-based action recognition. *arXiv preprint arXiv:2305.02324*, 2023.
- [Li *et al.*, 2023b] Tianyu Li, Li Chen, Xiangwei Geng, Huijie Wang, Yang Li, Zhenbo Liu, Shengyin Jiang, Yuting Wang, Hang Xu, Chunjing Xu, et al. Topology reasoning for driving scenes. *arXiv preprint arXiv:2304.05277*, 2023.
- [Liang *et al.*, 2020] Ming Liang, Bin Yang, Rui Hu, Yun Chen, Renjie Liao, Song Feng, and Raquel Urtasun. Learning lane graph representations for motion forecasting. In *ECCV*, 2020.
- [Liao *et al.*, 2022] Bencheng Liao, Shaoyu Chen, Xinggang Wang, Tianheng Cheng, Qian Zhang, Wenyu Liu, and Chang Huang. Maptr: Structured modeling and learning for online vectorized hd map construction. *arXiv preprint arXiv:2208.14437*, 2022.
- [Liu *et al.*, 2022] Yingfei Liu, Tiancai Wang, Xiangyu Zhang, and Jian Sun. Petr: Position embedding transformation for multi-view 3d object detection. *arXiv preprint arXiv:2203.05625*, 2022.
- [Mao *et al.*, 2023a] Jiageng Mao, Yuxi Qian, Hang Zhao, and Yue Wang. Gpt-driver: Learning to drive with gpt. *arXiv preprint arXiv:2310.01415*, 2023.
- [Mao *et al.*, 2023b] Jiageng Mao, Junjie Ye, Yuxi Qian, Marco Pavone, and Yue Wang. A language agent for autonomous driving. *arXiv preprint arXiv:2311.10813*, 2023.

- [Mo *et al.*, 2022] Xiaoyu Mo, Zhiyu Huang, Yang Xing, and Chen Lv. Multi-agent trajectory prediction with heterogeneous edge-enhanced graph attention network. *IEEE Transactions on Intelligent Transportation Systems*, 2022.
- [Phillion and Fidler, 2020] Jonah Phillion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *ECCV*, 2020.
- [Prakash *et al.*, 2021] Aditya Prakash, Kashyap Chitta, and Andreas Geiger. Multi-modal fusion transformer for end-to-end autonomous driving. In *CVPR*, 2021.
- [Salzmann *et al.*, 2020] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *ECCV*, 2020.
- [Shi and Rajkumar, 2020] Weijing Shi and Raj Rajkumar. Point-gnn: Graph neural network for 3d object detection in a point cloud. In *CVPR*, 2020.
- [Tong *et al.*, 2023] Wenwen Tong, Chonghao Sima, Tai Wang, Li Chen, Silei Wu, Hanming Deng, Yi Gu, Lewei Lu, Ping Luo, Dahua Lin, et al. Scene as occupancy. In *ICCV*, 2023.
- [Velivckovic *et al.*, 2017] Petar Velivckovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- [Wang and Solomon, 2021] Yue Wang and Justin M Solomon. Object dgcnn: 3d object detection using dynamic graphs. In *NeurIPS*, 2021.
- [Weng *et al.*, 2020] Xinshuo Weng, Yongxin Wang, Yunze Man, and Kris M Kitani. Gnn3dmot: Graph neural network for 3d multi-object tracking with 2d-3d multi-feature learning. In *CVPR*, 2020.
- [Weng *et al.*, 2021] Xinshuo Weng, Ye Yuan, and Kris Kitani. Ptp: Parallelized tracking and prediction with graph neural networks and diversity sampling. *IEEE Robotics and Automation Letters*, 2021.
- [Wu *et al.*, 2022] Penghao Wu, Xiaosong Jia, Li Chen, Junchi Yan, Hongyang Li, and Yu Qiao. Trajectory-guided control prediction for end-to-end autonomous driving: A simple yet strong baseline. In *NeurIPS*, 2022.
- [Wu *et al.*, 2023] Penghao Wu, Li Chen, Hongyang Li, Xiaosong Jia, Junchi Yan, and Yu Qiao. Policy pre-training for end-to-end autonomous driving via self-supervised geometric modeling. *arXiv preprint arXiv:2301.01006*, 2023.
- [Xu *et al.*, 2019] Hang Xu, Chenhan Jiang, Xiaodan Liang, and Zhenguo Li. Spatial-aware graph relation network for large-scale object detection. In *CVPR*, 2019.
- [Yan *et al.*, 2018] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *AAAI*, 2018.
- [Yang *et al.*, 2020] Jinrui Yang, Wei-Shi Zheng, Qize Yang, Ying-Cong Chen, and Qi Tian. Spatial-temporal graph convolutional network for video-based person re-identification. In *CVPR*, 2020.
- [Ye *et al.*, 2023] Tengju Ye, Wei Jing, Chunyong Hu, Shikun Huang, Lingping Gao, Fangzhen Li, Jingke Wang, Ke Guo, Wencong Xiao, Weibo Mao, et al. Fusionad: Multi-modality fusion for prediction and planning tasks of autonomous driving. *arXiv preprint arXiv:2308.01006*, 2023.
- [Zeng *et al.*, 2019] Wenyuan Zeng, Wenjie Luo, Simon Suo, Abbas Sadat, Bin Yang, Sergio Casas, and Raquel Urtasun. End-to-end interpretable neural motion planner. In *CVPR*, 2019.
- [Zhang *et al.*, 2021] Zhejun Zhang, Alexander Liniger, Dengxin Dai, Fisher Yu, and Luc Van Gool. End-to-end urban driving by imitating a reinforcement learning coach. In *ICCV*, 2021.