

Co-MTP: A Cooperative Trajectory Prediction Framework with Multi-Temporal Fusion for Autonomous Driving

Xinyu Zhang^{1*}, Zewei Zhou^{1*}, Zhaoyi Wang¹, Yangjie Ji¹, Yanjun Huang¹, Hong Chen¹

Abstract—Vehicle-to-everything technologies (V2X) have become an ideal paradigm to extend the perception range and see through the occlusion. Existing efforts focus on single-frame cooperative perception, however, how to capture the temporal cue between frames with V2X to facilitate the prediction task even the planning task is still underexplored. In this paper, we introduce the Co-MTP, a general cooperative trajectory prediction framework with multi-temporal fusion for autonomous driving, which leverages the V2X system to fully capture the interaction among agents in both history and future domains to benefit the planning. In the history domain, V2X can complement the incomplete history trajectory in single-vehicle perception, and we design a heterogeneous graph transformer to learn the fusion of the history feature from multiple agents and capture the history interaction. Moreover, the goal of prediction is to support future planning. Thus, in the future domain, V2X can provide the prediction results of surrounding objects, and we further extend the graph transformer to capture the future interaction among the ego planning and the other vehicles’ intentions and obtain the final future scenario state under a certain planning action. We evaluate the Co-MTP framework on the real-world dataset V2X-Seq, and the results show that Co-MTP achieves state-of-the-art performance and that both history and future fusion can greatly benefit prediction. Our code is available on our project website: <https://xiaomiaozhang.github.io/Co-MTP/>

I. INTRODUCTION

Autonomous vehicles (AVs) should interact with complex and dynamic environments to ensure safe and efficient operation [1]–[3]. Thus, the trajectory prediction has attracted the community’s attention [4]–[6]. The accurate prediction relies on the comprehensive perception information, especially in complex and dense scenarios [7]. The limited perception range and occlusion always lead to catastrophic results [8]–[10], however, most existing works assume that all surrounding objects are fully visible and have complete history trajectories [11]–[13], which is unreal in real vehicle onboard sensors. Thus, vehicle-to-everything (V2X) technology has been an ideal paradigm for sharing information among agents and seeing through occlusion. Nevertheless, existing V2X efforts focus on the single-frame cooperative perception [14]–[16]. Issues such as *How to capture the temporal cue between frames with V2X to facilitate the prediction task and even the planning task, and how to fully employ the V2X information in real-world environments with occlusion* are still underexplored and present deployment challenges.

*Equal contribution.

*Supported by the National Natural Science Foundation of China, Joint Fund for Innovative Enterprise Development (U23B2061)

¹School of Automotive Studies, Tongji University, Shanghai, China. {2332922, zhouzewe, 2233591, jiyj20, yanjun-huang, chen hong2019}@tongji.edu.cn

In the real world, each agent’s sensor differs in capability, accuracy, and coverage [17], [18], and the positioning of these sensors can lead to varied levels of perception and occlusion of the same object. The simplest method of trajectory fusion is association and stitching [19], but it cannot handle the diverse errors from different sensors and still faces the incomplete problem when multiple agents lose the same timestamp in dense scenarios. Thus, it is necessary to explore how cooperative prediction integrates and leverages incomplete history trajectories from multiple agents to obtain more comprehensive history information.

Furthermore, the goal of prediction is not merely to predict the motion of surrounding objects but to support the ego vehicle’s planning [1], [20]. Each planning candidate represents a driving intent, and the vehicle should evaluate each candidate by reasoning the future because the reaction of the environment varies depending on the ego’s actions. If the AV reasons the surroundings solely based on the isolated prediction, it will become too afraid to drive and will always consistently yield to other road users [21], [22]. On the contrary, proactive planning and prediction make AV consider their intention first and the reason for the other reaction. However, the ego planning trajectory is the only future feature in the future domain, leading to aggressive driving behaviors and danger. Thus, for planning-oriented prediction, it is necessary not only to incorporate history interactions but also to consider future interactions among surrounding objects.

To address these challenges, we present Co-MTP, a **Co**operative **M**ulti-**T**emporal fusion **P**rediction framework. First, the infrastructure, with elevated and stationary perspectives, can provide perception and prediction support to each agent in the scenario, *e.g.* AV, human driver, and pedestrian. Second, when AV drives into the infrastructure area, it can receive the infrastructure’s history and prediction information, and fuse them in both history and future domains with the planning information of the ego vehicle to facilitate the planning-oriented prediction. Our contributions can be summarized as follows:

- We propose Co-MTP, a general cooperative trajectory prediction framework with multi-temporal fusion across both history and future domains. Co-MTP is the first framework to fully fuse and exploit comprehensive temporal information in prediction with V2X.
- For the issue of incomplete trajectory in the history domain, we designed a heterogeneous graph to learn the fusion of the history feature from multiple agents with transformers.

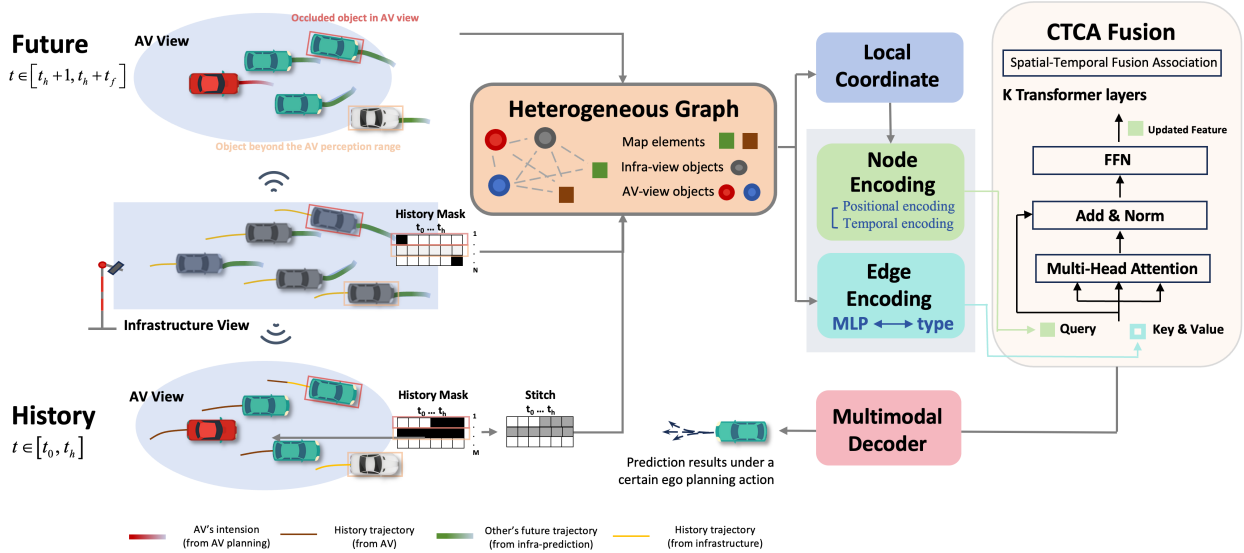


Fig. 1: **The overall architecture of Co-MTP.** In this framework, infrastructures share the history and their prediction results to ego AV. Then, we construct a heterogeneous scene graph with the processed trajectory data and map information, categorizing them according to the types of objects and map elements. Next, we initialize the features of nodes and edges in the relative coordinate system of each object. The CTCA Fusion is used to update the features of the nodes and edges selected by the STSA module over K Transformer layers. Finally, we take the nodes' hidden features from the last layer and input them into the Multimodal Decoder to obtain the multimodal trajectory prediction results.

- To facilitate the following planning, we further extended the heterogeneous graph to future domains to capture the future interaction among the ego planning and the other vehicles' actions from infrastructure prediction.
- We evaluate the Co-MTP framework on real-world dataset V2X-Seq. The results demonstrate its state-of-the-art performance, and both history and future fusion can greatly benefit prediction and planning.

II. RELATED WORK

Cooperative Prediction. The existing effort of V2X research focuses on cooperative perception, and cooperative prediction which considers the temporal information with V2X is still in its infancy [23]–[25]. V2VNet [26] is the first work to explore this field, which proposed an end-to-end perception and prediction model. However, its prediction head is simple and ignores the map information and complex interaction features. Yu, *et al.* [19] presented an open-source V2X sequence dataset for cooperative prediction, V2X-Seq. V2XPnP [25] provides an open-source general framework and V2X sequential dataset for end-to-end cooperative perception and prediction. Furthermore, the V2X-Graph [18] is proposed based on V2X-Seq, which provides a graph neural network (GNN) to learn the cooperative trajectory representations. However, how to fully employ the V2X temporal information, such as the different temporal domains, and fuse the cooperative trajectory still needs to be explored.

Planning-informed Prediction. The couple of prediction and planning is the focus for safe driving in interaction scenarios [1], [27]. PiP model [22] directly incorporated the planning information into prediction, and Liu, *et al.* [21] and Huang, *et al.* [20], further merge the prediction

module and general optimization-based planning module [28], [29] to construct a reaction framework. However, only considering ego planning in the future domain will inadvertently overestimate its priority on the road, leading to overly optimistic predictions and potentially hazardous behaviors [30]. Therefore, it is crucial to integrate the future intentions of other vehicles during prediction to ensure a comprehensive extraction of interaction features across both history and future domains. The other's intention is prediction targets in single-vehicle prediction, and the ego vehicle not only can access history information to see through occlusion, but also the future information as the prior knowledge.

Interaction-aware Prediction. Trajectory prediction is one of the most critical research areas in autonomous driving [4], which considers the perception results sequence and map information to capture the interaction context and predict the others' trajectories [31]. The interaction feature extractor has developed from CNN [22], [32] to GNN and Transformer [3], [33]. The Transformer uses multi-attention to capture the temporal feature across a sequence without relying on sequential order processing [5], [34], which is an advantage in handling incomplete trajectories and fusing the trajectory feature from different agents. Moreover, GNN builds the complex interaction context as a graph with nodes and edges [11], [35], which can facilitate the learning of complex interaction, especially in cooperative prediction involving multi-agent and multi-temporal features.

III. METHODOLOGY

Fig. 1 illustrates the Co-MTP framework. Both AVs and infrastructures can operate prediction individually to support the AV and scenario object separately. The input of the

vehicle-side prediction model is not limited to the history domain but also includes the prediction results from the infrastructure-side prediction model in the future domain.

A. Problem Definition

The input comprises two parts: history trajectories, AV’s future ground truth and map data. Consider a scenario where multiple viewpoints jointly observe a set of Z objects $\Omega = \{a_i | i = 1, 2, \dots, Z\}$ from time step t_0 to t_p . Among them, the AV perceives M objects $\Omega_{AV} = \{b_i | i = 1, 2, \dots, M\}$ and $b_i \in \Omega$, while the infrastructure perceives N agents $\Omega_{infra} = \{c_i | i = 1, 2, \dots, N\}$ and $c_i \in \Omega$. All objects’ features are given, including their bounding boxes (length, width, height) and types (pedestrian, bicycle, vehicle). Define the object trajectory i as $X_i = \{\mathbf{x}_i^{t_0}, \mathbf{x}_i^{t_1}, \dots, \mathbf{x}_i^{t_h}\}$ in a certain frequency, where \mathbf{x}_i^t includes coordinates, heading, velocity, etc. Due to occlusion and other reasons, the observation could be missing in some frames. Thus, each object’s trajectory has a corresponding mask vector $\mathbf{m}_i = \{m_i^{t_0}, m_i^{t_1}, \dots, m_i^{t_h}\}$, which indicates the data validity. If the data exists at time step t , then $m_i^t = 1$; otherwise, the value is 0. In addition, to simulate the planning trajectory from the real-world dataset, we follow the setting of [22], and use the quintic curve to interpolate the future ground truth to follow the planning mission goal. The map data includes roads and signals, and all roads are represented as polylines or polygons, such as road lines, crosswalks, and stop-lines.

The prediction goal is to use the aforementioned information to predict the future trajectories of the surrounding objects. In other words, the task is to learn a model $f(\cdot)$ that outputs $Y_i = \{\mathbf{y}_i^{t_h+1}, \mathbf{y}_i^{t_h+2}, \dots, \mathbf{y}_i^{t_h+t_f}\}$. Additionally, to capture the uncertainty, the predicted output could be multiple modes, each associated with a corresponding probability.

B. Scene Representing with Graph

Heterogeneous Graph. The driving scenario has a lot of heterogeneous elements, such as map, trajectories, different object types and traffic signals. To leverage the powerful scenario-representation capabilities, we build a heterogeneous graph for prediction scenario $\mathcal{G} = \{\mathcal{N}, \mathcal{E}\}$, where \mathcal{N} and \mathcal{E} present the set of nodes and edges separately. Objects and the map are represented as graph nodes, and the relationships between them are represented as edges between the nodes. To represent heterogeneity, we construct a type mapping function \mathcal{T} for the nodes and define the type of each edge based on its originating node.

Trajectory Data Preprocess. In Co-MTP framework, objects’ history trajectories come from both the AV and infrastructure perspectives, along with the corresponding mask matrices. Before learning and fusing the information from these perspectives, we follow the stitch method CBMOT [36] in V2X-Seq [19] to use the high-detection-score to interpolate the vehicle history trajectory with infrastructure history, and the masks also are reorganized. We found the prestitch of the incomplete trajectory can complement the temporal index and accelerate the learning of the following trajectory fusion. This prestitch method does not mean that

the infrastructure information not involved in the interpolation is discarded. All the original objects’ history trajectories from the infrastructure are encoded in the graph nodes. In the future domain, in addition to the AV’s planning trajectory, the infrastructure’s prediction results of the other objects are also included. Since our research does not focus on the infrastructure’s prediction task itself, we use the our Co-MTP model with only infrastructure information to output other agents’ predicted trajectories.

To reduce the computational burden, the number of edges for each node has the upper boundary. Whether the node \mathbf{n} can establish an edge with the node \mathbf{s} depends on the distance between them. For object nodes, we use the position observed at the last time step t_h to calculate the distance to others. For map nodes, the shortest distance from other nodes to the polyline or polygon is selected.

C. Relative Spatial-Temporal Encoding

Local Coordinate System. To facilitate interaction learning in a normalized form, each node requires a local coordinate system $(x_{\text{ref}}, y_{\text{ref}}, \theta_{\text{ref}})$ to provide a unique motion representation for each object, regardless of its global coordinates. All subsequent computations related to the node will be carried out within this local coordinate system. The position and heading of each object at the final time step t_h are taken as $(x_{\text{ref}}, y_{\text{ref}})$ and θ_{ref} . For map data, the center point coordinates are used as $(x_{\text{ref}}, y_{\text{ref}})$, and the direction of the longest line segment within the polyline or polygon is treated as θ_{ref} .

Graph Node Encoding. All the positions of objects from both AV and infrastructure are encoded with a shared-parameter MLP. Then, the position features are concatenated with other information (e.g. bounding box and speed) input into two distinct MLPs—one designed for the AV and the other for the infrastructure with different parameters. We employ two self-attention layers to capture the temporal features of each object in both history and future domains. Since the history trajectory may be incomplete, the mask vector \mathbf{m}_i is also input to the temporal encoding. For map data, we combine the MLP used for coordinate encoding with an improved PointNet to form the map position encoding method. Additionally, we use a set of Linear layers and GRUs to encode traffic signals. Thus, all nodes have been encoded into a set $\{\mathbf{n}^0\}$.

Graph Edge Encoding. All edges are first input into a shared-parameter MLP to encode the coordinate transformation $\Gamma_{\mathbf{s} \rightarrow \mathbf{n}}$. According to the type mapping function $f(\cdot)$, all edges are connected to their starting nodes and the combined input is fed into the corresponding type-specific MLP to generate the initial edge features.

$$\mathbf{e}_{\mathbf{s} \rightarrow \mathbf{n}}^0 = \text{MLP}_{\mathbf{s}}(\text{concat}[\mathbf{s}^0, \text{MLP}(\Gamma_{\mathbf{s} \rightarrow \mathbf{n}})]) \quad (1)$$

D. Cross-Temporal and Cross-Agent (CTCA) Fusion

To capture the complex trajectory features from multiple agents, we propose a cross-temporal graph with the spatial-temporal fusion association (STFA) module. A K-layer Transformer is employed to aggregate and update features for nodes and edges.




Fig. 2: Illustration of STFA in the heterogeneous graph. In addition to the object nodes with preprocessed data from the AV’s view, objects with raw data from the infrastructure’s view also participate as independent nodes. To model future interaction, we treat the AV’s planning and infrastructure’s prediction results as independent nodes, establishing edge relationships with historical nodes in the graph.

Spatial-Temporal Fusion Association. Fig. 2 illustrates the STFA module. It aims to fully integrate information from different views and temporal domains. AV and infrastructure can observe the same object in the same timestamp and have separate history information record, thus, it is necessary to establish a connection between these two agents’ information. We employ an implicit approach to extract the multi-view features across multi-temporal domain. When gathers information, the vehicle-side node \mathbf{v} , after trajectory interpolation, is continuously influenced by the infrastructure node \mathbf{i} . This not only reuses the infrastructure information that was not utilized during the interpolation but also learn the relationship between the two views. Furthermore, when aggregating agent node, future information should be considered. Similarly to the above approach, future nodes provide the agents with the necessary information for aggregation.

Transformer Layer. We utilize the multi-head attention layer (MHA) to capture the feature from various types of in-edges for each node, followed by an MLP to combine the information. During the aggregation phase, the node features are calculated as the query (\mathbf{Q}), and the in-edge features are used as the key (\mathbf{K}) and the value (\mathbf{V}).

$$\mathbf{n}_{type_e}^{k-1,'} = \text{MHA}_{type_e}(\mathbf{n}^{k-1}, \mathbf{e}_{s \rightarrow n}^{k-1}) \quad (2)$$

$$\mathbf{n}^{k-1,''} = \text{MLP}\left(\text{concat}\left(\left\{\mathbf{n}_{type_e}^{k-1,'} \mid type_e = 1, \dots, n_r\right\}\right)\right) \quad (3)$$

where $type_e$ means the type of $\mathbf{e}_{s \rightarrow n}^{k-1}$, and n_r is number of the in-edges’ types for node \mathbf{n} . After that, the features of each type of node are updated with a feed-forward network (FFN).

$$\mathbf{n}^k = \text{FFN}_n(\mathbf{n}^{k-1,''}) \quad (4)$$

where \mathbf{n}^k is the hidden feature of the k^{th} Transformer. FFN_n means the parameters depend on the type of node \mathbf{n} .

For edge feature aggregation and updating, we retrieve the outputs \mathbf{s}^{k-1} and $\Gamma_{s \rightarrow n}$ from the edge encoding stage and combine them with the edge features from the previous layer, $\mathbf{e}_{s \rightarrow n}^{k-1}$. These combined features are then input into an MLP to obtain the updated edge features.

$$\mathbf{e}_{s \rightarrow n}^k = \text{MLP}_{s \rightarrow n}\left(\text{concat}\left[\mathbf{s}^{k-1}, \Gamma_{s \rightarrow n}, \mathbf{e}_{s \rightarrow n}^{k-1}\right]\right) \quad (5)$$

E. Multimodal Decoder

The motion pattern of an agent is inherently multi-modal, where multiple future trajectory possibilities coexist. The hidden features output by the final layer of the transformer layer typically consist of a single set of feature embeddings. In this paper, we first replicate these single embeddings and their corresponding in-edge features K times. Then, we refine the node information through a cross-attention mechanism with processed in-edge features to extract richer feature representations and output K modes via a feed-forward network (FFN).

Traditionally, the K trajectory modes are treated as independent of one another. However, the modes of an object exhibit a competitive relationship, where the rise of one mode leads to the decline of others. Therefore, we input the refined features into a self-attention layer to capture the relationships between the K modes of the same object.

After this refinement process, we employ both regression and classification techniques to decode the agent’s trajectory. Specifically, a multi-layer perceptron (MLP) classifies the K modalities, assigning each a probability. However, simply classifying the trajectory can lead to issues such as trajectory drift, which may violate the agent’s history tend. To mitigate this, we reintroduce the original node features \mathbf{n}^0 and impose dynamic constraints on the regression results. These refined features are concatenated with the regression output and passed through an MLP, with temporal aspects handled by a series of refined convolutional neural networks (CNNs). For a single agent, the final predicted output trajectory can be expressed as follows:

$$Y_{cls} = \{c_k\}, k \in [1, K] \quad (6)$$

$$Y_{reg} = \{(\mathbf{y}_k^1, \mathbf{y}_k^2, \dots, \mathbf{y}_k^T)\}, k \in [1, K] \quad (7)$$

where c_k is probability of the k -th mode, \mathbf{y}_k^t is the position of the agent at the timestamp t , the shape of tensor Y_{cls} and Y_{reg} are $(N \times K)$ and $(N \times K \times T \times 2)$. N is the number of target agents, 2 means the 2-D coordinate (x, y) .

During the training, the total loss is calculated by the regression head and the classification head.

$$\begin{aligned} \mathcal{L} &= \lambda_1 \mathcal{L}_{cls} + \lambda_2 \mathcal{L}_{reg} \\ &= \lambda_1 \frac{1}{N} \sum_{n=1}^N \text{CrossEntropy}(\{c_k\}_n, I(k_n)) \\ &\quad + \lambda_2 \frac{1}{2NT} \sum_{n=1}^N \sum_{t=1}^T [l(x_{n,k}^t, x_n^{t*}) + l(y_{n,k}^t, y_n^{t*})] \end{aligned} \quad (8)$$

where λ_1 and λ_2 are the weights, which are set to 0.1 and 10. $I(k_n)$ is a hot vector, where the k_n -th element is 1 and others are 0. k_n is the mode with the smallest regression loss, calculated as the average deviation between the regression results and the ground truth:

$$k_n = \underset{k \in [1, K]}{\text{argmin}} \frac{1}{2T} \sum_{t=1}^T [l(x_{n,k}^t, x_n^{t*}) + l(y_{n,k}^t, y_n^{t*})] \quad (9)$$

where $l(\cdot)$ is smooth L1 loss.

TABLE I: Performance comparison on the V2X-Seq dataset.

Model	History Fusion	Future Fusion	Metrics ↓ (K=6)		
			minADE	minFDE	MR
TNT [13]	stitch	×	7.38	15.27	0.72
HiVT [33]	stitch	×	1.25	2.33	0.35
V2X-Graph [18]	learning	×	1.17	2.03	0.29
Co-HTTP	stitch	learning	0.83	1.33	0.19
Co-MTP	learning*	learning	0.76	1.15	0.16

* First stitches the AV-perceived trajectory with the infrastructure’s history trajectory, then fuses them through learning.

IV. EXPERIMENTS

A. Experimental Setup

Dataset and Metrics. We evaluate the Co-MTP framework on the real-world V2X dataset, V2X-Seq, which comprises over 50,000 trajectory fragments of vehicle-infrastructure cooperation. Each fragment lasts for 10s, with a recorded frequency of 10 Hz. The dataset labels the type (*e.g.* vehicle, bicycle, pedestrian, etc.) of each agent. To increase the sample size, we slice each original dataset fragment into 20 samples. Each sample contains 3s of history trajectories and 5s of future trajectories. According to standard evaluation protocol, we utilize metrics including minimum Average Displacement Error (minADE), minimum Final Displacement Error (minFDE), and Missing Rate (MR) for evaluation.

Implementation Details. All the trajectories are downsampled to 5Hz to accelerate the training. The map polylines and polygons are segmented, with each segment represented by 21 points. We select a map within 250m of the evaluation area and limit the number of edges for each node in the GNN to a maximum of 32. The prediction mode is set to 6 for fair comparison. 6 layers of Transformer are stacked for the encoder module. The dimension of the hidden feature is set to 256, and the number of heads in all MHA is 8. The learning rate is set to 2.5×10^{-3} initially and decreases over the training epochs. The AdamW optimizer is adopted with a weight decay of 1×10^{-6} . We train the model for 45 epochs with a batch size of 24 on a server with 4 NVIDIA 4090s.

B. Compared Methods

We consider the existing methods on the V2X-Seq dataset and a baselines we proposed for performance comparison.

- **PP-VIC** stitch the AV and infrastructure trajectory through CBMOT, and then test it with the SOAT prediction models TNT [13] and HiVT [33].
- **V2X-Graph** constructs a graph to fuse the history trajectories from AV and infrastructure, however, it just focuses on the history domain fusion.
- **Co-HTTP** is the baseline model, simplified from our Co-MTP model. It uses the CBMOT [36] method to stitch history trajectories from AV and infrastructure.

In the Table I, we can observe that our cooperative framework Co-MTP ranks first across minADE/minFDE/MR in the benchmark of V2X-Seq dataset, with a 35% reduction in minADE compared to V2X-Graph. Moreover, the performance of our proposed baseline is also noteworthy,

which preliminarily indicates that the heterogeneous graph network and cross-temporal fusion strategies are beneficial for prediction in cooperative scenarios.

C. Ablation Study

We investigate the effectiveness of multi-view data processing strategies and the decoder. Ablation studies evaluate variations of Co-MTP separately in both history and future time dimensions, and the results are presented in Table II.

History Fusion. To address the two key questions of what history information should be fused and how to fuse it, we design 4 models based on three conditions: whether the vehicle-side data is available, whether the infrastructure data is available, and which fusion method to adopt. First, the infrastructure has an elevated and static position and shows a better performance than the single-agent prediction. Moreover, as the history information from multiple views becomes more comprehensive, the prediction error gradually decreases. Furthermore, comparing variants 3 and 4 reveals that the fusion method involving full infrastructure information performs better, which means the learning method can fuse the trajectory feature better from noises, and the redundant information of infrastructure can benefit the fusion.

Future Fusion. We aim to demonstrate whether future fusion, specifically the driving intentions of the AV and the prediction output of the infrastructure, are beneficial to the prediction of the AV. Similarly, we design variations for Co-MTP and analyze them alongside the baseline Co-HTTP. Overall, all models incorporating future fusion outperform those that only consider a single history time dimension. Specifically, the accuracy of 5 and 6 is lower than that of Co-MTP, indicating that both the intention of the AV and the other’s future trajectories are crucial for the AV’s prediction task. Considering both AV planning and the other’s future trajectory from infrastructure’s prediction, the prediction performance is further improved in Co-HTTP and Co-MTP.

D. Robustness Assessment

Co-MTP utilizes V2X communication to fuse both history and future information. While this fusion enriches the temporal understanding, it also introduces common V2X practice challenges such as communication delay and noise. As the amount of communication data increases over the temporal domain, ensuring the model’s robustness against such disruptions becomes more critical. To evaluate this robustness, we designed experiments using the same Co-MTP model base, alongside two variants: Co-MTP-no fusion, which excludes the future fusion, and Co-HTTP-nofut, which simply stitches the trajectory without future information.

We conducted robustness assessments by introducing noise and communication delays, assuming a positional deviation of 0.2 meters and a time delay of 0.5 seconds. The results, presented in Table III and Table IV, demonstrate that all models—regardless of whether they communicate the history or future data, and whether they use stitching or learning methods—exhibit strong robustness to noise and delay. This indicates that the rich temporal information shared by multiple agents significantly contributes to system robustness,

TABLE II: Results of Model Ablation Study.

Time Dimension	Model Variation	Vehicle-side History	Infrastructure History	Fusion Method	AV's Planning	Other's Future	minADE ↓	minFDE ↓	MR ↓
History	1	✓	-	-	-	-	0.9887	1.5996	0.2224
	2	-	✓	-	-	-	0.9375	1.5123	0.2081
	3	✓	✓	stitch	-	-	0.9123	1.4647	0.2019
	4	✓	✓	stitch + full-infra	-	-	0.8518	1.3368	0.1828
Future	5	✓	✓	stitch + full-infra	✓	-	0.8378	1.2649	0.1720
	6	✓	✓	stitch + full-infra	-	✓	0.7992	1.1991	0.1659
	Co-HTTP	✓	✓	stitch	✓	✓	0.8258	1.3324	0.1916
	Co-MTP	✓	✓	stitch + full-infra	✓	✓	0.7647	1.1470	0.1594

*stitch represents the CBMOT method, while full-infra indicates that all infrastructure information is included as nodes in graph for fusion.

TABLE III: Noise Analysis

Model	Metrics		
	minADE ↓	minFDE ↓	MR ↓
Co-HTTP-nofut	0.9314 0.9123	1.5387 1.4647	0.2077 0.2019
Co-MTP-nofut	0.8637 0.8518	1.3571 1.3368	0.1847 0.1828
Co-MTP	0.7998 0.7647	1.2293 1.1470	0.1672 0.1594

* Results in small font represent prediction performance without noise.

TABLE IV: Time Delay Assessment

Model	Metrics		
	minADE ↓	minFDE ↓	MR ↓
Co-HTTP-nofut	0.9203 0.9123	1.5170 1.4647	0.2064 0.2019
Co-MTP-nofut	0.8696 0.8518	1.3584 1.3368	0.1810 0.1828
Co-MTP	0.7753 0.7647	1.1849 1.1470	0.1643 0.1594

* Results in small font represent prediction performance without delay.

particularly in handling time delays. The robustness to delay is attributed to the system's ability to effectively leverage temporal information, thus mitigating the impact of delay.

In the case of communication noise, the Co-MTP model shows a decline in accuracy. It can be attributed to the compounded errors affecting both history and future domain information, as noise introduces inaccuracies in both. Future research could explore further improvements in noise robustness. Nevertheless, it is noteworthy that despite the performance drop due to noise, Co-MTP continues to demonstrate superior accuracy compared to other models, underscoring its potential in V2X applications.

E. Qualitative Result

Figure 3 presents the prediction performance of Co-MTP on V2X-Seq in classic driving scenarios, such as following, highway driving, accelerating through an intersection, and waiting to turn. In each sub-figure, the predicted trajectories closely adhere to road structure constraints and traffic safety regulations. Specifically, in (d), the target vehicle intends to turn. However, due to the influence of the object's intention to go straight without decelerating, the target vehicle is forced to wait. This reinforces the point that future information is crucial for prediction.

V. CONCLUSION

To fully exploit comprehensive temporal information in prediction with V2X, we introduce the Co-MTP, the first gen-

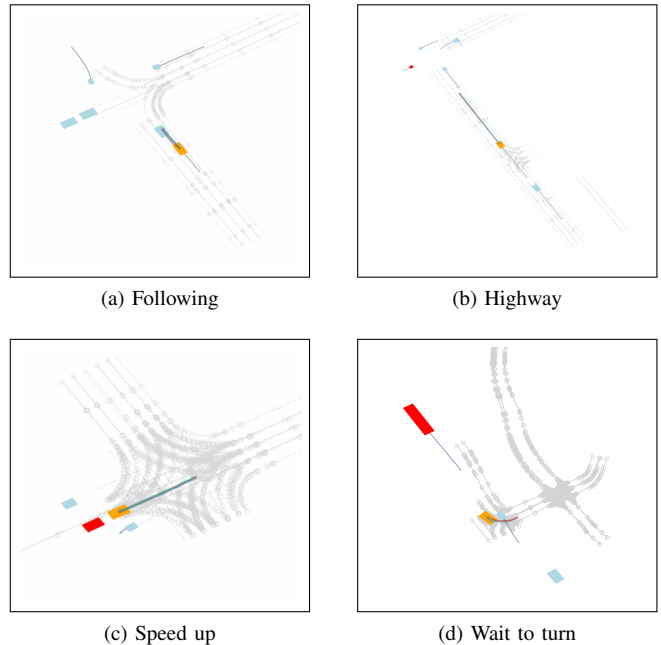


Fig. 3: Qualitative examples of Co-MTP on V2X-Seq dataset. The red box are AV, while the orange ones are the predicted targets and the blue ones are objects. The predicted trajectories are shown in green, the history ground-truth are shown in blue, and the future ground-truth are shown in brown.

eral cooperative trajectory prediction framework with multi-temporal fusion across both history and future domains. In the history domain, we develop a heterogeneous graph to learn the history fusion from multiple agents with a multi-layer Transformer to handle the incomplete history trajectory issue. To support the future planning of autonomous driving and avoid the overconfidence of the ego vehicle in the planning, we further extended the heterogeneous graph to future domains to capture the future interaction between ego planning and the other vehicles' actions from infrastructure prediction. Finally, we evaluate the Co-MTP framework on real-world dataset V2X-Seq, and the results show that Co-MTP achieves state-of-the-art performance and that both history and future fusion can greatly benefit prediction. Moreover, we investigate the robustness of V2X and test the Co-MTP performance with delays and noise.

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] K. Yuan, Y. Huang, S. Yang, Z. Zhou, Y. Wang, D. Cao, and H. Chen, “Evolutionary decision-making and planning for autonomous driving based on safe and rational exploration and exploitation,” *Engineering*, vol. 33, pp. 108–120, 2024.
- [3] Z. Zhou, Z. Yang, Y. Zhang, Y. Huang, H. Chen, and Z. Yu, “A comprehensive study of speed prediction in transportation system: From vehicle to traffic,” *iScience*, vol. 25, no. 3, 2022.
- [4] Y. Huang, J. Du, Z. Yang, Z. Zhou, L. Zhang, and H. Chen, “A survey on trajectory-prediction methods for autonomous driving,” *IEEE Transactions on Intelligent Vehicles*, vol. 7, no. 3, pp. 652–674, 2022.
- [5] S. Shi, L. Jiang, D. Dai, and B. Schiele, “Mtr++: Multi-agent motion prediction with symmetric scene modeling and guided intention querying,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [6] 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.
- [7] B. Lange, J. Li, and M. J. Kochenderfer, “Scene informer: Anchor-based occlusion inference and trajectory prediction in partially observable environments,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 14 138–14 145.
- [8] X. Han, Z. Meng, X. Xia, X. Liao, Y. He, Z. Zheng, Y. Wang, H. Xiang, Z. Zhou, L. Gao *et al.*, “Foundation intelligence for smart infrastructure services in transportation 5.0,” *IEEE Transactions on Intelligent Vehicles*, 2024.
- [9] Y. Ji, Z. Zhou, Z. Yang, Y. Huang, Y. Zhang, W. Zhang, L. Xiong, and Z. Yu, “Towards autonomous vehicles: a survey on cooperative vehicle-infrastructure system,” *iScience*.
- [10] Y. Ji, Z. Yang, Z. Zhou, Y. Huang, J. Cao, L. Xiong, and Z. Yu, “Optimization of roadside sensors placement for cooperative vehicle-infrastructure system,” in *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2023, pp. 4813–4819.
- [11] 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*, 2023.
- [12] W. Zeng, M. Liang, R. Liao, and R. Urtasun, “Lanercnn: Distributed representations for graph-centric motion forecasting,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 532–539.
- [13] H. Zhao, J. Gao, T. Lan, C. Sun, B. Sapp, B. Varadarajan, Y. Shen, Y. Shen, Y. Chai, C. Schmid *et al.*, “Tnt: Target-driven trajectory prediction,” in *Conference on Robot Learning*. PMLR, 2021, pp. 895–904.
- [14] R. Xu, H. Xiang, X. Xia, X. Han, J. Li, and J. Ma, “Opv2v: An open benchmark dataset and fusion pipeline for perception with vehicle-to-vehicle communication,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2583–2589.
- [15] H. Xiang, Z. Zheng, X. Xia, R. Xu, L. Gao, Z. Zhou, X. Han, X. Ji, M. Li, Z. Meng *et al.*, “V2x-real: a large-scale dataset for vehicle-to-everything cooperative perception,” *arXiv preprint arXiv:2403.16034*, 2024.
- [16] R. Xu, X. Xia, J. Li, H. Li, S. Zhang, Z. Tu, Z. Meng, H. Xiang, X. Dong, R. Song *et al.*, “V2v4real: A real-world large-scale dataset for vehicle-to-vehicle cooperative perception,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 13 712–13 722.
- [17] X. Cai, W. Jiang, R. Xu, W. Zhao, J. Ma, S. Liu, and Y. Li, “Analyzing infrastructure lidar placement with realistic lidar simulation library,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 5581–5587.
- [18] H. Ruan, H. Yu, W. Yang, S. Fan, Y. Tang, and Z. Nie, “Learning cooperative trajectory representations for motion forecasting,” *arXiv preprint arXiv:2311.00371*, 2023.
- [19] H. Yu, W. Yang, H. Ruan, Z. Yang, Y. Tang, X. Gao, X. Hao, Y. Shi, Y. Pan, N. Sun *et al.*, “V2x-seq: A large-scale sequential dataset for vehicle-infrastructure cooperative perception and forecasting,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 5486–5495.
- [20] Z. Huang, H. Liu, J. Wu, and C. Lv, “Differentiable integrated motion prediction and planning with learnable cost function for autonomous driving,” *IEEE transactions on neural networks and learning systems*, 2023.
- [21] J. Liu, W. Zeng, R. Urtasun, and E. Yumer, “Deep structured reactive planning,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 4897–4904.
- [22] H. Song, W. Ding, Y. Chen, S. Shen, M. Y. Wang, and Q. Chen, “Pip: Planning-informed trajectory prediction for autonomous driving,” in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*. Springer, 2020, pp. 598–614.
- [23] Z. Wu, Y. Wang, H. Ma, Z. Li, H. Qiu, and J. Li, “Cmp: Cooperative motion prediction with multi-agent communication,” *arXiv preprint arXiv:2403.17916*, 2024.
- [24] J. Araluce, A. Justo, A. Arizala, L. González, and S. Díaz, “Enhancing motion prediction by a cooperative framework,” in *2024 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2024, pp. 1389–1394.
- [25] Z. Zhou, H. Xiang, Z. Zheng, S. Z. Zhao, M. Lei, Y. Zhang, T. Cai, X. Liu, J. Liu, M. Bajji *et al.*, “V2XPnP: Vehicle-to-everything spatio-temporal fusion for multi-agent perception and prediction,” *arXiv preprint arXiv:2412.01812*, 2024.
- [26] T.-H. Wang, S. Manivasagam, M. Liang, B. Yang, W. Zeng, and R. Urtasun, “V2vnet: Vehicle-to-vehicle communication for joint perception and prediction,” in *Computer Vision – ECCV 2020*. Springer International Publishing, 2020.
- [27] J. L. V. Espinoza, A. Liniger, W. Schwarting, D. Rus, and L. Van Gool, “Deep interactive motion prediction and planning: Playing games with motion prediction models,” in *Learning for Dynamics and Control Conference*. PMLR, 2022, pp. 1006–1019.
- [28] M. Bhardwaj, B. Boots, and M. Mukadam, “Differentiable gaussian process motion planning,” in *2020 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2020, pp. 10 598–10 604.
- [29] Z. Zhou, Z. Yu, L. Xiong, D. Zeng, Z. Fu, Z. Li, and B. Leng, “A reliable path planning method for lane change based on hybrid pso-iaco algorithm,” in *2021 6th International Conference on Transportation Information and Safety (ICTIS)*. IEEE, 2021, pp. 1253–1258.
- [30] C. Tang, W. Zhan, and M. Tomizuka, “Interventional behavior prediction: Avoiding overly confident anticipation in interactive prediction,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 11 409–11 415.
- [31] 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.
- [32] N. Deo and M. M. Trivedi, “Convolutional social pooling for vehicle trajectory prediction,” in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2018, pp. 1468–1476.
- [33] 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.
- [34] Z. Zhou, Z. Wen, J. Wang, Y.-H. Li, and Y.-K. Huang, “Qcnext: A next-generation framework for joint multi-agent trajectory prediction,” *arXiv preprint arXiv:2306.10508*, 2023.
- [35] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, “Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data,” *IEEE transactions on neural networks and learning systems*, 2023.
- [36] N. Benbarka, J. Schröder, and A. Zell, “Score refinement for confidence-based 3d multi-object tracking,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 8083–8090.