

C2F-TP: A Coarse-to-Fine Denoising Framework for Uncertainty-Aware Trajectory Prediction

Zichen Wang¹, Hao Miao², Senzhang Wang^{1*}, Renzhi Wang¹, Jianxin Wang¹, Jian Zhang¹

¹Central South University

²Aalborg University

zawang0422@csu.edu.cn, haom@cs.aau.dk, szwang@csu.edu.cn,
rzwang516@gmail.com, jxwang@mail.csu.edu.cn, jjianzhang@csu.edu.cn

Abstract

Accurately predicting the trajectory of vehicles is critically important for ensuring safety and reliability in autonomous driving. Although considerable research efforts have been made recently, the inherent trajectory uncertainty caused by various factors including the dynamic driving intends and the diverse driving scenarios still poses significant challenges to accurate trajectory prediction. To address this issue, we propose C2F-TP, a coarse-to-fine denoising framework for uncertainty-aware vehicle trajectory prediction. C2F-TP features an innovative two-stage coarse-to-fine prediction process. Specifically, in the spatial-temporal interaction stage, we propose a spatial-temporal interaction module to capture the inter-vehicle interactions and learn a multimodal trajectory distribution, from which a certain number of noisy trajectories are sampled. Next, in the trajectory refinement stage, we design a conditional denoising model to reduce the uncertainty of the sampled trajectories through a step-wise denoising operation. Extensive experiments are conducted on two real datasets NGSIM and highD that are widely adopted in trajectory prediction. The result demonstrates the effectiveness of our proposal.

Code — <https://github.com/wangzc0422/C2F-TP>

Introduction

Vehicle trajectory prediction aims to predict the future trajectory of the target agent based on the historical trajectories of itself and its surrounding neighbours. Accurately predicting the future trajectory of vehicles is crucial for many autonomous driving applications, including optimal driving path planning, making accurate driving decisions in dynamic environments, and enhancing driving safety (Liu et al. 2021; Lun et al. 2024).

Traditionally, statistical models are employed to predict future trajectories based on the historical trajectories of individual agents (Hearst et al. 1998; Williams and Rasmussen 2006). Nonetheless, they do not consider the interactions between the target agent and the surrounding agents, which deteriorates the prediction performance. To address this issue, various deep learning based models (Alahi et al. 2016; Deo




Figure 1: Illustration of the dynamics and temporal correlation of inter-vehicle interactions.

and Trivedi 2018; Salzmann et al. 2020; Song et al. 2020; Zhao et al. 2021) are proposed to model the spatial interactions among vehicles. However, existing deep learning based approaches fall short in modeling the uncertainty of the trajectories (Chen et al. 2022; Wang et al. 2023), which is common in real-world driving scenarios (Liao et al. 2024). Generally, uncertainty in autonomous driving can be broadly categorized into uncertainty in driving behavior and in driving scenarios (Suk et al. 2024). How to design an uncertainty-aware trajectory prediction model is still not fully explored and remains as an open research problem.

However, achieving the uncertainty-aware trajectory prediction is non-trivial due to the following challenges. First, the inherent uncertainty in trajectories caused by various uncertainty factors in dynamic driving scenarios, make the accurate future trajectory prediction difficult (Li et al. 2023). Existing approaches mostly directly predict the coordinate points of future trajectories (Alahi et al. 2016; Xie et al. 2021) without considering such uncertainty factors. Although some recent methods are proposed for multimodal trajectory prediction (Deo and Trivedi 2018; Chen et al. 2022), they still cannot effectively model the complex uncertainty. Second, it is challenging to simultaneously model the dynamics and temporal correlations of vehicle interactions. Existing approaches either only learn a fixed interaction weight between two vehicles (Alahi et al. 2016; Deo and Trivedi 2018) or simply consider the dynamics of interactions by modeling them at each timestamp (Wang et al. 2023). For example, as shown in Figure 1, at time t , vehicle A is driving at a constant speed, and vehicles B and C intend to accelerate to merge into the middle lane. To avoid a collision, vehicle A may slow down to keep a safe distance from vehicles B and C , and vehicle C may delay the lane-changing considering vehicle B 's lane-changing behavior at

*Corresponding author

time $t + 1$. Third, it is challenging to model the multimodal driving intentions. Most existing methods (Deo and Trivedi 2018; Wang et al. 2023; Miao et al. 2024b) perform trajectory prediction by considering multiple intentions simultaneously. However, different motion intentions have diverse impacts on the future trajectories. Simply combining them together may lead to a suboptimal performance.

To address these challenges, we propose a coarse-to-fine denoising framework named C2F-TP for uncertainty-aware trajectory prediction. C2F-TP features a novel two-stage prediction process including a spatial-temporal interaction module and a refinement module. Specifically, in the spatial-temporal interaction module, C2F-TP aims to learn a multimodal future trajectory distribution based on historical trajectories. Next, a certain number of noisy trajectories are sampled from the learned distribution and fed into the refinement module. The refinement module adopts a conditional denoising model to generate stable and reliable future trajectories through stepwise denoising, greatly reducing the uncertainty of the input trajectories. In particular, in the spatial-temporal interaction module, motivated by (Wang et al. 2023) we propose an interaction pooling to simultaneously capture the dynamic and temporal correlations of interactions across vehicles. To adaptively learn the impact of different motion modes on trajectory prediction, we propose a re-weighted multi-modal trajectory predictor to embed each motion mode separately. The learned embeddings are then fused adaptively by considering their diverse importance to future trajectory prediction. Moreover, we learn a multimodal future trajectory distribution based on the fused embeddings. Finally, we sample k coarse-grained trajectories from the trajectory distribution and feed them into the refinement stage. The refinement module uses historical trajectories as a condition to refine these k coarse-grained trajectories, reducing their uncertainty through stepwise denoising and improving the accuracy of trajectory prediction.

The primary contributions of this paper are as follows.

- We for the first time propose a novel coarse-to-fine trajectory prediction model with a two-stage generation process, which considers the complex interactions across vehicles in trajectory prediction.
- To capture the complex spatial and temporal interactions, we design a spatial-temporal interaction module with a re-weighted multi-modal trajectory predictor to model the sophisticated trajectory distributions.
- To contend with the inherent stochastic nature of trajectories, we design a trajectory refinement stage based on the denoising diffusion model.

Related Work

Trajectory Prediction. Trajectory prediction attracts increasing interest due to the increasing availability of trajectory data (Krajewski et al. 2018; Deo and Trivedi 2018). Recently, significant advances have been achieved in developing deep learning based methods (Liu et al. 2021) for trajectory prediction. First, RNNs (Morton, Wheeler, and Kochenderfer 2016; Wang, Cao, and Philip 2020) have been applied

to predict the future movement paths of agents with historical trajectory data. However, in complex traffic environments, surrounding vehicles may significantly affect the prediction result due to the perception of interactions between agents. To address this issue, many recent studies (Alahi et al. 2016; Mohamed et al. 2020; Chen et al. 2022; Wang et al. 2023; Miao et al. 2024a) try to model the spatial interactions across agents by means of various emerging pooling techniques. Despite their advances, existing methods suffer from effective data uncertainty capturing.

Denoising diffusion probabilistic models. Denoising diffusion probabilistic models have achieved advanced performance on various applications (Sohl-Dickstein et al. 2015; Ho, Jain, and Abbeel 2020), such as image generation (Dhariwal and Nichol 2021) and audio synthesis (Kong et al. 2020). The diffusion model is first proposed by DPM (Sohl-Dickstein et al. 2015), which seeks to imitate the diffusion process in non-equilibrium statistical physics and reconstruct the data distribution using the denoising model. Then, a line of studies are proposed to improve the efficiency of the diffusion model by developing fast sampling techniques, such as DDPM (Ho, Jain, and Abbeel 2020) and DDIM (Song, Meng, and Ermon 2021). Recently, the diffusion model has been applied to trajectory data, which obtain good performance in trajectory-related applications, such as trajectory synthesis (Zhu et al. 2024) and trajectory prediction (Mao et al. 2023). However, these methods fall short in effectively modeling the interactions across agents.

Problem Definition and Preliminaries

Problem Definition. We consider trajectory prediction as a sequence generation problem which generates the future trajectory of an agent based on the historical trajectories of itself and surrounding agents. For a target agent, given its historical trajectories $\mathbf{X}_{\text{tar}} = \{X_{\text{tar}}^{t-T_h}, \dots, X_{\text{tar}}^{t-2}, X_{\text{tar}}^{t-1}\}$ and trajectories of surrounding N agents $\mathbb{X}_{\mathcal{N}} = \{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_i, \dots, \mathbf{X}_N\}$ over T_h time steps, we aim to predict its future trajectory $\mathbf{Y}_{\text{tar}} = \{Y_{\text{tar}}^{t+1}, Y_{\text{tar}}^{t+2}, \dots, Y_{\text{tar}}^{t+T_f}\}$ over T_f future time steps, where X_v^i is the two-dimensional coordinate (x_v^i, y_v^i) of the agent v at the i -th time step.

Preliminaries. In quantum mechanics, an entity (*e.g.*, electron, photon) is usually represented by a wave function (*e.g.*, de Broglie wave) containing both amplitude and phase (Arndt et al. 1999). Inspired by Wave-MLP (Tang et al. 2022), an agent is represented as a wave \tilde{z}_j with both amplitude and phase information, *i.e.*,

$$\tilde{z}_j = |\mathbf{z}_j| \odot e^{i\theta_j}, j = 1, 2, \dots, n, \quad (1)$$

where i is the imaginary unit, \odot is element-wise multiplication. The amplitude $|\mathbf{z}_j|$ represents the content of each agent. θ_j indicates the phase, which is the current location of an agent within a wave period. With both amplitude and phase, each agent \tilde{z}_j is represented in the complex-value domain. We view agent interaction as a superposition of waves, supposing $\tilde{z}_r = \tilde{z}_i + \tilde{z}_j$ is the aggregated results of agent




Figure 2: Interaction between two waves with different phases. Below is a superposition of two waves in the complex-valued domain, and the figure above shows how their projections along the real axis vary with phase.

wave \tilde{z}_i, \tilde{z}_j , its amplitude $|z_r|$ and phase θ_r can be calculated as follows:

$$|z_r| = \sqrt{|z_i|^2 + |z_j|^2 + 2|z_i| \odot |z_j| \odot \cos(\theta_j - \theta_i)}, \quad (2)$$

$$\theta_r = \theta_i + \text{atan2}(|z_j| \odot \sin(\theta_j - \theta_i), |z_i| + |z_j| \odot \cos(\theta_j - \theta_i)). \quad (3)$$

Inspired by WSiP (Wang et al. 2023), we consider the target and surrounding agents as waves, and use the superposition of waves to dynamically model the interactions between agents. The phase discrepancy between agents $|\theta_j - \theta_i|$ significantly modulates the aggregated amplitude of the result z_r , as depicted in Eq. (2) and illustrated in Figure 2, thereby simulating the variable influence exerted by surrounding agents on the target agent.

Methodology

Figure 3 shows the framework of C2F-TP, which is a coarse-to-fine denoising framework for uncertainty-aware vehicle trajectory prediction. C2F-TP contains a spatial-temporal interaction module and a refinement module. The former captures the dynamics and temporal correlation of inter-vehicle interactions and generates a multimodal future trajectory distribution for the target vehicle. Then we sample k noisy trajectories from the multimodal trajectory distribution. Next, the refinement module denoises the k trajectories based on a conditional denoising network and generates the final refined future trajectories.

Spatial-Temporal Interaction Module

In order to accurately capture the interactions between the target and surrounding agents, we propose the Spatial-Temporal Interaction Module to generate an initial future trajectory distribution. This module contains motion encoding, interaction pooling and re-weighted multimodal trajectory predictor. Next, we will describe the module in detail.

Motion Encoding. Following the previous approach (Deo and Trivedi 2018; Wang et al. 2023), we assume that the neighbors of the target agent (i.e., surrounding agents) are within ± 90 feet in the longitude direction and within four adjacent lanes centered on the target agent, as shown on the

left side of Figure 3. Given the historical trajectories of the target agent X_{tar} and the surrounding agents X_N , the module adopts a multi-layer perception (MLP) and a long short-term memory network (LSTM) to encode the high-level historical trajectory of the target agent enc_{tar} and the surrounding agents enc_{nbrs} , respectively.

Interaction Pooling. Interactions between agents are highly dynamic and time-dependent. Simultaneously capturing the dynamic and temporal correlations of interactions between agents helps to more accurately model the future trajectories of target agents. Motivated by WSiP (Wang et al. 2023), Interaction Pooling considers the agents in the road scene as waves, and uses the superposition of waves to dynamically model the interactions between agents. A multi-head self-attention is also adopted to capture the temporal correlations between social representations.

Inspired by Wave-MLP (Tang et al. 2022), for any agent a_i , given the hidden state of the current timestamp h_i of the agent, we use a plain fully connected layer (Plain-FC) to obtain the amplitude embedding z_i . Therefore, amplitude could represent the dynamics of the agent. Phase is used to modulate the aggregation of information from surrounding agents. We also use Plain-FC as follows to learn the phase embedding θ_i so that it can be dynamically adapted to the motion state of the agent,

$$\begin{aligned} z_i &= \text{Plain-FC}(h_i, W^z), \\ \theta_i &= \text{Plain-FC}(h_i, W^\theta), \end{aligned} \quad (4)$$

where W^z and W^θ are learnable weights. We design a Surrounding-FC module for aggregating different agent interaction information based on the token mixing module of Wave-MLP (Tang et al. 2022) as follows

$$\begin{aligned} \tilde{o}_j &= \text{Surrounding-FC}(\tilde{Z}, M_{pos}, W^t)_j, \\ j &= 1, 2, \dots, n, \end{aligned} \quad (5)$$

where $\tilde{Z} = [\tilde{z}_1, \tilde{z}_2, \dots, \tilde{z}_n]$ denotes n agent waves, M_{pos} denotes the position mask, and W^t is a learnable weight. The amplitude and phase information of different agents is aggregated by Surrounding-FC and the resulting \tilde{o}_j is a complex-valued representation of the aggregated features. Then we obtain the real-valued output o_j by summing the real and imaginary parts of \tilde{o}_j with the weights (Jacobs and Steck 2006) as follows

$$\begin{aligned} o_j &= \sum_k W_{jk}^t z_k \odot \cos \theta_k + W_{jk}^i z_k \odot \sin \theta_k, \\ j &= 1, 2, \dots, n, \end{aligned} \quad (6)$$

where W^t and W^i are learnable weights, \odot is the element-wise multiplication and j indicates the j -th output representation. In the above equation, the phase θ_k is dynamically adjusted according to the history states of various agents. Different agents with both the phase and amplitude information interact with each other through Surrounding-FC and obtain the social encoding.

The above social interaction feature is calculated for each timestamp t separately without considering the temporal

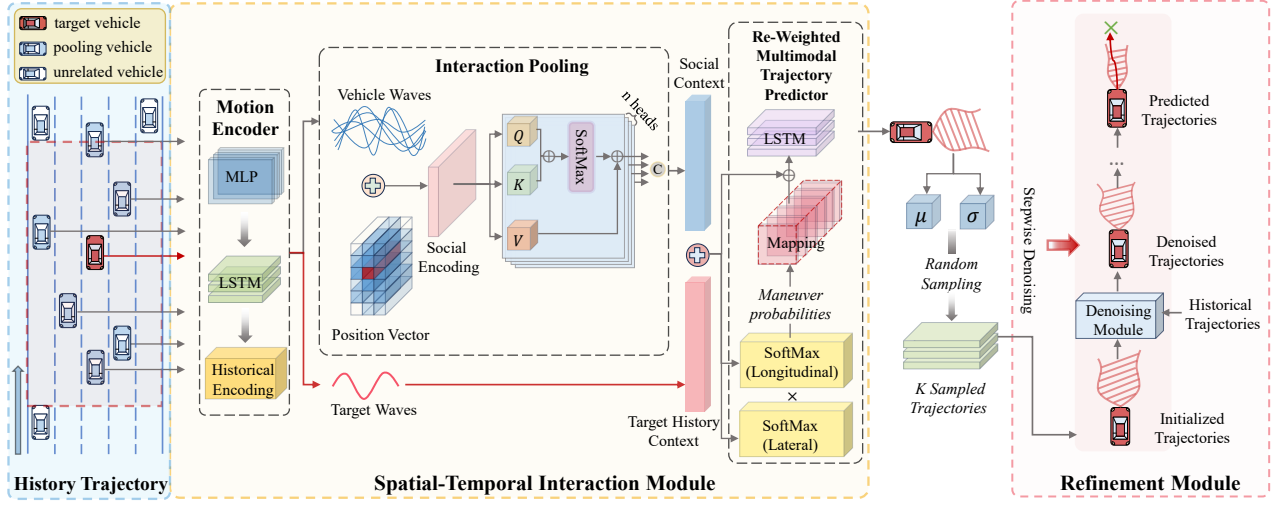


Figure 3: Framework of the proposed C2F-TP model, which contains a spatial-temporal interaction module and a refinement module. The spatial-temporal interaction module captures the dynamics and temporal correlations of inter-vehicle interactions and generates a multimodal trajectory distribution. A certain number of noisy trajectories are sampled from the multimodal trajectory distribution. The refinement module then denoises these noisy trajectories and finally generates accurate predicted trajectories.

correlation. Thus, we use a multi-head self-attention to capture the relationship between the social interaction representations across different timestamps. We denote the social encoding as H .

We take the sequence $H = (h^{t-T_h}, \dots, h^{t-2}, h^{t-1})$ as the input. Following self-attention mechanism, we calculate the query (Q), key (K), and value (V) matrices based on the sequence H . The temporal self-attention is represented as

$$\mathcal{M}_{score} = \text{softmax} \left((Q, K) / \sqrt{d_k} \right), \quad (7)$$

where $\mathcal{M}_{score} \in R^{T_h \times T_h}$ is the attentive score matrix whose entry \mathcal{M}_{score}^{ij} measures the temporal correlation between timestamps i and j . Then, the score matrix is used to aggregate temporal information from the corresponding values as follows

$$head = \mathcal{M}_{score} V. \quad (8)$$

The output of the multiple self-attention is as follows

$$Multi-Head(Q, K, V) = \text{Concat}(head_1, head_2, \dots, head_k), \quad (9)$$

where $Multi-Head(Q, K, V)$ are layer normalized to get the final social context $\tilde{H} = (\tilde{h}^{t-T_h}, \dots, \tilde{h}^{t-2}, \tilde{h}^{t-1})$ as shown in Figure 3.

Re-weighted Multimodal Trajectory Predictor. To predict the multimodal trajectory of a vehicle, we propose the Re-weighted Multimodal Trajectory Predictor to generate the future trajectory modalities \hat{Y} with probability of M . Note that $M = \{m_i \mid i = 1, 2, \dots, 6\}$ contains 6 modes, i.e., three lateral lane changing modes (turn left, turn right, and maintain lane), and two longitudinal speed shifting

modes (braking and maintaining speed). We begin by merging the social context \tilde{H} generated by the Interaction Pooling and the ego history context enc_{tar} generated by the Motion Encoder to obtain the interaction context \mathcal{C} , which is then input into this module. As shown in Figure 3, \mathcal{C} is fed to two soft-max layers to output the lateral and longitudinal maneuver probabilities $P(m_i \mid \mathbf{X})$, respectively. \mathbf{X} are the historical trajectories of agents in a scene.

Considering the future trajectories of the different modes vary, previous methods simply combine modal representations with interaction context for prediction (Deo and Trivedi 2018), which are unable to capture the relationship between modal representation pairs and interaction context at a fine-grained level. To address this problem, we design 6 mapping matrices to combine features at different historical timestamps in an adaptive manner inspired by (Chen et al. 2022).

We use softmax to generate the corresponding 6 weight matrices as follows for each modality $u_i^{t_f}$, where i denotes motion mode and t_f denotes the future prediction horizon.

$$u_i^{t_f} = \left(u_i^{-T_h, t_f}, \dots, u_i^{-2, t_f}, u_i^{-1, t_f} \right), \quad (10)$$

where $t_h = -T_h, \dots, -2, -1$ denotes the historical prediction horizons. Given the interaction context $\mathcal{C} = (c^{t-T_h}, \dots, c^{t-2}, c^{t-1})$, we use six $u_i^{t_f}$ for each mode to adaptively combine \mathcal{C}_{tar} separately as follows

$$\mathcal{V}^{t_f} = \sum_{t=-T_h}^{-1} c^t u_i^{t, t_f}. \quad (11)$$

Finally, the weighted modal mapping vectors \mathcal{V}^{t_f} are fed into an LSTM layer along with the interaction context \mathcal{C} ,

which can output the parameters of a bivariate Gaussian distribution \hat{Y} of the target agent in each maneuver modality. The predicted locations \hat{Y}_i^t at time step t are as follows

$$\hat{Y}_i^t \sim \mathcal{N}(\mu_i^t, \sigma_i^t, \rho_i^t), \quad (12)$$

where μ_i^t and σ_i^t are the means and variances of future locations respectively, and ρ_i^t is the correlation coefficient. We denote Ω as parameters of Gaussian distribution. The posterior probability of the future trajectories are as follows,

$$P(\mathbf{Y} | \mathbf{X}) = \sum_i P(m_i | \mathbf{X}) P_{\Omega}(\mathbf{Y} | m_i, \mathbf{X}). \quad (13)$$

Refinement Module

The highly dynamic interactions between vehicles make the trajectory uncertainty transmit among vehicles, increasing the challenge of accurately predicting future trajectories (Li et al. 2023). To address this issue, we sample out k noisy trajectories \hat{Y}_k based on the future trajectory distribution \hat{Y} from the first stage to model the uncertainty. In the inverse diffusion process, the refinement module reduces the uncertainty of the trajectories by progressively denoising them to generate more accurate and reliable future trajectories.

Specifically, the refinement module is a conditional denoising model that denoises the trajectory \hat{Y}_k based on historical trajectories (X_{tar} and \mathbb{X}_N) and generates optimized future trajectories by gradual denoising steps. This module contains a transformer-based context encoder $f_{context}(\cdot)$ to learn a social-temporal embedding and a noise estimation module $f_{\epsilon}(\cdot)$ to estimate the noise to reduce. The t -th denoising step works as follows

$$\begin{aligned} \chi &= f_{context}(X_{tar}, \mathbb{X}_N), \\ \epsilon_{\theta}^t &= f_{\epsilon}(\hat{Y}_k^{t+1}, \chi, t+1), \end{aligned} \quad (14)$$

$$\hat{Y}_k^t = \frac{1}{\sqrt{\alpha_t}} \left(\hat{Y}_k^{t+1} - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_{\theta}^t \right) + \sqrt{1 - \alpha_t} \mathbf{z}, \quad (15)$$

where α_t and $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$ are the parameters in the diffusion process and $\mathbf{z} \sim \mathcal{N}(\mathbf{z}; \mathbf{0}, \mathbf{I})$ is a noise. $f_{\epsilon}(\cdot)$ estimates the noise ϵ_{θ}^t in the noisy trajectory \hat{Y}_k implemented by multi-layer perceptions with the historical trajectories. Eq. (15) is a standard denoising process. This approach significantly reduces the uncertainty in trajectory prediction and further improves the accuracy of predicted trajectories.

Experiment

Dataset

The experiments are conducted on two datasets **NGSIM** (Deo and Trivedi 2018) and **highD** (Krajewski et al. 2018) that are widely adopted in trajectory prediction evaluation. The NGSIM dataset contains trajectories of real freeway traffic captured at 10 Hz over a time span of 45 minutes in 2015. The highD dataset consists of trajectories of 110,000 vehicles recorded at 25 Hz, which are collected at a segment of two-way roads around Cologne in Germany from 2017 to 2018. The dataset details are given in the associated code repository.

Evaluation Metrics

We employ Root Mean Square Error (RMSE), a widely used metric in trajectory prediction, to evaluate the spatial-temporal interaction stage as follows

$$RMSE(t) = \sqrt{\frac{1}{N} \sum_{i=1}^N (Y_i^t - \hat{Y}_i^t)^2}, \quad (16)$$

where N is the total number of test instances, Y_i^t and \hat{Y}_i^t are the ground-truth and predicted coordinates of agent a_i at time step t , respectively.

Additionally, at the end of the refinement stage, we adopt the Average Displacement Error (ADE) and Final Displacement Error (FDE) as follows, which are commonly used in denoising models, to assess the accuracy of the trajectories generated by our method.

$$\begin{aligned} ADE &= \frac{\sum_{i=1}^N \sum_{t=1}^{T_f} \|Y_i^t - \hat{Y}_i^t\|_2}{N \times T_f}, \\ FDE &= \frac{\sum_{i=1}^N \|Y_i^{T_f} - \hat{Y}_i^{T_f}\|_2}{N}. \end{aligned} \quad (17)$$

Baseline

In the spatial-temporal interaction stage, we compare the generated trajectory distribution with six baselines: *V-LSTM* (Graves and Graves 2012), *S-LSTM* (Alahi et al. 2016), *CS-LSTM* (Deo and Trivedi 2018), *PiP-noPlan* (Song et al. 2020), *STDAN* (Chen et al. 2022), and *WSiP* (Wang et al. 2023). We also compare the predicted final trajectory by the refinement module in the second stage with five state-of-the-art baselines: *V-LSTM* (Graves and Graves 2012), *S-LSTM* (Alahi et al. 2016), *CS-LSTM* (Deo and Trivedi 2018), *WSiP* (Wang et al. 2023) and *LED* (Mao et al. 2023). The detailed introduction of baselines is given in the associated code repository.

To gain insight into the effect of the key components of C2F-TP, we compare C2F-TP with its three variants.

- *w/o IP*: C2F-TP without the Interaction Pooling.
- *w/o RWMTP*: C2F-TP without the Re-weighted Multimodal Trajectory Predictor.
- *C2F-TP(C)*: C2F-TP without the Refinement module.

Implementation Details

We implement our model with the Pytorch framework on a GPU server with NVIDIA 3090 GPU. The parameters in the model are set as follows. We employ a 13×5 grid, which is defined around the target vehicle, where each column corresponds to a single lane, and the rows are separated by a distance of 15 feet. The hidden features of MLP layers are set to 32 with ReLu as the activation function. To train a coarse-to-fine framework, we consider a two-stage training strategy, where the first stage trains a denoising module and the second stage focuses on training a spatial-temporal interaction module. The details of the two-stage prediction process are given in the associated code repository. Each trajectory is split into segments over a horizon (i.e., 8s), which contains the past (3s) and future (5s) positions at 5Hz. We split the

Horizon	RMSE (NGSIM/HighD)							
	V-LSTM	S-LSTM	CS-LSTM	S-GAN	PiP-no Plan	STDAN	WSiP	C2F-TP (C)
1s	0.68/0.22	0.59/0.21	0.58/0.24	0.57/0.30	0.57/0.21	<u>0.42/0.15</u>	0.56/0.20	0.32/0.11
2s	1.66/0.65	1.29/0.65	1.27/0.68	1.32/0.78	1.24/0.62	<u>1.01/0.45</u>	1.23/0.60	0.92/0.41
3s	2.96/1.32	2.13/1.31	2.11/1.26	2.22/1.46	2.05/1.26	<u>1.69/0.94</u>	2.05/1.21	1.62/0.92
4s	4.56/2.22	3.21/2.16	3.19/2.15	3.26/2.34	3.07/2.14	<u>2.56/1.68</u>	3.08/2.07	2.44/1.64
5s	5.44/3.43	4.55/3.29	4.53/3.31	4.41/3.41	4.34/3.27	<u>3.67/2.58</u>	4.34/3.14	3.45/2.60
Average	3.06/1.57	2.35/1.52	2.34/1.53	2.36/1.66	2.25/1.50	<u>1.87/1.16</u>	2.25/1.44	1.75/1.14

Table 1: Results of the comparison between the initial trajectory distributions obtained in the first stage C2F-TP (C) and baselines on the NGSIM and highD datasets. We report the RMSE for 5s prediction horizon. Bold indicates the best result, and underline indicates the second-best result.

Dataset	Horizon	Metric (ADE/FDE) (m)						
		V-LSTM	S-LSTM	CS-LSTM	WSiP	LED	C2F-TP (C)	C2F-TP
NGSIM	1s	0.22/0.42	0.22/0.40	<u>0.21/0.40</u>	0.30/0.42	0.40/0.41	0.26/0.43	0.20/0.34
	2s	0.53/1.16	0.52/1.23	0.51/1.10	0.63/1.28	0.64/ 0.72	<u>0.50/0.95</u>	0.47/0.95
	3s	0.93/2.14	0.89/2.08	<u>0.88/1.81</u>	1.03/2.22	0.91/ 1.10	0.78/1.56	0.78/1.47
	4s	1.40/3.35	1.42/3.30	1.25/2.85	1.49/3.37	1.20/ <u>1.39</u>	<u>1.09/2.25</u>	1.08/1.35
	5s	1.94/4.76	1.89/4.37	1.69/4.08	2.03/4.71	1.52/ <u>2.02</u>	1.41/3.09	<u>1.45/1.36</u>
	Average	1.00/2.37	0.99/2.28	0.91/2.05	1.10/2.40	0.93/ <u>1.13</u>	<u>0.81/1.66</u>	0.79/1.09
HighD	1s	0.29/0.48	0.23/0.37	0.21/0.35	<u>0.15/0.23</u>	0.50/0.81	<u>0.15/0.26</u>	0.14/0.20
	2s	0.53/0.98	0.41/0.76	0.39/0.72	0.28/0.60	0.94/1.61	<u>0.27/0.49</u>	0.23/0.32
	3s	0.78/1.53	0.60/1.20	0.57/1.14	0.45/1.09	1.41/2.39	<u>0.40/0.88</u>	0.33/0.56
	4s	1.04/2.14	0.81/1.68	0.77/1.61	0.69/2.01	1.91/3.19	<u>0.57/1.47</u>	0.44/0.53
	5s	1.32/2.81	1.03/2.22	0.99/2.14	1.00/3.19	2.46/3.98	<u>0.78/2.24</u>	0.59/0.53
	Average	0.79/1.59	0.62/1.25	0.59/1.19	0.51/1.42	1.44/2.40	<u>0.43/1.07</u>	0.35/0.43

Table 2: Results of the comparison between baselines and our method (C2F-TP) on the NGSIM and highD datasets. We report two error metrics ADE and FDE for 5s prediction horizon. Bold indicates the best result, and underline indicates the second-best result.

dataset into training, validation, and testing sets with a splitting ratio of 7 : 2 : 1. Note that the baseline methods are set based on their original papers and the accompanying code.

Overall Performance Comparison

To study the effectiveness of the C2F-TP, we compare it with 6 baselines for the spatial-temporal interaction stage (C2F-TP (C)) and with 5 baselines for the refinement stage (C2F-TP), respectively. Based on the result shown in the Table 1 and Table 2, we can have the following observations.

Spatial-Temporal Interaction Stage. Table 1 shows the performance comparison among different methods on the two datasets for the spatial-temporal interaction stage. C2F-TP (C) achieves the best results among all baselines in most cases, performing better than the best among the baselines by up to 24% and 26.7% on NGSIM and highD, respectively in terms of RMSE. As a popular trajectory prediction model, STDAN performs the best among all the baselines in most cases due to its powerful ability to capture spatial and temporal correlations. Averagely, C2F-TP (C) performs better than STDAN by up to 6.4% and 2.1% on NGSIM and highD, respectively. This is because of the dynamic model-

ing of vehicle interactions with wave superposition, which facilitates trajectory prediction.

Refinement Stage. The sampled trajectories from the spatial-temporal interaction stage are fed into the refinement module. Table 2 reports the ADE and FDE results on both datasets when the prediction horizon is set to 5s. We observe that C2F-TP outperforms baselines on NGSIM in most cases, while it achieves the best performance on highD. More specifically, C2F-TP outperforms the best among the baselines by 6.7%-40.4% and 13.0%-75.2% for ADE and FDE on highD dataset, respectively. In addition, C2F-TP brings significant performance improvement on long-term trajectory prediction, especially on highD. This is because the refinement module captures the randomness and uncertainty of the trajectories, thereby increasing the accuracy and stability of the trajectory prediction through the gradual denoising process. Additionally, C2F-TP outperforms C2F-TP (C) by up to 76% due to the consideration of uncertainty, which generates more realistic predicted trajectories. We report the effect of denoising steps in the associated code repository.

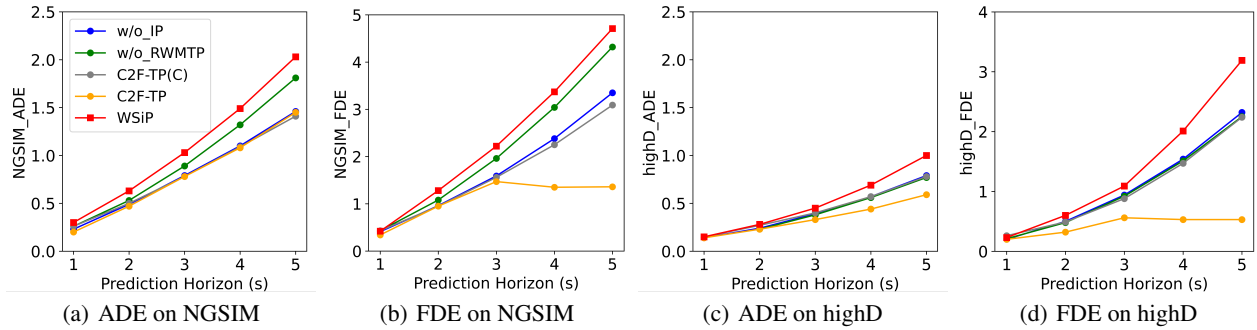


Figure 4: Comparison between C2F-TP and its variants.




Figure 5: Visualisation of S-LSTM, CS-LSTM and C2F-TP for three driving scenarios.

Ablation Study

To assess whether the components in C2F-TP all contribute to the performance of trajectory prediction, we compare C2F-TP with its three variants. The results are shown in Figure 4. We have the following observations from the result.

- Regardless of the datasets, C2F-TP always performs better than its counterparts without the interaction pooling, the re-weighted multi-trajectory predictor, and the refinement module. This shows these three components are all useful for effective trajectory prediction.
- w/o_RWMTP performs worst among all variants, which shows the importance of the re-weighted multimodal trajectory predictor. It offers evidence that the re-weighted multimodal trajectory predictor can learn comprehensive driving behaviors. C2F-TP and its variants consistently outperform WSiP, suggesting that modeling the time-dependent performance of interactions between intelligences captures a more complete characterization of the interactions and thus predicts more accurate trajectories.
- With an increase in prediction horizon, C2F-TP increasingly outperforms its variants due to its capabilities in long-term trajectory prediction.

Case Study

We visualize the trajectory prediction results for several cases in Figure 5 to intuitively show the effectiveness of C2F-TP. Three driving scenarios are selected including target keeping straight, merging to the left lane, and merging to the right lane. As shown in Figure 5(a), S-LSTM, CS-LSTM, and C2F-TP all achieve acceptable performance

when the target keeps going straight. C2F-TP can trace the right trajectory more accurately. Figure 5(b) shows that the target merges to the left lane. Both CS-LSTM and C2F-TP make correct direction predictions, but the predicted trajectory of C2F-TP is closer to the ground truth. Figure 5(c) shows the target merges to the right lane. The predictions of S-LSTM and CS-LSTM keep going straight, while C2F-TP predicts that the target will merge into the right lane. This is because of the re-weighted trajectory predictor, which considers multi-modal modeling enabling possible multiple trajectory prediction.

Conclusion

This paper proposed a coarse-to-fine trajectory prediction framework C2F-TP based on a novel two-stage generation process for trajectory prediction by considering the dynamic and temporal interactions across vehicles. Specifically, the spatial-temporal interaction stage employed an interaction pooling module to model the driving dynamics and time-dependence of inter-vehicle interactions. A re-weighted multimodal trajectory predictor was proposed to fuse specific interaction features based on specific modalities, facilitating the learning on the multimodal trajectory distributions. Then, a certain number of trajectories were sampled from the learned trajectory distribution, which was then fed into a refinement module based on a denoising diffusion model. The refinement module aimed to reduce the data uncertainty and generate the accurate trajectory predictions. An empirical study on two real datasets offered evidence of the effectiveness of C2F-TP.

Acknowledgments

This research was funded by the National Natural Science Foundation of China (No.62172443 and 62472446), Hunan Provincial Natural Science Foundation of China (No.2022JJ30053) and the High Performance Computing Center of Central South University.

References

- Alahi, A.; Goel, K.; Ramanathan, V.; Robicquet, A.; Fei-Fei, L.; and Savarese, S. 2016. Social lstm: Human trajectory prediction in crowded spaces. In *CVPR*, 961–971.
- Arndt, M.; Nairz, O.; Vos-Andraea, J.; Keller, C.; Van der Zouw, G.; and Zeilinger, A. 1999. Wave-particle duality of C60 molecules. *nature*, 401(6754): 680–682.
- Chen, X.; Zhang, H.; Zhao, F.; Hu, Y.; Tan, C.; and Yang, J. 2022. Intention-aware vehicle trajectory prediction based on spatial-temporal dynamic attention network for internet of vehicles. *TITS*, 23(10): 19471–19483.
- Deo, N.; and Trivedi, M. M. 2018. Convolutional social pooling for vehicle trajectory prediction. In *CVPR Workshop*, 1468–1476.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *NeurIPS*, 8780–8794.
- Graves, A.; and Graves, A. 2012. Long short-term memory. *Supervised sequence labelling with recurrent neural networks*, 37–45.
- Hearst, M. A.; Dumais, S. T.; Osuna, E.; Platt, J.; and Scholkopf, B. 1998. Support vector machines. *IEEE Intelligent Systems and their applications*, 13(4): 18–28.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *NeurIPS*, 6840–6851.
- Jacobs, K.; and Steck, D. A. 2006. A straightforward introduction to continuous quantum measurement. *Contemporary Physics*, 47(5): 279–303.
- Kong, Z.; Ping, W.; Huang, J.; Zhao, K.; and Catanzaro, B. 2020. Diffwave: A versatile diffusion model for audio synthesis. arXiv:2009.09761.
- Krajewski, R.; Bock, J.; Kloeker, L.; and Eckstein, L. 2018. The highd dataset: A drone dataset of naturalistic vehicle trajectories on german highways for validation of highly automated driving systems. In *ITSC*, 2118–2125.
- Li, Z.; Liang, H.; Wang, H.; Zheng, X.; Wang, J.; and Zhou, P. 2023. A multi-modal vehicle trajectory prediction framework via conditional diffusion model: A coarse-to-fine approach. *KBS*, 280: 110990.
- Liao, H.; Li, X.; Li, Y.; Kong, H.; Wang, C.; Wang, B.; Guan, Y.; Tam, K.; Li, Z.; and Xu, C. 2024. Characterized Diffusion and Spatial-Temporal Interaction Network for Trajectory Prediction in Autonomous Driving. *IJCAI*.
- Liu, J.; Mao, X.; Fang, Y.; Zhu, D.; and Meng, M. Q.-H. 2021. A survey on deep-learning approaches for vehicle trajectory prediction in autonomous driving. In *ROBIO*, 978–985.
- Lun, Y.; Miao, H.; Shen, J.; Wang, R.; Wang, X.; and Wang, S. 2024. Resisting tul attack: balancing data privacy and utility on trajectory via collaborative adversarial learning. *GeoInformatica*, 28(3): 381–401.
- Mao, W.; Xu, C.; Zhu, Q.; Chen, S.; and Wang, Y. 2023. Leapfrog diffusion model for stochastic trajectory prediction. In *CVPR*, 5517–5526.
- Miao, H.; Liu, Z.; Zhao, Y.; Zheng, K.; Zhang, Y.; and Jensen, C. S. 2024a. LightTR: A Lightweight Framework for Federated Trajectory Recovery. In *ICDE*, 4422–4434.
- Miao, H.; Zhong, X.; Liu, J.; Zhao, Y.; Zhao, X.; Qian, W.; Zheng, K.; and Jensen, C. S. 2024b. Task Assignment With Efficient Federated Preference Learning in Spatial Crowdsourcing. *TKDE*, 36(4): 1800–1814.
- Mohamed, A.; Qian, K.; Elhoseiny, M.; and Claudel, C. 2020. Social-stgcnn: A social spatio-temporal graph convolutional neural network for human trajectory prediction. In *CVPR*, 14424–14432.
- Morton, J.; Wheeler, T. A.; and Kochenderfer, M. J. 2016. Analysis of recurrent neural networks for probabilistic modeling of driver behavior. *TITS*, 18(5): 1289–1298.
- Salzmann, T.; Ivanovic, B.; Chakravarty, P.; and Pavone, M. 2020. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *ECCV*, 683–700.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2256–2265.
- Song, H.; Ding, W.; Chen, Y.; Shen, S.; Wang, M. Y.; and Chen, Q. 2020. Pip: Planning-informed trajectory prediction for autonomous driving. In *ECCV*, 598–614.
- Song, J.; Meng, C.; and Ermon, S. 2021. Denoising diffusion implicit models. *ICLR*.
- Suk, H.; Lee, Y.; Kim, T.; and Kim, S. 2024. Addressing uncertainty challenges for autonomous driving in real-world environments. In *Advances in Computers*, volume 134, 317–361.
- Tang, Y.; Han, K.; Guo, J.; Xu, C.; Li, Y.; Xu, C.; and Wang, Y. 2022. An image patch is a wave: Phase-aware vision mlp. In *CVPR*, 10935–10944.
- Wang, R.; Wang, S.; Yan, H.; and Wang, X. 2023. Wsip: Wave superposition inspired pooling for dynamic interactions-aware trajectory prediction. In *AAAI*, 4685–4692.
- Wang, S.; Cao, J.; and Philip, S. Y. 2020. Deep learning for spatio-temporal data mining: A survey. *TKDE*, 34(8): 3681–3700.
- Williams, C. K.; and Rasmussen, C. E. 2006. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA.
- Xie, X.; Zhang, C.; Zhu, Y.; Wu, Y. N.; and Zhu, S.-C. 2021. Congestion-aware multi-agent trajectory prediction for collision avoidance. In *ICRA*, 13693–13700.
- Zhao, H.; Gao, J.; Lan, T.; Sun, C.; Sapp, B.; Varadarajan, B.; Shen, Y.; Shen, Y.; Chai, Y.; Schmid, C.; et al. 2021. Tnt: Target-driven trajectory prediction. In *CoRL*, 895–904.
- Zhu, Y.; Ye, Y.; Zhang, S.; Zhao, X.; and Yu, J. 2024. Difftraj: Generating gps trajectory with diffusion probabilistic model. *NeurIPS*, 65168–65188.