

# IPCC-TP: Utilizing Incremental Pearson Correlation Coefficient for Joint Multi-Agent Trajectory Prediction

Dekai Zhu<sup>1,\*</sup>, Guangyao Zhai<sup>1,\*</sup>, Yan Di<sup>1,†</sup>, Fabian Manhardt<sup>2</sup>,  
Hendrik Berkemeyer<sup>3,4</sup>, Tuan Tran<sup>3</sup>, Nassir Navab<sup>1</sup>, Federico Tombari<sup>1,2</sup>, and Benjamin Busam<sup>1,5</sup>

<sup>1</sup> Technical University of Munich <sup>2</sup> Google

<sup>3</sup> Robert Bosch GmbH <sup>4</sup> University of Osnabrueck <sup>5</sup> 3Dwe.ai

{firstname.lastname,b.busam}@tum.de

## Abstract

Reliable multi-agent trajectory prediction is crucial for the safe planning and control of autonomous systems. Compared with single-agent cases, the major challenge in simultaneously processing multiple agents lies in modeling complex social interactions caused by various driving intentions and road conditions. Previous methods typically leverage graph-based message propagation or attention mechanism to encapsulate such interactions in the format of marginal probabilistic distributions. However, it is inherently sub-optimal. In this paper, we propose IPCC-TP, a novel relevance-aware module based on Incremental Pearson Correlation Coefficient to improve multi-agent interaction modeling. IPCC-TP learns pairwise joint Gaussian Distributions through the tightly-coupled estimation of the means and covariances according to interactive incremental movements. Our module can be conveniently embedded into existing multi-agent prediction methods to extend original motion distribution decoders. Extensive experiments on nuScenes and Argoverse 2 datasets demonstrate that IPCC-TP improves the performance of baselines by a large margin.

## 1. Introduction

Trajectory prediction refers to predicting the future trajectories of one or several target agents based on past trajectories and road conditions. It is an essential and emerging subtask in autonomous driving [22,27,36,44] and industrial robotics [20,34,47].

Previous methods [6, 14, 17, 32, 48] concentrate on Single-agent Trajectory Prediction (STP), which leverages past trajectories of other agents and surrounding road conditions as additional cues to assist ego-motion estimation. De-




Figure 1. Schematic illustration of a driving scene with five agents ①–⑤. (a) IPCC-TP models future interactions for  $t$  steps based on joint Gaussian distribution among one-dimensional increments of agents' displacement  $d$ . (b) The values of IPCC represented by  $\rho$  on the axis indicate the correlation between agents' trajectories. In the predicted future, ① and ② interact intensively with ① yielding to ② ( $\rho_{12} < 0$ ). ④'s movement does not interfere ⑤ ( $|\rho_{45}| \approx 0$ ). ③ is closely following ④ ( $\rho_{34} \leq +1$ ).

spite the significant progress made in recent years, its application is limited. The reason is that simultaneously processing all traffic participants remains unreachable, although it is a common requirement for safe autonomous driving.

A straightforward idea to deal with this issue is to directly utilize STP methods on each agent respectively and finally attain a joint collision-free prediction via pairwise collision detection. Such methods are far from reliably handling Multi-agent Trajectory Prediction (MTP) since the search space for collision-free trajectories grows exponentially as the number of agents increases, *i.e.*,  $N$  agents with  $M$  modes yield  $M^N$  possible combinations, making such search strategy infeasible with numerous agents. By in-

\* Equal contribution. † Corresponding author.

incorporating graph-based message propagation or attention mechanism into the framework to model the future interactions among multiple agents, [16,30,42] aggregate all available cues to simultaneously predict the trajectories for multiple agents in the format of a set of marginal probability distributions. While these frameworks improve the original STP baseline to some extent, the estimation based on a set of marginal distributions is inherently sub-optimal [33]. To address this issue, a predictor that can predict a joint probability distribution for all agents in the scene is necessary. An intuitive strategy for designing such a predictor is to directly predict all parameters for means and covariance matrices, which define the joint probability distributions [38]. To predict a trajectory of  $T$  time-steps with  $N$  agents and  $M$  modes up to  $4N^2TM$  covariance parameters are needed. However, such redundant parameterization fails to analyze the physical meaning of covariance parameters and also makes it hard to control the invertibility of the covariance matrix.

In this paper, we take a further step to explicitly interpret the physical meanings of the individual variances and Pearson Correlation Coefficients (PCC) inside the covariance matrix. Instead of directly modeling the locations of agents at each sampled time, we design a novel movement-based method which we term Incremental PCC (IPCC). It helps to simulate the interactions among agents during driving. IPCC models the increment from the current time to a future state at a specific step. Thereby individual variances capture the local uncertainty of each agent in the Bird-Eye-View without considering social interactions (Figure 1.a), while IPCC indicates the implicit driving policies of two agents, e.g., one agent follows or yields to another agent or the two agents' motions are irrelevant (Figure 1.b). Compared to position-based modeling, IPCC has two advantages. First, IPCC models social interactions during driving in a compact manner. Second, since position-based modeling requires respective processes along the two axes, IPCC can directly model the motion vector, which reduces memory cost for covariance matrix storage by a factor of four and facilitates deployment on mobile platforms, such as autonomous vehicles and service robots.

Based on IPCC, we design and implement IPCC-TP, a module for the MTP task. IPCC-TP can be conveniently embedded into existing methods by extending the original motion decoders to boost their performance. Experiments show that IPCC-TP effectively captures the pairwise relevance of agents and predicts scene-compliant trajectories.

In summary, the main contributions of this work are threefold:

- We introduce a novel movement-based method, namely IPCC, that can intuitively reveal the physical meaning of pairwise motion relevance and facilitate deployment by reducing memory cost.

- Based on IPCC, we propose IPCC-TP, which is compatible with state-of-the-art MTP methods and possess the ability to model future interactions with joint Gaussian distributions among multiple agents.
- Experiment results show that methods enhanced with IPCC-TP outperform the original methods by a large margin and become the new state-of-the-art methods.

## 2. Related Work

In this section, we first briefly review the general frameworks of learning-based trajectory prediction and then introduce the literature closely related to our approach.

**General trajectory prediction architecture.** Trajectory prediction for moving agents is important for the safe planning of autonomous driving [10,15,42]. One branch of trajectory prediction consists of the approaches [1,5,6,14,17,26,29,32,35,48] based on typical neural networks (CNN, GNN [12,21,25], RNN [9,19]). Cui et al. [10] propose to use agents' states and features extracted by CNN from Bird-Eye-View images to generate a multimodal prediction. Similarly, CoverNet [32] uses CNN to encode the scene context while it proposes dozens of trajectory anchors based on the bicycle model and converts prediction from position regression to anchor classification. A hierarchical GNN named VectorNet [15] encapsulates the sequential features of map elements and past trajectories with instance-wise subgraphs and models interactions with a global graph. TNT [48] and DenseTNT [17] are extensions of [15] that focus on predicting reasonable destinations. Social LSTM [1] models the trajectories of individual agents from separate LSTM networks and aggregates the LSTM hidden cues to model their interactions. CL-SGR [41] considers the sample replay model in a continuous trajectory prediction scenario setting to avoid catastrophic forgetting. The other branch [16,30,42] models the interaction among the agents based on the attention mechanism. They work with the help of Transformer [39], which achieves huge success in the fields of natural language processing [11] and computer vision [4,13,28,43,46]. Scene Transformer [30] mainly consists of attention layers, including self-attention layers that encode sequential features on the temporal dimension, self-attention layers that capture interactions on the social dimension between traffic participants, and cross-attention layers that learn compliance with traffic rules.

**Prediction of multi-agent interaction.** Most aforementioned methods are proposed for the STP task. Recently, more work has been concentrated on the MTP task. M2I [37] uses heuristic-based labeled data to train a module for classifying the interacting agents as pairs of influencers and reactors. M2I first predicts a marginal trajectory for the influencer and then a conditional trajectory for the reactor. But M2I can only deal with interactive scenes consisting

of two interacting agents. Models such as [16, 30, 42] are proposed for general multi-agent prediction. Scene Transformer [30] and AutoBots [16] share a similar architecture, while the latter employs learnable seed parameters in the decoder to accelerate inference. AgentFormer [42] leverages the sequential representation of multi-agent trajectories by flattening the state features across agents and time. Although these methods utilize attention layers intensively in their decoders to model future interactions and achieve state-of-the-art performance, their predictions are suboptimal since the final prediction is not a joint probability distribution between agents but only consists of marginal probability distributions. To the best of our knowledge, [38] is the only work that describes the future movement of all agents with a joint probability distribution. This work investigates the collaborative uncertainty of future interactions, mathematically speaking, the covariance matrix per time step. They use MLP decoders to predict mean values and the inverse of covariance matrices directly. We refer to [38] as CUM (Collaborative Uncertainty based Module) for convenience and provide methodological and quantitative comparisons between CUM and our module in Sec. 3 and 4.

### 3. Methodology

In this section, we first introduce the background knowledge of MTP. Then we propose the IPCC method, which models the relevance of the increments of target agents. Based on this, we design IPCC-TP and illustrate how this module can be plugged into the state-of-the-art models [16, 42]. In the end, we compare our module with CUM.

#### 3.1. Preliminary

**Problem Statement.** Given the road map  $\mathcal{G}$  and  $N$  agents with trajectories of past  $T_{obs}$  steps, MTP is formulated as predicting future trajectories of all agents for  $T$  steps. For observed time steps  $t \in (-T_{obs}, 0]$ , the past states at  $t$  step are  $\mathcal{H}_t = \{\mathcal{H}_i^t | i = 1, \dots, N\}$ , where  $\mathcal{H}_i^t$  contains the state information, *i.e.*  $\{x, y\}$  position of agent  $i$  at this time step. And the past trajectories are denoted as  $\mathcal{H} = \{\mathcal{H}_t | -T_{obs} < t \leq 0\}$ . Similarly, for future time steps  $t \in (0, T]$ , the future states at step  $t$  are  $\mathcal{F}_t = \{\mathcal{F}_i^t | i = 1, \dots, N\}$ , where  $\mathcal{F}_i^t$  contains  $\{x, y\}$  position of agent  $i$  at this step. And the future trajectories are  $\mathcal{F} = \{\mathcal{F}_t | 0 < t \leq T\}$ . Considering that a specific current state could develop to multiple interaction patterns, most MTP methods predict  $M$  interaction modes in the future. For simplicity, we only analyze the cases of  $M = 1$  in this section and our method can be conveniently extended to predict multiple modes.

**Marginal Probability Modeling.** In learning-based MTP methods, the training objective is to find the optimal parameter  $\theta^*$  to maximize the Gaussian likelihood  $P$ :

$$\theta^* = \arg \max_{\theta} P(\mathcal{F} | \mathcal{H}, \mathcal{G}, \theta). \quad (1)$$

Thereby the core of MTP lies in constructing an effective joint probability distribution  $P$  to model the generative future trajectories of all agents. For simplicity, existing methods such as AutoBots [16] and AgentFormer [42] approximate the high-dimensional multivariate distribution  $P = \{P_i^t\}$  in Eq. (1) by aggregating multiple marginal probability distributions. Each distribution  $P_i^t$  describes  $i$ -th agent's position at a specific time  $t$ . The objective function Eq. (1) is then transformed as:

$$\theta^* = \arg \max_{\theta} \sum_{t=1}^T \sum_{i=1}^N P_i^t(\mathcal{F}_i^t | \mathcal{H}, \mathcal{G}, \theta), \quad (2)$$

where the marginal prediction  $P_i^t \sim \mathcal{N}(\mu_i^t, \Sigma_i^t)$  is a 2D Gaussian distribution defined by:

$$\begin{aligned} \mu_i^t &= [\mu_i^{tx}, \mu_i^{ty}] \\ \Sigma_i^t &= \begin{bmatrix} (\sigma_i^{tx})^2 & \rho_i^{txy} \sigma_i^{tx} \sigma_i^{ty} \\ \rho_i^{txy} \sigma_i^{tx} \sigma_i^{ty} & (\sigma_i^{ty})^2 \end{bmatrix}, \end{aligned} \quad (3)$$

where  $\mu_i^t$  denotes the mean position of the  $i$ -th agent at time step  $t$ ,  $\Sigma_i^t$  is its corresponding covariance matrix.  $\rho_i^{txy}, \rho_i^{tyx}$  in  $\Sigma_i^t$  are the PCCs of agent  $i$  on  $x, y$  axis respectively.

However, such marginal approximation is inherently sub-optimal [33], resulting in unsatisfactory results on real-world benchmarks [3, 8, 23, 24, 31]. Although existing methods typically employ attention-based mechanism to help each agent attends to its interactive counterparts, social interactions in motion between different pairs of agents is in fact not explicitly analyzed and modeled, which limits the capability of networks to predict scene compliant and collision-free trajectories.

#### 3.2. PCC: Position Correlation Modeling

The joint probability distribution among multiple agents in the scene contains correlations between the motions of different agents, which are essential parts of precisely modeling future interactions. Theoretically, using such a scene-level modeling method would lead to a more optimal estimation of future interactions [33] than stacking a set of marginal probability distributions to approximate the joint distribution.

For time step  $t$ , the joint prediction among agents in the scene  $\hat{\mathcal{F}}_t = P_t(\cdot | \mathcal{H}, \mathcal{M}, \theta) \sim \mathcal{N}(M_t, \Sigma_t)$  is defined by:

$$M_t = [\mu_1^t, \dots, \mu_N^t], \mu_i^t = [\mu_i^{tx}, \mu_i^{ty}], \quad (4)$$




Figure 2. The movement from position  $\mu_i^0$  to  $\mu_i^t$  in the x-y coordinate system corresponds to an incremental movement of  $\mu_i^{t\delta}$  in IPCC method (solid black line: incremental movement, block dash line: actual trajectory).  $\theta_i^t$  is the approximate yaw angle, while  $\phi_i^t$  is the actual yaw angle. When  $\theta_i^t$  is close to  $\phi_i^t$ , the prediction  $\hat{\mathcal{F}}_t^*$  derived from IPCC method is an approximation to  $\hat{\mathcal{F}}_t$  in the x-y coordinate system.

$$\Sigma_t = \begin{bmatrix} \sigma_{11}^t & \cdots & \sigma_{1N}^t \\ \vdots & \ddots & \vdots \\ \sigma_{N1}^t & \cdots & \sigma_{NN}^t \end{bmatrix}, \quad (5)$$

$$\sigma_{ij}^t = \begin{bmatrix} \rho_{ij}^{txx} \sigma_i^{tx} \sigma_j^{tx} & \rho_{ij}^{txy} \sigma_i^{tx} \sigma_j^{ty} \\ \rho_{ij}^{tyx} \sigma_i^{ty} \sigma_j^{tx} & \rho_{ij}^{tyy} \sigma_i^{ty} \sigma_j^{ty} \end{bmatrix},$$

where  $M_t \in \mathbb{R}^{2N}$  contains the mean positions of all agents, and  $\Sigma_t \in \mathbb{R}^{2N \times 2N}$  is the corresponding covariance matrix, and  $\rho_{ij}^{txx} = \rho_{ij}^{tyy} = 1$  when  $i = j$ . The objective function Eq. (1) then turns into:

$$\theta^* = \arg \max_{\theta} \sum_{t=1}^T P_t(\mathcal{F}_t | \mathcal{H}, \mathcal{G}, \theta). \quad (6)$$

Thereby, a straightforward strategy for designing a predictor based on such joint probability modeling is to directly predict all parameters in means and covariance matrices. However, such redundant parameterization still fails to capture the interaction patterns in motion between agents since PCCs inside the covariance matrix reflect the *agent-to-agent* position correlations rather than motion correlations. Moreover, such strategy also makes it hard to control the invertibility of the covariance matrix.

### 3.3. IPCC: Motion Interaction Modeling.

Since the joint Gaussian distribution defined by Eq. (4) and Eq. (5) is not feasible in real applications, we propose a novel motion interaction modeling method, IPCC. For Eq. (5), we observe that each pair of agents requires estimations of PCCs:  $\{\rho_{ij}^{txx}, \rho_{ij}^{txy}, \rho_{ij}^{tyx}, \rho_{ij}^{tyy}\}$ . These representations are redundant since they respectively model the

pairwise position correlation between agent  $i$  and  $j$  along the 2D axis from the bird-eye view. In the following part, we demonstrate that this correlation can be merged into a single motion-aware parameter IPCC by modeling the one-dimensional increments from current time step ( $t=0$ ) to any future time step.

**From PCC to IPCC.** For future time step  $t > 0$ , we predict the increments  $\hat{\mathcal{F}}_{\Delta}^t$  as shown in Fig. 2, which is the one-dimensional displacements with respect to the current positions  $M_0$ , for all agents in the scene.  $\hat{\mathcal{F}}_{\Delta}^t \sim \mathcal{N}(M_{\Delta}^t, \Sigma_{\Delta}^t)$  is defined by:

$$M_{\Delta}^t = [\mu_1^{t\delta}, \dots, \mu_N^{t\delta}], \mu_i^{t\delta} \in \mathbb{R}_0^+ \quad (7)$$

$$\Sigma_{\Delta}^t = P_{\Delta}^t \odot \begin{bmatrix} (\sigma_1^{t\delta})^2 & \cdots & \sigma_1^{t\delta} \sigma_N^{t\delta} \\ \vdots & \ddots & \vdots \\ \sigma_N^{t\delta} \sigma_1^{t\delta} & \cdots & (\sigma_N^{t\delta})^2 \end{bmatrix}, \quad (8)$$

$$P_{\Delta}^t = \begin{bmatrix} 1 & \cdots & \rho_{1N}^{t\delta} \\ \vdots & \ddots & \vdots \\ \rho_{N1}^{t\delta} & \cdots & 1 \end{bmatrix},$$

where  $P_{\Delta}^t, \Sigma_{\Delta}^t \in \mathbb{R}^{N \times N}$ .  $\Sigma_{\Delta}^t$  denotes the covariance matrix and IPCC matrix  $P_{\Delta}^t$  can be extracted from  $\Sigma_{\Delta}^t$  as in Eq. (8).  $P_{\Delta}^t$  preserves all pairwise correlation parameters for the increments.  $\odot$  denotes element-wise multiplication. Thereby  $\rho_{ij}^{t\delta} \in [-1, +1]$  naturally encapsulates the motion relevance between  $i$ -th and  $j$ -th agents, which conveys various interaction patterns: **(1)** The more it approaches +1, the more similar behaviours the two agents share, e.g., two vehicles drive on the same lane towards the same direction. **(2)** The more it approaches -1, the more one vehicle is likely to yield if the other one is driving aggressively, e.g., two vehicles are at a merge. **(3)** The closer it is to 0 the less two agents impact each other. Such explicit modeling of motion interactions enforces the network to effectively predict the trajectory of each agent by attending to different interaction patterns from other agents.

Since we aim to build a plug-in module that can be conveniently embedded into other existing methods, we project the predicted IPCC parameters back to the original  $x, y$  axis. Therefore, the training losses of baselines can be directly inherited and extended for supervision.

**IPCC Projection.** Suppose that for agent  $i$ , the projections of mean displacement  $\mu_i^{t\delta}$  on x- and y-axis, namely  $\mu_i^{t\delta_x}$  and  $\mu_i^{t\delta_y}$ , are already obtained, we can estimate the approximate yaw angle for agent  $i$  at step  $t$  via

$$\theta_i^t = \arctan(\mu_i^{t\delta_y}, \mu_i^{t\delta_x}). \quad (9)$$

Similarly,  $\Theta_t = [\theta_1^t, \dots, \theta_N^t]$  at time  $t$  is available. The reconstructed future positions  $\hat{\mathcal{F}}_t^*$  in the global coordinate

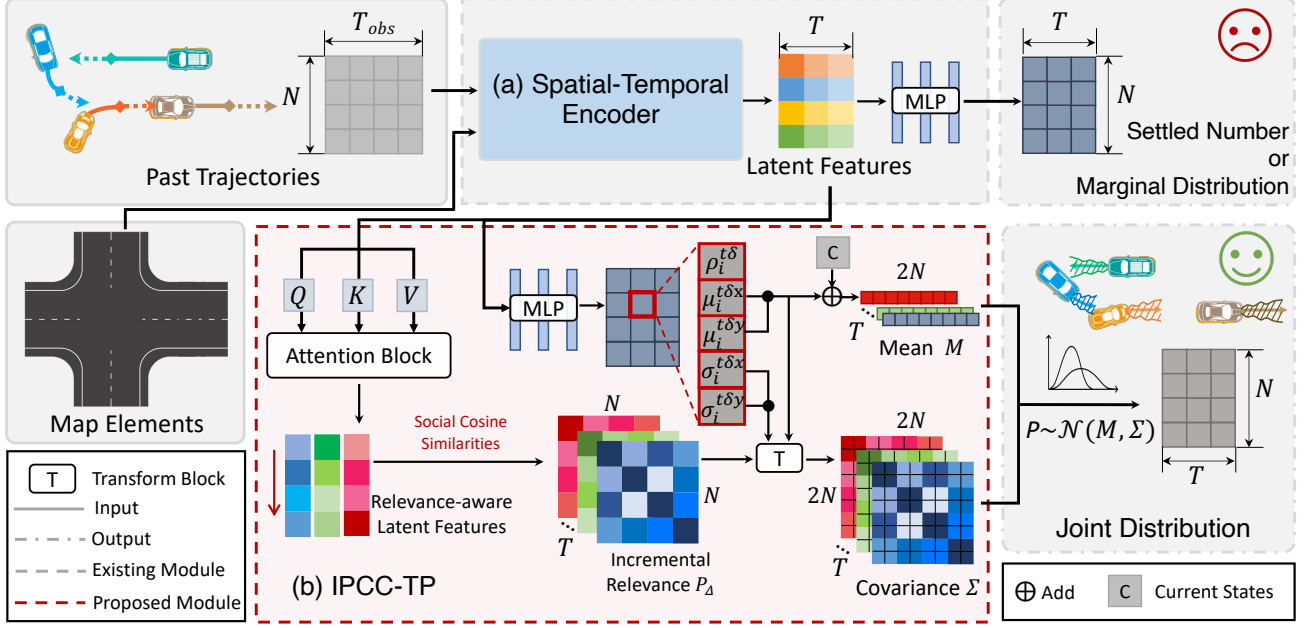


Figure 3. The overall framework of **IPCC-TP**. The inputs of existing MTP models include past trajectories and map elements. Latent features extracted by (a) Spatial-Temporal encoder are passed to (b) IPCC-TP to predict the future interactions represented by joint Gaussian distributions among the agents. IPCC-TP consists of an *Attention Block* and a *Transform Block*. *Attention Block* maps the original latent features to relevance-aware latent features, which are later used for calculating the incremental relevance  $P_{\Delta}$  via cosine similarity measurement. The mean positions are obtained by adding the current positions of agents and the mean values in marginal distributions. Using incremental relevance  $P_{\Delta}$  and  $\mu_i^{tx}, \mu_i^{ty}, \sigma_i^{tx}, \sigma_i^{ty}$  in marginal distributions, *Transform Block* estimates the yaw angle and reconstructs the covariance matrices in the x-y coordinate system. Defined by mean position  $M$  and covariance matrices  $\Sigma$ , the joint Gaussian distributions  $P \sim \mathcal{N}(M, \Sigma)$  for modeling the future interactions among the agents for  $T$  steps are predicted by IPCC-TP.

can be expressed as a sum of the current states  $M_{t_0}$  and increments.  $\hat{\mathcal{F}}_t^* \sim \mathcal{N}(M_t^*, \Sigma_t^*)$  is defined by

$$M_t^* = C \cdot M_{\Delta_r}^t + M_{t_0},$$

$$M_{\Delta_r}^t = [\mu_1^{t\delta}, \mu_1^{t\delta}, \dots, \mu_N^{t\delta}, \mu_N^{t\delta}], \quad (10)$$

$$C = \text{diag}([\cos \theta_1^t, \sin \theta_1^t, \dots, \cos \theta_N^t, \sin \theta_N^t]),$$

and

$$\Sigma_t^* = C^T \Sigma_{\Delta_r}^t C = \begin{bmatrix} \sigma_{11}^{t\delta} & \dots & \sigma_{1N}^{t\delta} \\ \vdots & \ddots & \vdots \\ \sigma_{N1}^{t\delta} & \dots & \sigma_{NN}^{t\delta} \end{bmatrix},$$

$$\Sigma_{\Delta_r}^t = \begin{bmatrix} (\sigma_1^{t\delta})^2 & (\sigma_1^{t\delta})^2 & \dots & \sigma_1^{t\delta} \sigma_N^{t\delta} & \sigma_1^{t\delta} \sigma_N^{t\delta} \\ (\sigma_1^{t\delta})^2 & (\sigma_1^{t\delta})^2 & \dots & \sigma_1^{t\delta} \sigma_N^{t\delta} & \sigma_1^{t\delta} \sigma_N^{t\delta} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \sigma_N^{t\delta} \sigma_1^{t\delta} & \sigma_N^{t\delta} \sigma_1^{t\delta} & \dots & (\sigma_N^{t\delta})^2 & (\sigma_N^{t\delta})^2 \\ \sigma_N^{t\delta} \sigma_1^{t\delta} & \sigma_N^{t\delta} \sigma_1^{t\delta} & \dots & (\sigma_N^{t\delta})^2 & (\sigma_N^{t\delta})^2 \end{bmatrix}, \quad (11)$$

$$\sigma_{ij}^{t\delta} = \rho_{ij}^{t\delta} \sigma_i^{t\delta} \sigma_j^{t\delta} \begin{bmatrix} \cos \theta_i^t \cos \theta_j^t & \cos \theta_i^t \sin \theta_j^t \\ \sin \theta_i^t \cos \theta_j^t & \sin \theta_i^t \sin \theta_j^t \end{bmatrix},$$

where  $M_{\Delta_r}^t$  and  $\Sigma_{\Delta_r}^t$  are the replicated augments of  $M_{\Delta}^t$  and  $\Sigma_{\Delta}^t$  respectively, where  $M_{\Delta_r}^t \in \mathbb{R}^{2N}$  and  $\Sigma_{\Delta_r}^t, C \in \mathbb{R}^{2N \times 2N}$ . Note that when the approximate yaw angles  $\Theta_t = [\theta_1^t, \dots, \theta_N^t]$  are equal to the actual yaw angles  $\Phi_t = [\phi_1^t, \dots, \phi_N^t]$ ,  $\hat{\mathcal{F}}_t^*$  and  $\hat{\mathcal{F}}_t$  are identical distributions, which leads to:

$$M_t^* = M_t, \quad \Sigma_t^* = \Sigma_t, \quad (12)$$

and

$$\sigma_{ij}^t = \sigma_{ij}^{t\delta}. \quad (13)$$

According to Eq. (5), (11), (12), and (13), we obtain  $\rho_{ij}^{txx}, \rho_{ij}^{txy}, \rho_{ij}^{tyx}, \rho_{ij}^{tyy}$  with the signum function  $\text{sgn}(\cdot)$ ,

$$\begin{bmatrix} \rho_{ij}^{txx} & \rho_{ij}^{txy} \\ \rho_{ij}^{tyx} & \rho_{ij}^{tyy} \end{bmatrix} = \rho_{ij}^{t\delta} \cdot \text{sgn} \left( \begin{bmatrix} \cos \theta_i^t \cos \theta_j^t & \cos \theta_i^t \sin \theta_j^t \\ \sin \theta_i^t \cos \theta_j^t & \sin \theta_i^t \sin \theta_j^t \end{bmatrix} \right) \quad (14)$$

So far, we have shown that when the approximate yaw angles are equal to the actual yaw angles,  $\rho_{ij}^{txx}, \rho_{ij}^{txy}, \rho_{ij}^{tyx}, \rho_{ij}^{tyy}$  can be obtained via Eq. (14). For short-term trajectory prediction (no longer than 6s), the vehicles are unlikely to have a sharp turn in such a short time, thus the angle  $\theta_i^t$  based on the incremental movement is



close to the actual yaw angle  $\phi_i^t$ , and  $\hat{\mathcal{F}}_t^*$  is a suitable approximation to  $\hat{\mathcal{F}}_t$ .

### 3.4. IPCC-TP: a Plugin Module

Based on the IPCC method, we design and implement a plugin module named IPCC-TP, which can be easily inserted into existing models. The pipeline of IPCC-TP is shown in Figure. 3. Previous methods use spatial-temporal encoders to extract latent features  $\mathcal{L} \in \mathbb{R}^{N \times T \times d}$ , which are subsequently sent to MLP decoders to generate deterministic predictions [42] or marginal distributions [16] (marked with the sad face). In contrast, our IPCC-TP first sends the latent features  $\mathcal{L}_t \in \mathbb{R}^{N \times d}$  at time step  $t$  to the *Attention Block*. This block is designated for modeling the future interactions via which the latent features  $\mathcal{L}_t$  are mapped to the relevance-aware latent features  $\mathcal{L}_t^{rel} \in \mathbb{R}^{N \times d}$ . Considering PCC is in the range of [-1, 1], we compute pairwise motion relevance based on cosine similarity. In this way, if agent  $i$  and agent  $j$  share a higher motion relevance, their relevance-aware feature,  $\mathcal{L}_{it}^{rel}, \mathcal{L}_{jt}^{rel} \in \mathbb{R}^d$ , will have a higher cosine similarity. For future time step  $t > 0$ , IPCC-TP computes the social cosine similarities as  $P_{\Delta}^t$  in Eq. (8) for each pair of agents. Then with the approximated yaw angles  $\Theta_t$ ,  $P_{\Delta}^t$  is used to predict the  $\rho$  parameters in  $\Sigma^t$  according to Eq. (14). Once the  $\rho$  parameters in  $\Sigma^t$  are obtained, together with  $\{\mu_i^{tx}, \mu_i^{ty}, \sigma_i^{tx}, \sigma_i^{ty}, \rho_i^{txy}\}$ , which are already available in the marginal probability distributions, we can extend the marginal probability distributions to a joint Gaussian distribution  $\hat{\mathcal{F}}_t$  defined in Eq. (4) and Eq. (5). As shown in Figure. 2, the approximated yaw angles are computed via Eq. (9). In order to stabilize the training, we use Tikhonov Regularization [18] in IPCC-TP by adding a small value  $\Delta_{reg}$  to the main diagonal of covariance matrices  $\Sigma_t$  to obtain  $\Sigma_t^{reg}$ ,

$$\Sigma_t^{reg} = \Sigma_t + \Delta_{reg} \cdot I, \quad (15)$$

where  $I \in \mathbb{R}^{2N \times 2N}$ , an identity matrix.

Thereby we adopted the following loss functions for training, following [16, 42]. In the training of the IPCC-TP embedded models, we replace the negative log-likelihood (NLL) loss of [16, 42] with the scene-level NLL loss. The scene-level NLL loss is calculated as

$$\mathcal{L}_{NLL-scene} = \sum_{t=1}^T 0.5[\ln(|\Sigma_t|) + (\mathcal{F}_t - M_t)^T \Sigma_t^{-1} (\mathcal{F}_t - M_t) + 2N \ln(2\pi)]. \quad (16)$$

Thus the loss function for training the enhanced AgentFormer is

$$\mathcal{L}^{Ag} = \mathcal{L}_{kld} + \mathcal{L}_{sample} + \mathcal{L}_{NLL-scene}, \quad (17)$$

where  $\mathcal{L}_{kld}$  term is the Kullback-Leibler divergence for the CVAE [2] in AgentFormer,  $\mathcal{L}_{sample}$  term is the mean square

error between the ground truth and the sample from the latent  $z$  distribution. And the loss function for training the enhanced AutoBots is

$$\mathcal{L}^{Ab} = \mathcal{L}_{ME} + \mathcal{L}_{kld}^* + \mathcal{L}_{NLL-scene}, \quad (18)$$

where  $\mathcal{L}_{ME}$  is a mode entropy regularization term to penalize large entropy of the predicted distribution,  $\mathcal{L}_{kld}^*$  is the Kullback-Leibler divergence between the approximating posterior and the actual posterior.

### 3.5. Methodology Comparison to CUM

CUM [38] is another work that tries to predict the joint probabilistic distribution for the MTP task. Compared with our IPCC-TP, the main difference lies in the interpretation and implementation of the covariance matrix. In order to avoid the inverse calculation for covariance matrices  $\Sigma_t$ , CUM uses MLP decoders to predict all the parameters in  $\Sigma_t^{-1}$ . In contrast, our module instead predicts  $\rho$  parameters in  $P_{\Delta}^t$ , which can intuitively encapsulate the physical meaning of motion relevance in future interactions, and then predicts  $\Sigma_t$  based on Eq. (14). Experiments demonstrate the superiority of IPCC-TP.

## 4. Experiments

### 4.1. Experimental Settings

**Datasets.** We apply IPCC-TP to AgentFormer [42] and AutoBots [16] and evaluate the performance on two public datasets, *nuScenes* [3] and *Argoverse 2* [40]. *nuScenes* collects 1K scenes with abundant annotations for the research of autonomous driving. For the prediction task, observable past trajectories are 2s long, and predicted future trajectories are 6s long. The sample rate is 2 Hz. *Argoverse 2* contains 250K scenes in total. In each scene, observable past trajectories are 5s long, and predicted future trajectories are 6s long. The sample rate is 10 Hz. Considering the vast number of scenes in *Argoverse 2*, we randomly pick 8K scenes for training and 1K for testing.

**Implementation Details.** For both AgentFormer and AutoBots, we use the official implementations as the baselines in our experiments. For IPCC-TP, we implement the attention block inside it with a self-attention layer followed by a 2-layer MLP. We set the Tikhonov regularization parameter  $\Delta_{reg}$  as 1e-4 in the experiments. More details of the training are provided in the supplementary material.

**Metrics.** For evaluation, considering that the widely used metrics minADE and minFDE are designed for the ego-motion prediction in STP task, we choose to use *Minimum Joint Average Displacement Error (minJointADE)* and *Minimum Joint Final Displacement Error (minJointFDE)* in the multi-agent prediction task of INTERPRET Challenge [7, 45]. These two metrics are proposed specifically

for the MTP task; their calculation formula is listed in the supplementary material. According to the rule of challenges on the nuScenes and Argoverse 2 benchmarks, we set the number of modes  $M$  as 5 in the evaluation on nuScenes and 6 in the evaluation on Argoverse 2.

## 4.2. Training Details

**AgentFormer.** In the evaluation on nuScenes, we use the official pre-trained model as the baseline and the default hyperparameters for training the enhanced model. In the evaluation on Argoverse 2 dataset, since there are much more future steps and agents, we set the dimension as 24 and the number of heads as 4 in the attention layers to reduce VRAM cost. In both evaluations, we use Adam optimizer with a learning rate of  $1.0e-4$ .

**AutoBots.** In the evaluation on nuScenes, we use the default hyperparameters for AutoBots baseline and the enhanced AutoBots. In the evaluation on Argoverse 2 dataset, we set the dimension as 32 in the attention layers also for reducing the VRAM cost. In both evaluations, we use Adam optimizer with a learning rate of  $7.5e-4$ .

## 4.3. Experimental Results

We report and analyze the qualitative and quantitative results against AgentFormer and AutoBots on nuScenes and Argoverse 2. After that, we compare the performance of CUM and IPCC-TP, and illustrate why our module can surpass CUM on the baseline. More results can be found in the supplementary material.

**Evaluation on nuScenes.** We enhance AgentFormer and AutoBots with IPCC-TP for the evaluation on nuScenes. We use the official pre-trained AgentFormer as a baseline, which is reported to perform minADE(5) of  $1.86m$  and minFDE(5) of  $3.89m$  on nuScenes. The experiment results are listed in Table 1 from which we can observe that IPCC-TP improves the original AgentFormer and AutoBots by  $0.33m$  (4.22%) and  $0.23m$  (3.54%) regarding minJointFDE(5) respectively. The bottom part of Figure 4 shows a case of comparisons between the original AgentFormer and the AgentFormer embedded with IPCC-TP, which illustrates that AgentFormer boosted by IPCC-TP can predict more compliant and logical future interactions with fewer collisions than the original one can do.

In Sec. 3.3, we have illustrated how  $\rho$  in  $P_{\Delta}^t$  (8) represents pairwise motion relevance. Here we show visualizations of  $P_{\Delta}^t$  in different interaction patterns in Figure 5. In Scene 1, Vehicle ① is entering the intersection while Vehicle ②–④ have almost left the intersection, and they are driving with a similar speed towards the same direction. IPCC-TP predicts  $P_{\Delta}^t$  consisting of all positive values, and it also predicts a larger correlation between two vehicles that are closer to each other, e.g., ① has a larger correlation with ②

Model	minJ-ADE(5)	minJ-FDE(5)
AgentFormer [42]	3.47	7.82
AgentFormer+IPCC-TP	<b>3.29</b>	<b>7.49</b>
AutoBots [16]	3.43	6.50
AutoBots+IPCC-TP	<b>3.22</b>	<b>6.27</b>

Table 1. Evaluation on nuScenes.

Model	minJ-ADE(6)	minJ-FDE(6)
AgentFormer [42]	2.74	5.88
AgentFormer+IPCC-TP	<b>2.10</b>	<b>4.64</b>
AutoBots [16]	2.72	5.76
AutoBots+IPCC-TP	<b>2.20</b>	<b>4.79</b>

Table 2. Evaluation on Argoverse 2.

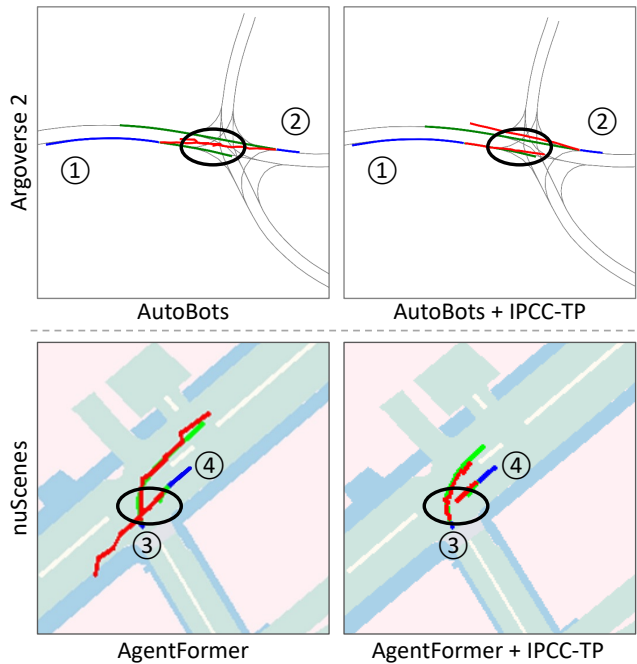


Figure 4. Comparisons between baselines and baselines + IPCC-TP on Argoverse 2 and nuScenes (blue line: past trajectories, green line: ground truth, red line: predictions). **Top:** Vehicle ① and ② drive towards different directions, and they should keep their lanes without interfering with each other. **Below:** Vehicle ③ and ④ drive aggressively, and ④ is about to stop. Black circles show that the results of baselines are abnormal, while they become reasonable after plus IPCC-TP.

than the correlation with ④. In Scene 2, IPCC-TP predicts a negative correlation between two vehicles, which describes the interaction mode that ② is yielding to ①.

**Evaluation on Argoverse 2.** We repeat the experiment on Argoverse 2 dataset with AgentFormer and AutoBots. The

Model	minJ-ADE(6)	minJ-FDE(6)
AutoBots+CUM [38]	2.25	4.85
AutoBots+IPCC-TP	<b>2.20</b>	<b>4.79</b>

Table 3. Comparison to CU-based module on Argoverse 2

Model	minJ-ADE(5)	minJ-FDE(5)
AgentFormer+CUM [38]	3.43	7.75
AgentFormer+IPCC-TP	<b>3.29</b>	<b>7.49</b>

Table 4. Comparison to CU-based module on nuScenes

experiment results are listed in Table. 2. In this evaluation, IPCC-TP improves the original AgentFormer and AutoBots by  $1.24m$  (**21.09%**) and  $0.97m$  (**16.84%**) regarding minJointFDE(6) respectively. One case for comparison between the original AutoBots and the AutoBots embedded with IPCC-TP is illustrated at the top part of Figure. 4. In this instance, AutoBots plus IPCC-TP predicts reasonable and collision-free interactions. In contrast, the predicted trajectories from the original AutoBots deviate significantly from the center lines, resulting in two vehicles crashing into each other. The prediction of collisions also reveals that the initial baseline is inefficient in modeling future interactions.

**Quantitative Comparison to CUM.** We have compared CUM with our IPCC-TP methodologically in Sec. 3.5. In this subsection, we conduct a quantitative comparison between these two plugin modules on both Argoverse 2 and nuScenes using AutoBots and AgentFormer as baselines respectively. The experiment results are listed in Table. 3 and Table. 4. In the evaluation, IPCC-TP is superior to CUM by **1.24%** on AutoBots and by **3.35%** on AgentFormer regarding minJointFDE. The reason is that, compared with CUM, IPCC-TP not only models the future interaction more intuitively but also boosts the performance of MTP models more efficiently with less memory cost.

#### 4.4. Ablation Study

Training the enhanced models involves the inverse computation of  $\Sigma_t$  in (16). In order to improve the robustness of the inverse computation, we use Tikhonov Regularization (15) on  $\Sigma_t$  during training the models in the aforementioned experiments. In this subsection, we conduct an ablation experiment on the magnitude of Tikhonov Regularization parameter  $\Delta_{reg}$  (15). Based on AutoBots, we repeat training the enhanced models on Argoverse 2 with  $\Delta_{reg}$  in different magnitudes. The experiment results are listed in Table. 5. Without any regulation, the training fails at the first iteration. When the regulation parameter is  $1e-5$ , the training can run for the first epoch but still fails soon. In this experiment,  $1e-4$  is the most appropriate regulation we found for IPCC-TP. When we set the regulation to a rela-

$\Delta_{reg}$	minJ-ADE(6)	minJ-FDE(6)
$1.0e-3$	4.01	8.59
$1.0e-4$	<b>2.20</b>	<b>4.79</b>
$1.0e-5$	14.42	28.24
0	-	-

Table 5. Different parameter magnitudes in Tikhonov Regularization




Figure 5. Visualization of correlation matrices  $P_{\Delta}^t$  in 2 scenes of nuScenes dataset (blue line: past trajectories, light-green line: ground truth, red line: predictions, green area: driveable area). **Scene 1:** Vehicle ②–④ are about to leave the intersection while ① is just entering it. ②–④ are driving towards the same direction. **Scene 2:** Vehicle ① and ② meet at a T-junction. ② should yields to ① because the latter enters the junction first.

tively large value such as  $1e-3$ , the performance of the enhanced model is even worse than the original model. We additionally conduct an ablation study on the Attention Block. More details can be found in the supplementary.

## 5. Conclusion

In this paper, we propose a new interaction modeling method named IPCC, which is based on one-dimensional motion increments relative to the current positions of the agents. Compared with the traditional methods modeling the future interactions based on the x-y coordinate system, the IPCC method can intuitively reveal the pairwise motion relevance, and furthermore, with fewer parameters need to be predicted. Based on the IPCC method, we design and implement IPCC-TP. This module can be conveniently embedded into state-of-the-art MTP models to model future interactions more precisely by extending the original predictions in the format of marginal probability distributions to a joint prediction represented by a multivariate Gaussian distribution. Experiments demonstrate that IPCC-TP can boost the performance of MTP models on two real-world autonomous driving datasets.

## Acknowledgment

Guangyao Zhai and Yan Di gratefully acknowledge the financial support from China Scholarship Council.



## References

- [1] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–971, 2016. 2
- [2] Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew M Dai, Rafal Jozefowicz, and Samy Bengio. Generating sentences from a continuous space. In *20th SIGNLL Conference on Computational Natural Language Learning, CoNLL 2016*, pages 10–21. Association for Computational Linguistics (ACL), 2016. 6
- [3] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 3, 6
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 2
- [5] Sergio Casas, Cole Gulino, Renjie Liao, and Raquel Urtasun. Spagmn: Spatially-aware graph neural networks for relational behavior forecasting from sensor data. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9491–9497. IEEE, 2020. 2
- [6] Yuning Chai, Benjamin Sapp, Mayank Bansal, and Dragomir Anguelov. Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. In *Conference on Robot Learning*, pages 86–99. PMLR, 2020. 1, 2
- [7] INTERPRET Challenge. <http://challenge.interaction-dataset.com>, 2021. 6
- [8] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, et al. Argoverse: 3d tracking and forecasting with rich maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8748–8757, 2019. 3
- [9] Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *NIPS 2014 Workshop on Deep Learning, December 2014*, 2014. 2
- [10] Henggang Cui, Vladan Radosavljevic, Fang-Chieh Chou, Tsung-Han Lin, Thi Nguyen, Tzu-Kuo Huang, Jeff Schneider, and Nemanja Djuric. Multimodal trajectory predictions for autonomous driving using deep convolutional networks. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2090–2096. IEEE, 2019. 2
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, 2019. 2
- [12] Yan Di, Ruida Zhang, Zhiqiang Lou, Fabian Manhardt, Xiangyang Ji, Nassir Navab, and Federico Tombari. Gpv-pose: Category-level object pose estimation via geometry-guided point-wise voting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6781–6791, 2022. 2
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 2
- [14] Liangji Fang, Qinhong Jiang, Jianping Shi, and Bolei Zhou. Tpnnet: Trajectory proposal network for motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6797–6806, 2020. 1, 2
- [15] Jiyang Gao, Chen Sun, Hang Zhao, Yi Shen, Dragomir Anguelov, Congcong Li, and Cordelia Schmid. Vectornet: Encoding hd maps and agent dynamics from vectorized representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11525–11533, 2020. 2
- [16] Roger Girgis, Florian Golemo, Felipe Codevilla, Martin Weiss, Jim Aldon D’Souza, Samira Ebrahimi Kahou, Felix Heide, and Christopher Pal. Latent variable sequential set transformers for joint multi-agent motion prediction. In *International Conference on Learning Representations*, 2022. 2, 3, 6, 7
- [17] Junru Gu, Chen Sun, and Hang Zhao. Densentnt: End-to-end trajectory prediction from dense goal sets. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15303–15312, 2021. 1, 2
- [18] DE Hilt and DE Seegrist. Ridge: a computer program for calculating ridge regression estimates. *USDA Forest Service Research Note NE (USA)*. no. 236., 1977. 6
- [19] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 2
- [20] Nikolay Jetchev and Marc Toussaint. Trajectory prediction: learning to map situations to robot trajectories. In *Proceedings of the 26th annual international conference on machine learning*, pages 449–456, 2009. 1
- [21] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017. 2
- [22] Xin Kong, Xuemeng Yang, Guangyao Zhai, Xiangrui Zhao, Xianfang Zeng, Mengmeng Wang, Yong Liu, Wanlong Li, and Feng Wen. Semantic graph based place recognition for 3d point clouds. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8216–8223. IEEE, 2020. 1
- [23] Parth Kothari, Sven Kreiss, and Alexandre Alahi. Human trajectory forecasting in crowds: A deep learning perspective. *IEEE Transactions on Intelligent Transportation Systems*, 2021. 3
- [24] Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. Crowds by example. In *Computer graphics forum*, volume 26, pages 655–664. Wiley Online Library, 2007. 3
- [25] Yujia Li, Richard Zemel, Marc Brockschmidt, and Daniel Tarlow. Gated graph sequence neural networks. In *Inter-*

- national Conference on Learning Representations (ICLR)*, 2016. 2
- [26] Ming Liang, Bin Yang, Rui Hu, Yun Chen, Renjie Liao, Song Feng, and Raquel Urtasun. Learning lane graph representations for motion forecasting. In *European Conference on Computer Vision*, pages 541–556. Springer, 2020. 2
- [27] Liang Liu, Guangyao Zhai, Wenlong Ye, and Yong Liu. Unsupervised learning of scene flow estimation fusing with local rigidity. In *IJCAI*, pages 876–882, 2019. 1
- [28] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 2
- [29] Jean Mercat, Thomas Gilles, Nicole El Zoghby, Guillaume Sandou, Dominique Beauvois, and Guillermo Pita Gil. Multi-head attention for multi-modal joint vehicle motion forecasting. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9638–9644. IEEE, 2020. 2
- [30] Jiquan Ngiam, Vijay Vasudevan, Benjamin Caine, Zhengdong Zhang, Hao-Tien Lewis Chiang, Jeffrey Ling, Rebecca Roelofs, Alex Bewley, Chenxi Liu, Ashish Venugopal, et al. Scene transformer: A unified architecture for predicting future trajectories of multiple agents. In *International Conference on Learning Representations*, 2022. 2, 3
- [31] Stefano Pellegrini, Andreas Ess, Konrad Schindler, and Luc Van Gool. You’ll never walk alone: Modeling social behavior for multi-target tracking. In *2009 IEEE 12th international conference on computer vision*, pages 261–268. IEEE, 2009. 3
- [32] Tung Phan-Minh, Elena Corina Grigore, Freddy A Boulton, Oscar Beijbom, and Eric M Wolff. Covnet: Multimodal behavior prediction using trajectory sets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14074–14083, 2020. 1, 2
- [33] Simon Prince. Probability part i: Probability introduction to probability 1.1 random variables. 2, 3
- [34] Christoph Rösmann, Malte Oeljeklaus, Frank Hoffmann, and Torsten Bertram. Online trajectory prediction and planning for social robot navigation. In *2017 IEEE International Conference on Advanced Intelligent Mechatronics (AIM)*, pages 1255–1260. IEEE, 2017. 1
- [35] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *European Conference on Computer Vision*, pages 683–700. Springer, 2020. 2
- [36] Yongzhi Su, Yan Di, Guangyao Zhai, Fabian Manhardt, Jason Rambach, Benjamin Busam, Didier Stricker, and Federico Tombari. Opa-3d: Occlusion-aware pixel-wise aggregation for monocular 3d object detection. *IEEE Robotics and Automation Letters*, 2023. 1
- [37] Qiao Sun, Xin Huang, Junru Gu, Brian C Williams, and Hang Zhao. M2i: From factored marginal trajectory prediction to interactive prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6543–6552, 2022. 2
- [38] Bohan Tang, Yiqi Zhong, Ulrich Neumann, Gang Wang, Siheng Chen, and Ya Zhang. Collaborative uncertainty in multi-agent trajectory forecasting. *Advances in Neural Information Processing Systems*, 34:6328–6340, 2021. 2, 3, 6, 8
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2
- [40] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ramesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, Deva Ramanan, Peter Carr, and James Hays. Argoverse 2: Next generation datasets for self-driving perception and forecasting. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS Datasets and Benchmarks 2021)*, 2021. 6
- [41] Ya Wu, Ariyan Bighashdel, Guang Chen, Gijds Dubbelman, and Pavol Jancura. Continual pedestrian trajectory learning with social generative replay. *IEEE Robotics and Automation Letters*, 8(2):848–855, 2023. 2
- [42] Ye Yuan, Xinshuo Weng, Yanlan Ou, and Kris M Kitani. Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9813–9823, 2021. 2, 3, 6, 7
- [43] Guangyao Zhai, Dianye Huang, Shun-Cheng Wu, HyunJun Jung, Yan Di, Fabian Manhardt, Federico Tombari, Nassir Navab, and Benjamin Busam. Monograspnet: 6-dof grasping with a single rgb image. In *IEEE International Conference on Robotics and Automation*. IEEE, 2023. 2
- [44] Guangyao Zhai, Xin Kong, Jinhao Cui, Yong Liu, and Zhen Yang. Flowmot: 3d multi-object tracking by scene flow association. *arXiv preprint arXiv:2012.07541*, 2020. 1
- [45] Wei Zhan, Liting Sun, Di Wang, Haojie Shi, Aubrey Clause, Maximilian Naumann, Julius Kümmerle, Hendrik Königshof, Christoph Stiller, Arnaud de La Fortelle, and Masayoshi Tomizuka. INTERACTION Dataset: An INTERNATIONAL, Adversarial and Cooperative moTION Dataset in Interactive Driving Scenarios with Semantic Maps. *arXiv:1910.03088 [cs, eess]*, 2019. 6
- [46] Chenyangguang Zhang, Zhiqiang Lou, Yan Di, Federico Tombari, and Xiangyang Ji. Sst: Real-time end-to-end monocular 3d reconstruction via sparse spatial-temporal guidance. *arXiv preprint arXiv:2212.06524*, 2022. 2
- [47] Zhen Zhang, Jiaqing Yan, Xin Kong, Guangyao Zhai, and Yong Liu. Efficient motion planning based on kinodynamic model for quadruped robots following persons in confined spaces. *IEEE/ASME Transactions on Mechatronics*, 26(4):1997–2006, 2021. 1
- [48] Hang Zhao, Jiyang Gao, Tian Lan, Chen Sun, Ben Sapp, Balakrishnan Varadarajan, Yue Shen, Yi Shen, Yuning Chai, Cordelia Schmid, et al. Tnt: Target-driven trajectory prediction. In *Conference on Robot Learning*, pages 895–904. PMLR, 2021. 1, 2