

# Adapting to Length Shift: FlexiLength Network for Trajectory Prediction

Yi Xu Yun Fu

Northeastern University, USA

xu.yi@northeastern.edu, yunfu@ece.neu.edu

## Abstract

Trajectory prediction plays an important role in various applications, including autonomous driving, robotics, and scene understanding. Existing approaches mainly focus on developing compact neural networks to increase prediction precision on public datasets, typically employing a standardized input duration. However, a notable issue arises when these models are evaluated with varying observation lengths, leading to a significant performance drop, a phenomenon we term the *Observation Length Shift*. To address this issue, we introduce a general and effective framework, the *FlexiLength Network (FLN)*, to enhance the robustness of existing trajectory prediction techniques against varying observation periods. Specifically, *FLN* integrates trajectory data with diverse observation lengths, incorporates *FlexiLength Calibration (FLC)* to acquire temporal invariant representations, and employs *FlexiLength Adaptation (FLA)* to further refine these representations for more accurate future trajectory predictions. Comprehensive experiments on multiple datasets, i.e., *ETH/UCY*, *nuScenes*, and *Argoverse 1*, demonstrate the effectiveness and flexibility of our proposed *FLN* framework.

## 1. Introduction

The goal of trajectory prediction is to predict the future locations of agents conditioned on their past observed states, a key step in understanding their motion and behavior patterns. This task is essential yet challenging in many real-world applications such as autonomous driving [21, 23, 59, 77], robotics [10, 36], and scene understanding [11, 69]. Recent advancements in deep learning have significantly improved the accuracy of trajectory predictions [2, 4, 8, 9, 17, 27, 38, 44, 47, 50, 52, 65, 66, 76]. However, achieving such promising prediction performance typically requires complex models and extensive computational resources. Additionally, many state-of-the-art methods are developed on public datasets [7, 55, 61, 74] with a fixed observation length, failing to consider the potential discrepancies between training conditions and diverse real-




Figure 1. The Observation Length Shift phenomenon is a common issue in the trajectory prediction task. The AgentFormer model is trained with a standard observation length of 8 timesteps and tested at varying lengths to compare with Isolated Training (IT). For each dataset, the bar groups from left to right correspond to observation lengths of 2, 4, and 6 timesteps, respectively.

world scenarios. Consequently, the fixed training setups are often too rigid to meet in practice, especially when faced with environmental changes or observational variables.

Recently, some studies have explored the issue of training-testing discrepancy in trajectory prediction. For example, work [68] proposes a solution for scenarios where observations might be incomplete due to occlusions or limited viewpoints, by developing a method that combines trajectory imputation and prediction. Other research efforts [5, 24, 26, 67] focus on tackling the domain shift challenge, where the testing environment differs from the training. Additionally, studies [34, 40, 54] address situations with limited observations, such as having only one or two data points available. These approaches try to mitigate the training-testing gap from various angles. Acknowledging these efforts, the importance of addressing the training-testing gap is highlighted. In our paper, we concentrate on the less-explored issue of observation length discrepancy.

For certain well-known architectures such as RNN [22] and Transformer [57], the ability to process inputs of varying lengths presents a straightforward solution: train the model with the standard input length and then evaluate at different input lengths. This solution is contrasted with Isolated Training (IT), where the model is trained at the specific observation length and evaluated with the same input length. We conducted experiments using a Transformer-

based method AgentFormer [73] on dataset ETH/UCY [31, 46] with results illustrated in Fig. 1. The results reveal a clear performance discrepancy between these two solutions, suggesting performance declines when not evaluated at the training length. We term this phenomenon the Observation Length Shift. It should be noted that while the performance of IT marginally improves, it requires multiple training rounds for different observation lengths. This observation raises two crucial questions:

1. **What are the underlying reasons behind this performance discrepancy?**
2. **Can this issue be resolved without substantial modifications or increased computational expense?**

Prompted by these questions and given the popularity of Transformer-based models in the trajectory prediction task, we conducted an in-depth analysis of a Transformer-based trajectory prediction method, as detailed in Sec. 3. This investigation revealed two main causes, which we summarized as positional encoding deviation and normalization shift. Building on these findings, we introduce a robust and adaptable framework, the FlexiLength Network (FLN), to overcome the Observation Length Shift problem. FLN effectively merges trajectory data from various observation lengths, integrates FlexiLength Calibration (FLC) to create time-invariant representations, and applies FlexiLength Adaptation (FLA) to enhance these representations, thereby improving the trajectory prediction precision. Notably, FLN requires only one-time training, without substantial modifications, but offers the flexibility to different observation lengths, thus addressing our second question.

Despite some studies focusing on positional encoding deviation [13, 42] and normalization shift [71, 72], our work addresses the Observation Length Shift, a commonly overlooked issue in trajectory prediction. We propose the FlexiLength Network (FLN) as a practical and non-trivial solution, especially valuable in real-world scenarios. Unlike Isolated Training (IT), which is time and memory-expensive for adapting to various observation lengths, FLN offers a more resource-efficient alternative. Even though modern architectures accommodate various input lengths, they often suffer performance discrepancies, a gap that FLN effectively bridges while preserving the original model design. Furthermore, FLN can be applied to nearly all existing models that are based on the Transformer architecture. The key contributions of our work are organized as follows:

- We identify and thoroughly analyze the Observation Length Shift phenomenon in trajectory prediction, pinpointing the factors that lead to performance degradation and guiding the direction of our solution.
- We introduce the FlexiLength Framework (FLN), a robust and efficient framework requiring only one-time training. It includes FlexiLength Calibration (FLC) and FlexiLength Adaptation (FLA) for learning

temporal invariant representations, boosting prediction performance with different observation lengths.

- We validate the effectiveness and generality of FLN through comprehensive experiments on multiple datasets. It demonstrates consistent superiority over Isolated Training (IT) with various observation lengths.

## 2. Related Work

**Trajectory Prediction.** The objective is to forecast the future trajectories of agents based on their historical observations. The primary focus has been on understanding the social dynamics among these agents, leading to the development of various methodologies, utilizing Graph Neural Networks (GNNs) as seen in works like [3, 28, 33, 39, 51, 53, 62]. The inherently uncertain and diverse nature of future trajectories has led to the development with generative models, including Generative Adversarial Networks (GANs) [1, 19, 28, 32, 48], Conditional Variational Autoencoders (CVAEs) [25, 37, 49, 63, 64], and Diffusion models [18, 27, 38]. These models have achieved promising results on various datasets. However, most of these methods prioritize prediction accuracy while overlooking the discrepancies that often arise in real-world applications, resulting in decreased adaptability under varying conditions. In contrast, our work focuses on the variation in observation lengths, a challenge that many existing methods face. It is noteworthy that some methods based on RNN [12, 25, 49] and Transformer [16, 56, 70, 73] can handle varying input lengths. However, these approaches show limitations under observation length shifts. Our work delves into this specific challenge, proposing a comprehensive solution that not only enhances the robustness of existing methods but also improves their performance on varying observation lengths.

**Training-Testing Discrepancy.** The gap between training and testing can present in different ways, either in input types or environmental conditions. Recent studies have begun to uncover various aspects of discrepancy in trajectory prediction. Some methods [5, 24, 26, 67] have focused on the environmental shifts between training and testing, offering innovative solutions. Another approach [60] deals with the absence of High-Definition (HD) maps during testing, employing Knowledge Distillation (KD) to transfer map knowledge from a more informed teacher model to a student model. Another aspect recently explored is the completeness of observations. Studies like [34, 40, 54] assume only limited observed trajectory data, specifically one or two frames, is accessible for forecasting. This assumption has led to the concept of instantaneous trajectory prediction. Other studies [30, 68] have addressed the challenges posed by incomplete observations due to occlusions or viewing limitations in real environments. Another recent work [45] attempts to diminish internal discrepancies across datasets




Figure 2. ADE<sub>5</sub> and FDE<sub>5</sub> results for the AgentFormer model, which is trained on the nuScenes dataset using a standard observation length of 4 timesteps, and tested at shorter observation lengths of 2 and 3 timesteps. These results are compared to those obtained through Isolated Training (IT).

collected in arbitrary temporal configurations. While these methods have recognized various aspects of training-testing discrepancies, the issue of observation length shift remains relatively unexplored. Our proposed framework stands out in its generality and effectiveness, improving the ability to handle trajectory data across different observation lengths.

**Test-Time Adaptation.** The objective of test-time adaptation is to adapt models to new data during testing. A common approach includes integrating an auxiliary task [14, 15, 43] that employs straightforward self-supervision techniques on the target domain, which typically requires modifications to the training process in the source domain to include this auxiliary task. However, recent strategies [35, 58] enable straightforward adjustments in the target domain, avoiding the need to modify the training process. Unlike these methods, our proposed framework only requires one-time training, eliminating the necessity for adjustments during training or additional tuning during inference. Moreover, our framework functions independently, without the need for evaluation data during the training phase.

### 3. A Close Look at Performance Degradation

**Observation Length Shift.** The Observation Length Shift phenomenon becomes apparent when models, trained with a specific observation length, are tested using varying lengths, as illustrated in Fig. 1. To investigate further, we train the AgentFormer model [73] using the nuScenes dataset [6] and evaluate its performance across different observation lengths, with the results presented in Fig. 2. These results reveal noticeable performance gaps when the model is evaluated with different observation lengths compared to Isolated Training. Notably, the performance degradation is smaller when the testing observation length is close to the training observation length.

**Positional Encoding Deviation.** Previous studies [20, 41] have demonstrated that Transformer models struggle to generalize to input lengths they haven’t encountered before, due to a lack of inductive bias. Other works [13, 42] identify positional encoding as a critical factor that limits the gener-




Figure 3. Layer Normalization statistics in two different layers of the Transformer encoder within the AgentFormer model, isolatedly trained on the Eth dataset at observation lengths of 2, 6, and 8 timesteps.

alization ability to new lengths and suggest model structure adjustments to enhance the positional encoding. Focusing on the Transformer-based models, we analyze the positional encoding in the AgentFormer. The positional encoding feature  $\tau_n^t(k)$  for agent  $n$  at timestep  $t$  is calculated as follows:

$$\tau_n^t(k) = \begin{cases} \sin((t + H)/10000^{k/d_t}), & k \text{ is even} \\ \cos((t + H)/10000^{(k-1)d_t}), & k \text{ is odd} \end{cases}, \quad (1)$$

where  $k$  represents the feature index,  $d_t$  is the feature dimension, and  $H$  is the observation length. It is obvious that simply removing observations without compensation will cause the length change, subsequently altering the positional encoding. Consequently, we hypothesize that the positional encoding deviation is one crucial factor to the performance gap. Although other Transformer-based models [75, 76] use a fully learnable positional embedding, our experiments in Sec. 5 indicate that this type of positional embedding remains a key reason.

**Normalization Shift.** Previous works [71, 72] have studied the issue of normalization shift in the context of image and video classification. They found that altering network widths impacts Batch Normalization (BN), leading to variations in the means and variances of aggregated features and thereby affecting feature aggregation consistency. Given that Layer Normalization (LN) is widely used in Transformer-based models and operates similarly to BN, we hypothesize that the discrepancy in LN statistics could be another factor affecting the generalization ability to different lengths. For example, if a model is trained using an observation length of  $H$  and later evaluated using a length of  $L$ , the input to the LN would be the intermediate feature  $f^L$  during testing, resulting in the following output:

$$y^L = \gamma^H \frac{f^L - \mu^H}{\sqrt{\sigma^{H^2} + \epsilon}} + \beta^H, \quad (2)$$

where  $\mu^H, \sigma^{H^2}$  are derived from the data distribution with input length  $H$ , and  $\gamma^H, \beta^H$  are parameters learned during

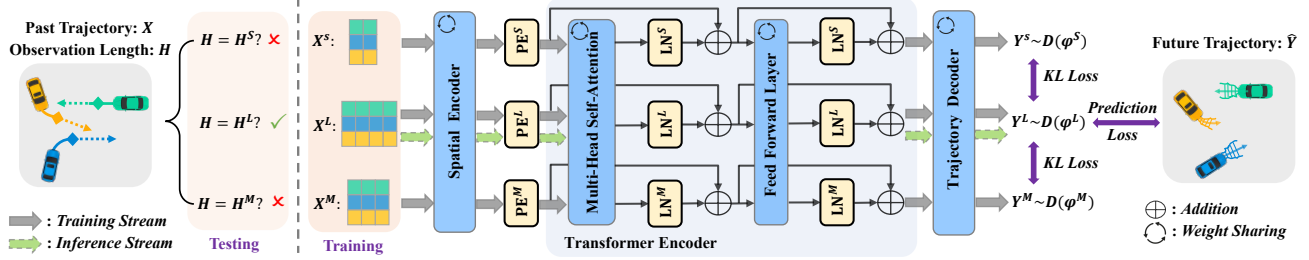


Figure 4. Illustration of our FlexiLength Network (FLN). The map encoding branch is omitted for simplicity. During training, with inputs of varying observation lengths  $H^S$ ,  $H^M$ , and  $H^L$ , we utilize FlexiLength Calibration (FLC) to acquire temporal invariant representations. Furthermore, FlexiLength Adaptation (FLA) is employed to align these invariant representations with different sub-networks, thereby augmenting the model capabilities. During inference, the sub-network with the closest match in observation length is activated.

training with the same input length. We train the AgentFormer model using observation lengths  $H$  and  $L$  and then analyze the feature statistics in the Transformer encoder’s two Layer Normalizations (LNs). Clear statistical differences at varying observation lengths can be observed in Fig. 3. Importantly, both  $\mu$  and  $\sigma$  are data-dependent, indicating that the differences arise directly from the data itself. Therefore, we believe that the variations in LN statistics are another reason for the Observation Length Shift phenomenon in Transformer-based models.

## 4. FlexiLength Network

### 4.1. Problem Formulation

Despite some variations in terminology, the fundamental problem setup for multi-agent trajectory prediction remains the same across existing methods. Essentially, the objective is to predict the future trajectories of multiple agents,  $\mathbf{Y} \in \mathbb{R}^{N \times T \times F}$ , based on their past observed states  $\mathbf{X} \in \mathbb{R}^{N \times H \times F}$ . Here,  $N$  represents the number of agents,  $H$  is the length of observation,  $T$  is the prediction duration and  $F$  typically stands for 2, denoting coordinate dimensions. Common approaches aim to develop a generative model, denoted as  $p_\theta(\mathbf{Y}|\mathbf{X}, \mathbf{I})$ , where  $\mathbf{I}$  represents the contextual information such as High-definition (HD) maps.

As previously mentioned, these methods often face challenges when there is a variation in observation lengths during evaluation. Hence, our objective is to develop a model  $p_\theta(\mathbf{Y}|\mathbf{X}, \mathbf{I})$  that can be evaluated across a range of input lengths  $\mathbf{H}$ , and achieves similar or even better performance compared to Isolated Training (IT), where  $\mathbf{H}$  denotes a set of different observation lengths.

### 4.2. FlexiLength Calibration

As observed in Sec. 3, the performance gap is smaller when the evaluation length is closer to the training length. This observation inspired us to incorporate observations of varying lengths during training and to develop corresponding

sub-networks for feature extraction from these sequences. Given an observation trajectory sequence  $\mathbf{X}$ , we can generate sequences of varying lengths either through truncation or a sliding window approach. In this way, we can collect three sequences:  $X^S$ ,  $X^M$ , and  $X^L$ , corresponding to Short, Medium, and Long lengths, with corresponding lengths  $H^S$ ,  $H^M$ , and  $H^L$ . During training, these three sequences are each fed into their corresponding sub-networks  $F^S(\cdot)$ ,  $F^M(\cdot)$ , and  $F^L(\cdot)$ , for processing.

Fig. 4 showcases our FlexiLength Network (FLN) developed atop the Transformer-based architecture, where the map branch is omitted for simplicity. We categorize the components of typical Transformer-based models into several key parts: a Spatial Encoder for spatial feature extraction, a Positional Encoder (PE) for embedding positional information, a Transformer Encoder for temporal dependency modeling, and a Trajectory Decoder for generating trajectories. During training, the three inputs  $X^S$ ,  $X^M$ , and  $X^L$  are each processed by their respective sub-networks  $F^S(\cdot)$ ,  $F^M(\cdot)$ , and  $F^L(\cdot)$ . During inference, the sub-network matching the observation length is exclusively activated. This approach sets up a computational stream in the FLN framework, enabling effective evaluation across various observation lengths.

**Sub-Network Weight Sharing.** Given the three trajectory samples  $X^S$ ,  $X^M$ , and  $X^L$ , we employ shared weights for the spatial encoder, temporal encoder, and trajectory decoder across all sub-networks, to find a group of parameters  $\theta$  to extract the spatio-temporal features as follows:

$$Y^* \sim \mathcal{D}(\psi^*), \quad \mathcal{D}(\psi^*) = F^*(X^*; \theta), \quad (3)$$

where  $*$   $\in$   $\{S, M, L\}$ , and  $\mathcal{D}(\psi^*)$  denotes the distribution parameterized by  $\psi^*$ , that  $Y^*$  adheres to, such as a bivariate Gaussian or Laplace mixture distribution [76].

This design contrasts with Isolated Training (IT) by being more efficient in terms of parameter usage, as it maintains only a single set of weights applicable to various lengths. Additionally, this shared-weight strategy enhances



performance by implicitly providing the model with the prior knowledge that these three observed sequences are part of the same trajectory. This insight makes the model more resilient to observation length shifts.

**Temporal Distillation.** Trajectory prediction models generally perform better with longer observations  $X^L$  as more movement information is contained. Consequently, we regard  $Y^L \sim \mathcal{D}(\psi^L)$  as the most ‘accurate’ among the three predictions. To update the parameters of  $F^L(\cdot)$ , we employ a negative log-likelihood loss, represented as follows:

$$\mathcal{L}_{reg} = -\log(\mathbb{P}(\hat{Y}|\psi^L)), \quad (4)$$

where  $\hat{Y}$  is the ground truth trajectory. Note that as multiple agents always exist within a single trajectory, the above equation is a simplified version and omits the calculation for the average value per agent at each time step.

For updating the parameters of the other two sub-networks,  $F^S(\cdot)$  and  $F^M(\cdot)$ , a direct approach might involve calculating a corresponding negative log-likelihood loss as shown in Eq. (4). However, it has potential drawbacks. Since the weights are shared across all three sub-networks, the optimal parameters for  $F^L(\cdot)$  may not be as effective for  $F^S(\cdot)$  or  $F^M(\cdot)$ . Additionally, optimizing the log-likelihood loss for  $\psi^S$  and  $\psi^M$  could also hinder the performance of  $F^L(\cdot)$ . This is because their inputs  $X^S$  and  $X^M$ , with less movement information, might lead to a poorer fit. To address this issue, we employ the KL divergence loss [29] to calibrate and incorporate  $\psi^S$  and  $\psi^M$  into the computational process as follows:

$$\mathcal{L}_{kl} = \text{KL}(\mathcal{D}(\psi^L) || \mathcal{D}(\psi^M)) + \text{KL}(\mathcal{D}(\psi^L) || \mathcal{D}(\psi^S)). \quad (5)$$

Updating the weights of shared sub-networks through Eq. (5) aims to align the predicted distributions of the ‘student’ networks ( $\mathcal{D}(\psi^M)$  and  $\mathcal{D}(\psi^S)$ ) closely with the predicted distribution of the ‘teacher’ network ( $\mathcal{D}(\psi^L)$ ). This process facilitates the transfer of valuable knowledge from  $F^L(\cdot)$  to both  $F^M(\cdot)$  and  $F^S(\cdot)$ . The entire set of parameters for our FLN is updated as follows:

$$\mathcal{L}_{FLN} = \mathcal{L}_{reg} + \lambda_1 \cdot \mathcal{L}_{kl}, \quad (6)$$

where  $\lambda$  serves to balance these two terms during backpropagation, we have set  $\lambda = 1$  in our implementation.

Considering both Sub-Network Weight Sharing and Temporal Distillation,  $\mathcal{L}_{reg}$  offers guidance for the prediction task, while  $\mathcal{L}_{kl}$  offers intra-trajectory knowledge to network training. This will promote  $\mathcal{D}(\psi^L)$ ,  $\mathcal{D}(\psi^M)$ , and  $\mathcal{D}(\psi^S)$  exhibit high similarity, as the length shift does not alter the distribution of one specific trajectory. Through this approach, FLN not only learns temporal invariant representations but also ensures ease of implementation, as it does not require any changes to the feature extractors.

### 4.3. FlexiLength Adaptation

We developed FLC to guide our FLN in learning temporal invariant representations. Additionally, we introduce FlexiLength Adaptation (FLA) to optimize the fit of these invariant features across different sub-networks, thereby further improving their representational capacity.

**Independent Positional Encoding.** Our analysis in Sec. 3 identifies positional encoding as one factor that can confuse the model regarding observation length. To counter this, we implement independent positional encoding for each sub-network. Taking the positional encoding from the AgentFormer [73] as an example, we define the positional encoding for inputs with diverse observation lengths as follows:

$$\tau_n^t(k) = \begin{cases} \sin((t + H^*)/10000^{k/d_t}), & k \text{ is even} \\ \cos((t + H^*)/10000^{(k-1)d_t}), & k \text{ is odd} \end{cases}, \quad (7)$$

where  $* \in L, M, S$ . While recent Transformer-based models [75, 76] have shifted to a fully learnable positional embedding instead of the traditional sinusoidal pattern, we adopt this learnable positional embedding across each of our sub-networks. As a result, for each input sequence  $X^S$ ,  $X^M$ , and  $X^L$ , we define unique learnable positional encoding  $\tau^S$ ,  $\tau^M$ , and  $\tau^L$ . Note that the implementation of independent positional encoding for each sub-network is a resource-efficient solution, introducing only a negligible increase in parameters and computation.

**Specialized Layer Normalization.** As outlined in Sec. 3, normalization shift is another reason for the performance gap when observation length shifts. Denote the intermediate features for  $X^S$ ,  $X^M$ , and  $X^L$  as  $f^S$ ,  $f^M$ , and  $f^L$ , respectively. We introduce the specialized layer normalization for each input sequence as follows:

$$y^* = \gamma^* \frac{f^* - \mu^*}{\sqrt{\sigma^{*2} + \epsilon}} + \beta^*, \quad (8)$$

where  $* \in \{S, M, L\}$ . This specialized normalization allows for the independent learning of  $\gamma$  and  $\beta$ , and the calculation of  $\mu$  and  $\sigma$  specific to each sequence during training. Moreover, this approach is efficient, as normalization typically involves a simple transformation with less than 1% of total model parameters.

## 5. Experiments

### 5.1. Settings

**Baselines.** Our proposed FLN is a general framework that is compatible with those Transformer-based methods. We have chosen two widely recognized open-source methods, AgentFormer [73] and HiVT [75], to integrate with our FLN. Furthermore, for a more comprehensive comparison,

Method	Note	#Param	ADE <sub>5</sub> /FDE <sub>5</sub> ↓ K = 5 Samples			ADE <sub>10</sub> /FDE <sub>10</sub> ↓ K = 10 Samples		
			2 Timesteps	3 Timesteps	4 Timesteps	2 Timesteps	3 Timesteps	4 Timesteps
AFormer [73]	-	6.67M	2.26/4.53	2.00/4.10	1.86/3.89	1.72/3.06	1.53/2.95	1.45/2.86
♦AFormer-Mixed	$\rho^S=\rho^M=\rho^L=0.5$	6.67M	2.14/4.38	1.98/4.05	1.94/3.91	1.68/3.06	1.50/2.91	1.57/2.93
♦AFormer-Mixed	$\rho^S=0.75, \rho^M=\rho^L=0.5$	6.67M	2.10/4.38	1.98/4.02	2.05/4.14	1.65/3.01	1.50/2.90	1.66/3.11
♦AFormer-Tuning	4Ts→2Ts	6.67M	2.10/4.32	2.15/4.31	2.28/4.43	1.60/2.98	1.81/3.34	1.79/3.12
♦AFormer-Tuning	4Ts→3Ts	6.67M	2.32/4.67	1.93/3.94	2.32/4.96	1.98/3.33	1.48/2.92	2.12/3.97
♦AFormer-Joint	-	6.67×3M	<u>1.98/4.19</u>	1.95/3.97	1.90/4.02	1.54/3.04	1.50/2.91	1.49/2.96
♦AFormer-IT	-	6.67×3M	2.02/4.23	1.93/3.97	<u>1.86/3.89</u>	<u>1.52/3.00</u>	1.47/2.89	<u>1.45/2.86</u>
<b>AFormer-FLN</b>	$\rho^S=\rho^M=\rho^L=1$	6.68M	<b>1.92/3.91</b>	<b>1.88/3.89</b>	<b>1.83/3.78</b>	<b>1.47/2.90</b>	<b>1.43/2.82</b>	<b>1.32/2.73</b>

Table 1. Comparison with baseline models on nuScenes, evaluated using ADE<sub>5</sub>/FDE<sub>5</sub> and ADE<sub>10</sub>/FDE<sub>10</sub> metrics. The best results are highlighted in bold and the second-best results are underlined.

Method	Note	#Param	ADE <sub>6</sub> /FDE <sub>6</sub> ↓ K = 6 Samples		
			10 Timesteps	20 Timesteps	30 Timesteps
HiVT-64 [75]	-	662K	1.23/1.48	0.91/1.37	0.69/1.04
♦HiVT-64-Mixed	$\rho^S=\rho^M=\rho^L=0.5$	662K	1.14/1.57	0.86/1.22	0.80/1.12
♦HiVT-64-Mixed	$\rho^S=0.75, \rho^M=\rho^L=0.5$	662K	1.08/1.49	0.88/1.22	0.84/1.13
♦HiVT-64-Tuning	30Ts→10Ts	662K	<u>0.90/1.38</u>	1.14/1.48	1.18/1.31
♦HiVT-64-Tuning	30Ts→20Ts	662K	1.33/1.61	0.87/1.20	1.12/1.29
♦HiVT-64-Joint	-	662×3K	0.98/1.45	0.89/1.24	0.74/1.09
♦HiVT-64-IT	-	662×3K	0.92/1.43	<u>0.81/1.17</u>	<u>0.69/1.04</u>
<b>HiVT-64-FLN</b>	$\rho^S=\rho^M=\rho^L=1$	680K	<b>0.81/1.25</b>	<b>0.72/1.08</b>	<b>0.65/0.98</b>

Table 2. Comparison with baseline models on the Argoverse 1 validation set, evaluated using ADE<sub>6</sub>/FDE<sub>6</sub> metrics. The best results are highlighted in bold and the second-best results are underlined.

we include four additional baseline models: (1) **Mixed Sampling**: In each training iteration, we assign three probabilities  $\rho^S$ ,  $\rho^M$ , and  $\rho^L$  to take trajectory data with observation lengths  $H^S$ ,  $H^M$ , and  $H^L$  for training. (2) **Fine-tuning**: The model is trained using the observation length  $H^L$ , and then fine-tuned using another length ( $H^S$  or  $H^M$ ) until convergence. (3) **Joint**: We expand the training dataset by including trajectory samples from all three observation lengths ( $H^S$ ,  $H^M$ , and  $H^L$ ) and then train the original model without any structural changes. (4) **Isolated Training**: The model is trained exclusively with trajectory data of a single observation length,  $H^S$ ,  $H^M$ , or  $H^L$ .

**Datasets.** We use three following datasets: (1) The ETH/UCY dataset [31, 46] is a primary benchmark for pedestrian trajectory prediction, including five datasets, Eth, Hotel, Univ, Zara1, and Zara2, with densely interactive trajectories sampled at 2.5Hz. (2) The nuScenes dataset [74] is a large autonomous driving dataset featuring 1000 scenes, each annotated at 2Hz, and includes HD maps with 11 semantic classes. (3) Argoverse 1 [7] contains 323,557 real-world driving sequences sampled at 10Hz, complemented with HD maps for trajectory prediction.

**Evaluation Protocol.** Across the three datasets, we use the metrics Average Displacement Error (ADE<sub>K</sub>) and Final Displacement Error (FDE<sub>K</sub>) for comparison, where  $K$  indicates the number of trajectories to be predicted. Each dataset follows its own evaluation protocol: (1) For

ETH/UCY, the leave-one-out setting is standard, with the task to predict 12 future time steps from 8 observed steps. Here,  $K$  is commonly set to 20. (2) Within the nuScenes dataset, as used in AgentFormer, only vehicle data is considered, predicting 12 future steps from 4 observed steps.  $K$  is usually set to 5 and 10. (3) In Argoverse 1, sequences are segmented into 5-second intervals, with predictions made for 30 future steps (3 seconds) based on 20 observed steps (2 seconds) involving multiple agents. The validation set is used for our evaluation purpose, with  $K$  being 6.

**Implementation Details.** For each dataset, we define three different observation lengths (in timestep)  $\mathbf{H} = \{H^S, H^M, H^L\}$  for training, where  $H^L$  is set as the default length in the standard evaluation protocol, considering the potential unavailability of data beyond this length. Specifically, we use  $\mathbf{H} = \{2, 6, 8\}$  for ETH/UCY,  $\mathbf{H} = \{2, 3, 4\}$  for nuScenes, and  $\mathbf{H} = \{10, 20, 30\}$  for Argoverse 1. The prediction length  $T$  remains as defined in the standard settings. The prototype models are trained at observation length  $H^L$  and evaluated at  $H^S$  and  $H^M$ . Isolated Training (IT) refers to training and evaluating the model separately at observation length  $H^S$ ,  $H^M$ , and  $H^L$ . More details and experiments are provided in the Supplementary Material.

## 5.2. Main Results

**Comparison with Baselines.** Tab. 1 presents the results using the AgentFormer model on the nuScenes dataset, where




Figure 5. Performance on five ETH/UCY datasets using the AgentFormer model, measured with  $ADE_{20}$ . These results are compared with those of the baseline model and Isolated Training (IT), showcasing notable improvements achieved by our FLN.

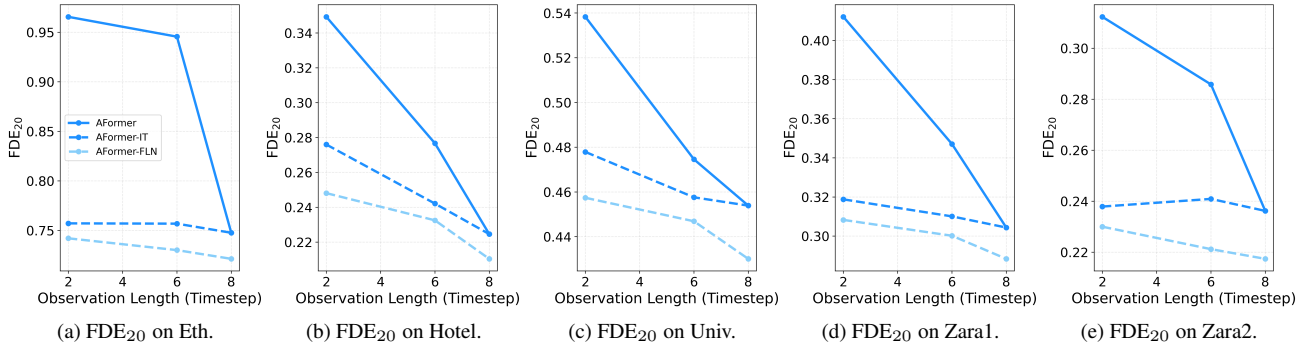


Figure 6. Performance on five ETH/UCY datasets using the AgentFormer model, measured with  $FDE_{20}$ . These results are compared with those of the baseline model and Isolated Training (IT), showcasing notable improvements achieved by our FLN.

our FLN framework outperforms all baseline models, showcasing its robustness in adapting to different observation lengths. It is noticeable that Mixed Sampling helps mitigate the Observation Length Shift issue, as indicated by the improvement at observation lengths of 2 and 3 timesteps, compared to the results achieved with the standard training protocol. However, this improvement comes at the cost of reduced accuracy at the 4-timestep observation length. Increasing the probability for length  $H^S$  improves the performance at that length, but conversely, the performance at  $H^L$  becomes worse. This pattern suggests a trade-off: improved results at shorter lengths result in diminished performance at longer lengths. One interesting observation is that the performance at  $H^M$  remains relatively stable, possibly because the quantity of sequence data for  $H^M$  and  $H^L$  remains unchanged, thus maintaining a consistent relationship between these two data clusters. Fine-tuning enables the baseline method to achieve comparable performance with IT, neglecting information from other lengths. As for the joint baseline, its performance is similar to IT, and even better than IT at the 2-timestep length when  $K = 5$ . However, a significant issue with both Joint and IT baselines is the increased complexity of the models due to repeated training sessions. In contrast, our FLN requires only a sin-

gle training session with minimal additional computational overhead, yet it allows for evaluation at various lengths and obtains superior results.

**Performance across Datasets.** Additionally, the evaluation on the ETH/UCY benchmark, as illustrated in Fig. 5 and Fig. 6, reveals that our proposed FLN outperforms IT at different observation lengths. We also extend our FLN to the HiVT method, which utilizes learnable embeddings for positional encoding, and evaluate it on the Argoverse 1 validation set. The results, presented in Tab. 2, show that FLN consistently achieves superior performance across three different observation lengths, affirming its robustness in adapting to various observation lengths.

### 5.3. Generality Study

We have demonstrated that FLN surpasses IT at the training observation lengths, while not evaluating the model at lengths outside the training range. Inspired by our findings in Sec. 3, where the performance gap narrows with observation length closer to the standard input length, we introduce an efficient inference approach enabling FLN to be evaluated at any length. When encountering an unseen length  $H^I$ , we first calculate its difference from  $H^S$ ,  $H^M$ , and  $H^L$ , and then activate the sub-network with the least difference in length. In cases where two sub-networks have the same

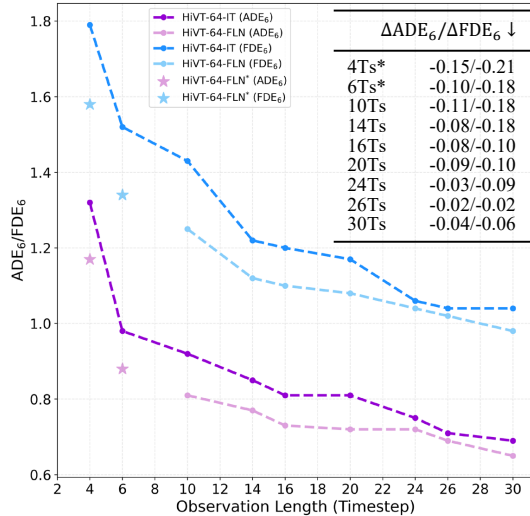


Figure 7. Performance of FLN and IT at different observation lengths on the Argoverse 1 validation set. The improvements are listed in the table. Results from observation lengths outside the range of 10 to 30 are denoted with a \*.

length difference, we default to the one with the longer input length. This approach allows for the evaluation of FLN across a spectrum of unseen observation lengths.

Since we set the longest length  $H^L$  as the standard observation length, it is challenging to generate trajectory sequences with observation lengths longer than  $H^L$  due to their pre-segmentation. The results on the Argoverse 1 validation set with observation length shorter than 30 timesteps are illustrated in Fig. 7. These results clearly demonstrate that FLN consistently outperforms IT at all observation lengths, even those beyond the 10-30 range. This validates the strong generality of our FLN in handling observation lengths that are not included during training.

#### 5.4. Ablation Study

**Length Combination.** We implement different combinations of observation lengths on Argoverse 1, with the results presented in Tab. 3. It is observed that FLN-2 surpasses IT in terms of  $ADE_6$  and  $FDE_6$  at the length of 10 and 16 timesteps, and achieves comparable results at the length of 30. However, its performance at 20-timestep length is worse than IT, likely due to the absence of this sequence length in the training data. Both FLN-3 and FLN-4 outperform IT at all observation lengths. This improvement can be attributed to the existence of trajectory data with intermediate observation lengths, which helps mitigate the Observation Length Shift within this range. Although FLN-4 shows better results than FLN-3 at lengths of 10, 16, and 20 timesteps, it also costs additional time and resources for training.

**Model Design.** We conducted a detailed analysis of each component within our FLN on the nuScenes dataset, with results presented in Tab. 4. We first remove Weight Sharing

Method	Combination	$ADE_6/FDE_6 \downarrow$ K = 6 Samples			
		10Ts	16Ts	20Ts	30Ts
HiVT-64-IT [75]	-	0.92/1.43	0.81/1.20	0.81/1.17	0.69/1.04
HiVT-64-FLN-2	10/30	0.88/1.37	0.80/1.20	0.84/1.21	0.69/1.02
HiVT-64-FLN-3	10/20/30	0.81/1.25	0.73/1.10	<b>0.72/1.08</b>	<b>0.65/0.98</b>
HiVT-64-FLN-4	10/16/20/30	<b>0.79/1.19</b>	<b>0.70/1.02</b>	<b>0.72/1.06</b>	0.68/0.99

Table 3. Performance of FLN on the Argoverse 1 validation set using different combinations of observation lengths. The best results are highlighted in bold.

Method	Note	$ADE_5/FDE_5 \downarrow$ K = 5 Samples		
		2Ts	3Ts	4Ts
AFormer-IT [73]	-	2.02/4.23	1.93/3.97	1.86/3.89
AFormer-FLN	w/o WS	1.94/3.95	1.89/3.92	1.92/3.94
AFormer-FLN	w/o TD	2.23/4.39	1.98/4.01	1.85/3.81
AFormer-FLN	w/o IPE	2.13/4.33	2.02/4.07	1.95/3.98
AFormer-FLN	w/o SLN	2.12/4.30	1.98/4.02	1.93/3.95
AFormer-FLN	-	1.92/3.91	1.88/3.89	1.83/3.78

Table 4. Ablation study of FLN on nuScenes. WS, TD, IPE, and SLN denote Sub-Network Weight Sharing, Temporal Distillation, Independent Positional Encoding, and Specialized Layer Normalization, respectively. The best results are highlighted in bold.

(w/o WS), creating unique sub-networks for each length. The results are better than IT at observation lengths of 2 and 3 timesteps, but worse at length of 4. This suggests that shared weights enable FLN to more effectively capture temporal-invariant features across diverse input lengths. We excluded Temporal Distillation (w/o TD) from FLN, meaning not utilizing the KL loss for optimization. This leads to the performance worse than IT at observation lengths of 2 and 3 timesteps, suggesting the necessity of our TD design. We also tested shared Positional Encoding (w/o IPE) and shared Layer Normalization (w/o SLN). Both resulted in a performance decrease, validating the significance of our Independent Positional Encoding (IPE) and Specialized Layer Normalization (SLN) designs.

## 6. Conclusion

In this paper, we tackle the critical challenge of Observation Length Shift in trajectory prediction by introducing the FlexiLength Network (FLN). This novel framework, incorporating FlexiLength Calibration (FLC) and FlexiLength Adaptation (FLA), offers a general solution for handling varying observation lengths and requires only one-time training. Our thorough experiments on datasets such as ETH/UCY, nuScenes, and Argoverse 1 demonstrate that FLN not only improves prediction accuracy and robustness over a range of observation lengths but also consistently outperforms Isolated Training (IT).

**Limitation and Future Work.** One limitation is that FLN will increase training time due to handling several input sequences per training iteration. Going forward, we will focus on enhancing the training efficiency of FLN.



## References

- [1] Javad Amirian, Jean-Bernard Hayet, and Julien Pettré. Social ways: Learning multi-modal distributions of pedestrian trajectories with GANs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 2964–2972, 2019. 2
- [2] Görkay Aydemir, Adil Kaan Akan, and Fatma Güney. Adapt: Efficient multi-agent trajectory prediction with adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8295–8305, 2023. 1
- [3] Inhwan Bae, Jin-Hwi Park, and Hae-Gon Jeon. Learning pedestrian group representations for multi-modal trajectory prediction. In *Proceedings of the European Conference on Computer Vision*, pages 270–289, 2022. 2
- [4] Inhwan Bae, Jean Oh, and Hae-Gon Jeon. Eigentrajectory: Low-rank descriptors for multi-modal trajectory forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10017–10029, 2023. 1
- [5] Mohammadhossein Bahari, Saeed Saadatnejad, Ahmad Rahimi, Mohammad Shaverdikondori, Amir Hossein Shahidzadeh, Seyed-Mohsen Moosavi-Dezfooli, and Alexandre Alahi. Vehicle trajectory prediction works, but not everywhere. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17123–17133, 2022. 1, 2
- [6] 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 multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11621–11631, 2020. 3
- [7] 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. 1, 6
- [8] Guangyi Chen, Zhenhao Chen, Shunxing Fan, and Kun Zhang. Unsupervised sampling promoting for stochastic human trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17874–17884, 2023. 1
- [9] Hao Chen, Jiase Wang, Kun Shao, Furui Liu, Jianye Hao, Chenyong Guan, Guangyong Chen, and Pheng-Ann Heng. Traj-mae: Masked autoencoders for trajectory prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8351–8362, 2023. 1
- [10] Yujiao Cheng, Liting Sun, Changliu Liu, and Masayoshi Tomizuka. Towards efficient human-robot collaboration with robust plan recognition and trajectory prediction. *IEEE Robotics and Automation Letters*, 5(2):2602–2609, 2020. 1
- [11] Hyung-gun Chi, Kwonjoon Lee, Nakul Agarwal, Yi Xu, Karthik Ramani, and Chiho Choi. Adamsformer for spatial action localization in the future. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17885–17895, 2023. 1
- [12] Luigi Filippo Chiara, Pasquale Coscia, Sourav Das, Simone Calderara, Rita Cucchiara, and Lamberto Ballan. Goal-driven self-attentive recurrent networks for trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2518–2527, 2022. 2
- [13] Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Lukasz Kaiser. Universal transformers. In *Proceedings of the International Conference on Learning Representations*, 2019. 2, 3
- [14] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *Proceedings of the Advances in Neural Information Processing Systems*, 2014. 3
- [15] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *Proceedings of the International Conference on Learning Representations*, 2018. 3
- [16] Francesco Giuliari, Irtiza Hasan, Marco Cristani, and Fabio Galasso. Transformer networks for trajectory forecasting. In *Proceedings of the IEEE International Conference on Pattern Recognition*, pages 10335–10342, 2020. 2
- [17] Junru Gu, Chenxu Hu, Tianyuan Zhang, Xuanyao Chen, Yilun Wang, Yue Wang, and Hang Zhao. Vip3d: End-to-end visual trajectory prediction via 3d agent queries. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5496–5506, 2023. 1
- [18] Tianpei Gu, Guangyi Chen, Junlong Li, Chunze Lin, Yongming Rao, Jie Zhou, and Jiwen Lu. Stochastic trajectory prediction via motion indeterminacy diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17113–17122, 2022. 2
- [19] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social GAN: Socially acceptable trajectories with generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2255–2264, 2018. 2
- [20] Michael Hahn. Theoretical limitations of self-attention in neural sequence models. *Transactions of the Association for Computational Linguistics*, 8:156–171, 2020. 3
- [21] Christopher Hazard, Akshay Bhagat, Balarama Raju Budharaju, Zhongtao Liu, Yunming Shao, Lu Lu, Sammy Omari, and Henggang Cui. Importance is in your attention: agent importance prediction for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2532–2535, 2022. 1
- [22] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. 1
- [23] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhai Wang, et al. Planning-oriented autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17853–17862, 2023. 1
- [24] Manh Huynh and Gita Alaghband. Online adaptive temporal memory with certainty estimation for human trajectory prediction. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, pages 940–949, 2023. 1, 2

- [25] Boris Ivanovic and Marco Pavone. The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2375–2384, 2019. 2
- [26] Boris Ivanovic, James Harrison, and Marco Pavone. Expanding the deployment envelope of behavior prediction via adaptive meta-learning. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2023. 1, 2
- [27] Chiyu Jiang, Andre Cornman, Cheolho Park, Benjamin Sapp, Yin Zhou, Dragomir Anguelov, et al. Motiondiffuser: Controllable multi-agent motion prediction using diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9644–9653, 2023. 1, 2
- [28] Vineet Kosaraju, Amir Sadeghian, Roberto Martín-Martín, Ian Reid, Hamid Rezafofighi, and Silvio Savarese. Socialbigat: Multimodal trajectory forecasting using bicycle-gan and graph attention networks. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 137–146, 2019. 2
- [29] Solomon Kullback. *Information theory and statistics*. Courier Corporation, 1997. 5
- [30] Bernard Lange, Jiachen Li, and Mykel J Kochenderfer. Scene informer: Anchor-based occlusion inference and trajectory prediction in partially observable environments. *arXiv preprint arXiv:2309.13893*, 2023. 2
- [31] Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. Crowds by example. In *Computer graphics forum*, pages 655–664. Wiley Online Library, 2007. 2, 6
- [32] Jiachen Li, Hengbo Ma, and Masayoshi Tomizuka. Conditional generative neural system for probabilistic trajectory prediction. In *Proceedings of the International Conference on Intelligent Robots and Systems*, 2019. 2
- [33] Lihuan Li, Maurice Pagnucco, and Yang Song. Graph-based spatial transformer with memory replay for multi-future pedestrian trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2231–2241, 2022. 2
- [34] Rongqing Li, Changsheng Li, Dongchun Ren, Guangyi Chen, Ye Yuan, and Guoren Wang. Bcdiff: Bidirectional consistent diffusion for instantaneous trajectory prediction. In *Proceedings of the Advances in Neural Information Processing Systems*, 2023. 1, 2
- [35] Yuejiang Liu, Parth Kothari, Bastien Van Delft, Baptiste Bellot-Gurlet, Taylor Mordan, and Alexandre Alahi. Ttt++: When does self-supervised test-time training fail or thrive? In *Proceedings of the Advances in Neural Information Processing Systems*, pages 21808–21820, 2021. 3
- [36] Mohammad Mahdavian, Payam Nikdel, Mahdi TaherAhmadi, and Mo Chen. Stpotr: Simultaneous human trajectory and pose prediction using a non-autoregressive transformer for robot follow-ahead. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 9959–9965, 2023. 1
- [37] Karttikeya Mangalam, Harshayu Girase, Shreyas Agarwal, Kuan Hui Lee, Ehsan Adeli, Jitendra Malik, and Adrien Gaidon. It is not the journey but the destination: End-point conditioned trajectory prediction. In *Proceedings of the European Conference on Computer Vision*, pages 759–776, 2020. 2
- [38] Weibo Mao, Chenxin Xu, Qi Zhu, Siheng Chen, and Yanfeng Wang. Leapfrog diffusion model for stochastic trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5517–5526, 2023. 1, 2
- [39] Abdullah Mohamed, Kun Qian, Mohamed Elhoseiny, and Christian Claudel. Social-STGCNN: A social spatio-temporal graph convolutional neural network for human trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14424–14432, 2020. 2
- [40] Alessio Monti, Angelo Porrello, Simone Calderara, Pasquale Coscia, Lamberto Ballan, and Rita Cucchiara. How many observations are enough? Knowledge distillation for trajectory forecasting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6553–6562, 2022. 1, 2
- [41] Hyoungwook Nam, Segwang Kim, and Kyomin Jung. Number sequence prediction problems for evaluating computational powers of neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4626–4633, 2019. 3
- [42] Hyoungwook Nam, Seung Byum Seo, Vikram Sharma Malthody, Noor Michael, and Lan Li. I-bert: Inductive generalization of transformer to arbitrary context lengths. *arXiv preprint arXiv:2006.10220*, 2020. 2, 3
- [43] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *Proceedings of the European Conference on Computer Vision*, pages 69–84. Springer, 2016. 3
- [44] Daehee Park, Hobin Ryu, Yunseo Yang, Jegyeong Cho, Jiwon Kim, and Kuk-Jin Yoon. Leveraging future relationship reasoning for vehicle trajectory prediction. In *Proceedings of the International Conference on Learning Representations*, 2023. 1
- [45] Daehee Park, Jaewoo Jeong, and Kuk-Jin Yoon. Improving transferability for cross-domain trajectory prediction via neural stochastic differential equation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024. 2
- [46] Stefano Pellegrini, Andreas Ess, Konrad Schindler, and Luc J. Van Gool. You’ll never walk alone: Modeling social behavior for multi-target tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 261–268, 2009. 2, 6
- [47] Luke Rowe, Martin Ethier, Eli-Henry Dykhne, and Krzysztof Czarnecki. Fjmp: Factorized joint multi-agent motion prediction over learned directed acyclic interaction graphs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13745–13755, 2023. 1
- [48] Amir Sadeghian, Vineet Kosaraju, Ali Sadeghian, Noriaki Hirose, Hamid Rezafofighi, and Silvio Savarese. SoPhie: An attentive GAN for predicting paths compliant to social and

- physical constraints. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1349–1358, 2019. 2
- [49] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *Proceedings of the European Conference on Computer Vision*, pages 683–700, 2020. 2
- [50] Ari Seff, Brian Cera, Dian Chen, Mason Ng, Aurick Zhou, Nigamaa Nayakanti, Khaled S Refaat, Rami Al-Rfou, and Benjamin Sapp. Motionlm: Multi-agent motion forecasting as language modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8579–8590, 2023. 1
- [51] Liushuai Shi, Le Wang, Chengjiang Long, Sanping Zhou, Mo Zhou, Zhenxing Niu, and Gang Hua. SgcN: Sparse graph convolution network for pedestrian trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8994–9003, 2021. 2
- [52] Liushuai Shi, Le Wang, Sanping Zhou, and Gang Hua. Trajectory unified transformer for pedestrian trajectory prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9675–9684, 2023. 1
- [53] Jianhua Sun, Qinhong Jiang, and Cewu Lu. Recursive social behavior graph for trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 660–669, 2020. 2
- [54] Jianhua Sun, Yuxuan Li, Liang Chai, Hao-Shu Fang, Yong-Lu Li, and Cewu Lu. Human trajectory prediction with momentary observation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6467–6476, 2022. 1, 2
- [55] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2446–2454, 2020. 1
- [56] Li-Wu Tsao, Yan-Kai Wang, Hao-Siang Lin, Hong-Han Shuai, Lai-Kuan Wong, and Wen-Huang Cheng. Social-ssl: Self-supervised cross-sequence representation learning based on transformers for multi-agent trajectory prediction. In *Proceedings of the European Conference on Computer Vision*, pages 234–250, 2022. 2
- [57] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 5998–6008, 2017. 1
- [58] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *Proceedings of the International Conference on Learning Representations*, 2021. 3
- [59] Jingke Wang, Tengju Ye, Ziqing Gu, and Junbo Chen. Ltp: Lane-based trajectory prediction for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17134–17142, 2022. 1
- [60] Yuning Wang, Pu Zhang, Lei Bai, and Jianru Xue. Enhancing mapless trajectory prediction through knowledge distillation. *arXiv preprint arXiv:2306.14177*, 2023. 2
- [61] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, et al. Argoverse 2: Next generation datasets for self-driving perception and forecasting. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, 2021. 1
- [62] Chenxin Xu, Maosen Li, Zhenyang Ni, Ya Zhang, and Siheng Chen. Groupnet: Multiscale hypergraph neural networks for trajectory prediction with relational reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6498–6507, 2022. 2
- [63] Chenxin Xu, Weibo Mao, Wenjun Zhang, and Siheng Chen. Remember intentions: Retrospective-memory-based trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6488–6497, 2022. 2
- [64] Pei Xu, Jean-Bernard Hayet, and Ioannis Karamouzas. Socialvae: Human trajectory prediction using timewise latents. In *Proceedings of the European Conference on Computer Vision*, 2022. 2
- [65] Yi Xu, Jing Yang, and Shaoyi Du. Cf-lstm: Cascaded feature-based long short-term networks for predicting pedestrian trajectory. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 12541–12548, 2020. 1
- [66] Yi Xu, Dongchun Ren, Mingxia Li, Yuehai Chen, Mingyu Fan, and Huaxia Xia. Tra2tra: Trajectory-to-trajectory prediction with a global social spatial-temporal attentive neural network. *IEEE Robotics and Automation Letters*, 6(2):1574–1581, 2021. 1
- [67] Yi Xu, Lichen Wang, Yizhou Wang, and Yun Fu. Adaptive trajectory prediction via transferable gnn. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6520–6531, 2022. 1, 2
- [68] Yi Xu, Armin Bazarjani, Hyung-gun Chi, Chiho Choi, and Yun Fu. Uncovering the missing pattern: Unified framework towards trajectory imputation and prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9632–9643, 2023. 1, 2
- [69] Yiteng Xu, Peishan Cong, Yichen Yao, Runnan Chen, Yue-nan Hou, Xinge Zhu, Xuming He, Jingyi Yu, and Yuexin Ma. Human-centric scene understanding for 3d large-scale scenarios. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20349–20359, 2023. 1
- [70] Cunjun Yu, Xiao Ma, Jiawei Ren, Haiyu Zhao, and Shuai Yi. Spatio-temporal graph transformer networks for pedestrian trajectory prediction. In *Proceedings of the European Conference on Computer Vision*, pages 507–523, 2020. 2
- [71] Jiahui Yu and Thomas S Huang. Universally slimmable networks and improved training techniques. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1803–1811, 2019. 2, 3
- [72] Jiahui Yu, Linjie Yang, Ning Xu, Jianchao Yang, and Thomas Huang. Slimmable neural networks. In *Proceed-*

*ings of the International Conference on Learning Representations*, 2019. [2](#), [3](#)

- [73] Ye Yuan, Xinshuo Weng, Yanglan 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](#), [5](#), [6](#), [8](#)
- [74] Wei Zhan, Liting Sun, Di Wang, Haojie Shi, Aubrey Clause, Maximilian Naumann, Julius Kummerle, Hendrik Konigshof, Christoph Stiller, Arnaud de La Fortelle, et al. Interaction dataset: An international, adversarial and cooperative motion dataset in interactive driving scenarios with semantic maps. *arXiv preprint arXiv:1910.03088*, 2019. [1](#), [6](#)
- [75] Zikang Zhou, Luyao Ye, Jianping Wang, Kui Wu, and Kejie Lu. Hivt: Hierarchical vector transformer for multi-agent motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8823–8833, 2022. [3](#), [5](#), [6](#), [8](#)
- [76] Zikang Zhou, Jianping Wang, Yung-Hui Li, and Yu-Kai Huang. Query-centric trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17863–17873, 2023. [1](#), [3](#), [4](#), [5](#)
- [77] Yanliang Zhu, Dongchun Ren, Yi Xu, Deheng Qian, Mingyu Fan, Xin Li, and Huaxia Xia. Simultaneous past and current social interaction-aware trajectory prediction for multiple intelligent agents in dynamic scenes. *ACM Transactions on Intelligent Systems and Technology*, 13(1):1–16, 2021. [1](#)