

## MIT Open Access Articles

*Learning Socio-Temporal Graphs for  
Multi-Agent Trajectory Prediction*

The MIT Faculty has made this article openly available. *Please share* how this access benefits you. Your story matters.

**Citation:** Li, Yuke, Chen, Lixiong, Chen, Guangyi, Chan, Ching-Yao, Zhang, Kun et al. 2024. "Learning Socio-Temporal Graphs for Multi-Agent Trajectory Prediction."

**As Published:** <https://doi.org/10.1145/3688865.3689481>

**Publisher:** ACM|Proceedings of the 5th International Workshop on Human-centric Multimedia Analysis

**Persistent URL:** <https://hdl.handle.net/1721.1/157626>

**Version:** Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

**Terms of Use:** Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



# Learning Socio-Temporal Graphs for Multi-Agent Trajectory Prediction

Yuke Li \*  
Boston College  
Boston, U.S.A.

Lixiong Chen †  
University of Oxford  
Oxford, U.K.

Guangyi Chen  
Carnegie Mellon University;  
Mohamed bin Zayed University of  
Artificial Intelligence  
Pittsburgh, U.S.A.

Ching-Yao Chan  
University of California Berkeley  
Berkeley, U.S.A.

Kun Zhang  
Carnegie Mellon University;  
Mohamed bin Zayed University of  
Artificial Intelligence  
Pittsburgh, U.S.A.

Stefano Anzellotti  
Boston College  
Boston, U.S.A.

Donglai Wei  
Boston College  
Boston, U.S.A.

## ABSTRACT

In order to predict a pedestrian’s trajectory in a crowd accurately, one has to take into account her/his underlying socio-temporal interactions with other pedestrians consistently. Unlike existing work that represents the relevant information separately, partially, or implicitly, we propose a complete representation for it to be fully and explicitly captured and analyzed. In particular, we introduce a Directed Acyclic Graph-based structure, which we term **Socio-Temporal Graph** (STG), to explicitly capture pair-wise socio-temporal interactions among a group of people across both space and time. Our model is built on a time-varying generative process, whose latent variables determine the structure of the STGs. We design an attention-based model named **STGformer** that affords an end-to-end pipeline to learn the structure of the STGs for trajectory prediction. Our solution achieves overall state-of-the-art prediction accuracy in two large-scale benchmark datasets. Our analysis shows that a person’s past trajectory is critical for predicting another person’s future path. Our model learns this relationship with a strong notion of socio-temporal localities. Statistics show that utilizing this information explicitly for prediction yields a noticeable performance gain with respect to the trajectory-only approaches.

\*Equal contribution

†Equal contribution

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
*MM’24, October 28 - November 1, 2024, Melbourne, Australia.*

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 979-8-4007-1192-3/24/10  
<https://doi.org/10.1145/3688865.3689481>

## CCS CONCEPTS

• **Computing methodologies** → **Computer vision representations.**

## KEYWORDS

Computer Vision, Human-Centric Trajectory Analysis

### ACM Reference Format:


Yuke Li , Lixiong Chen , Guangyi Chen, Ching-Yao Chan, Kun Zhang, Stefano Anzellotti, and Donglai Wei. 2024. Learning Socio-Temporal Graphs for Multi-Agent Trajectory Prediction. In *Proceedings of the 5th International Workshop on Human-centric Multimedia Analysis (HCMA ’24), October 28-November 1 2024, Melbourne, VIC, Australia* *Proceedings of the 32nd ACM International Conference on Multimedia (MM’24), October 28-November 1, 2024, Melbourne, Australia.* ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3688865.3689481>

## 1 INTRODUCTION

When a person walks in a crowd, in addition to the location of their destination, this person’s motion path is also affected by other agents’ movements, e.g. following the lead, taking a detour to avoid collisions, or making a stop at a rendezvous. Accurately forecasting an agent’s motion trajectories depends on how accurately their surrounding agents’ activities are captured and modeled. In vision-based applications, including autonomous driving

In a scene of only pedestrians, there exists a hidden graph that governs the socio-temporal interactions among the agents. Relative information is either discarded, partially represented, or implicitly encoded by various models proposed to date, such as pooling

To this end, we introduce a structure named **Socio-Temporal Graphs** (STGs, see Fig. 1). An STG instance not only considers each agent’s past trajectories or how they interact at each time instant, but also explicitly describes how one’s past trajectory affects the prediction for another one’s future trajectory. Also, in contrast to existing work that lets such information be hard-coded into the model



**Figure 1: (a) Given the observed trajectories of multiple agents at time  $t=0-2$ , (b) actions at  $t-1$ , the socio-temporal models for multi-agent trajectory forecasting at  $t$ . (a) Most previous models assume a fixed structure, while (b) our model learns the socio-temporal graph (STG).**

We design an attention-based model, named **STGformer**, that realizes STG and allows it to be learned through an end-to-end pipeline. In particular, when training the model by maximizing the data likelihood, we introduce two separate modules, one representing the prior distribution of the latent sequence and the other representing the posterior of the distribution conditioned on observations. By minimizing the corresponding KL divergence between these two distributions, the latent prior can be learned recursively, so the shape of the distribution is effectively regularized without explicit parametric restrictions.

Our solution achieves state-of-the-art results on Stanford Drone and ETH/UCY trajectory prediction benchmarks with respect to the trajectory-only approaches. Interestingly, our analysis shows that a person’s past trajectory is critical for predicting another person’s future path, and explicitly utilizing this fact can yield a noticeable performance gain. Furthermore, our experiment indicates that this information exhibits strong spatial-temporal localities.

**Contributions.** (1) We introduce the STG structure and formulate multi-agent trajectory forecasting based on learning it with a latent sequential generative model. (2) We design an attention-based model named STGformer that allows STG to be effectively learned in an end-to-end fashion. (3) We provide analysis to identify the efficacy of learning STG regarding performance gain.

## 2 RELATED WORK

**Multi-agent trajectory prediction.** Current trajectory forecasting approaches exploit a variety of deep neural network models [7]. For instance, some methods adopt a set of LSTMs to characterize the movement of pedestrians [1, 9]. By contrast, NSP-SFM [51] formulates the trajectory forecasting with neural differential equations [32]. PECNet [24] infers the endpoints to assist in trajectory forecasting using a Conditional Variational AutoEncoder (CVAE) [7]. LB-EBM [26] generates predictions with a latent energy-based model [5], while MID [8] uses diffusion processes [34].

Other approaches rely on graphs. Among these, some harness a graph neural network to associate all agents in a fully-connected fashion [11, 12, 33], while others [36, 40, 44, 48, 50] replace the graph neural network with attention mechanisms [37]. CausalHTP [4] and CausalMR [20] analyze the causal effects with the aid

of a pre-defined causal graph [27] that models the effects from the confounding environmental elements to the agents.

Most of these studies share the idea of implicitly assuming a fully-connected underlying structure to consider the interactions among the agents per time step. However, the connectivity among agents might be falsely addressed. For example, one agent might be too far from another agent, so considering the motion of one does not help predict the motion of the other. Learning these socio-temporal dependencies and pruning irrelevant edges might reduce overfitting and improve predictions.

**Modeling socio-temporal interactions.** Recent studies have attempted to capture socio-temporal correlations in the context of pedestrian trajectory forecasting. Some studies model socio-temporal interactions implicitly. Among them, a subset first assign a network to each person to obtain temporal cues, and subsequently leverage another network for social cues [1, 4, 9, 12, 13, 21, 24, 25, 29, 42, 46, 48, 51]. by contrast, some recent approaches [41, 50] tie the social and temporal correlations jointly due to their co-occurrence.

New studies introduced learned time-invariant structures to tackle different tasks, such as physical motion modeling, or to infer the bioinformatic signaling [19, 22, 45]. DA-Former [10] leverages a greedy strategy to search for an optimal graph for the machine translation problem. Following these works, a set of recent techniques explicitly learn a static graph [16, 43]. Nevertheless, more than the static representation might be required to characterize the time-varying socio-temporal correlations. EvolveGraph [17] partially addresses this issue with a learned dynamic local social structure. Our work takes inspiration from the aforementioned recent advances, introducing a novel approach to learning time-varying STGs for human trajectory forecasting.


**Transformer-based trajectory prediction.** With the success of Transformer models in both fields of natural language processing (NLP) and computer vision, many methods [6, 8, 18, 47, 49, 50] employ the Transformer to build trajectory prediction systems. Despite the strong representation ability of the attention model to learn the relations, the Transformer-based framework implicitly builds the fully-connected graph of all agents. This paper shows that we can explicitly learn STG, which reduces ambiguity and significantly improves trajectory prediction.

## 3 METHOD

Let  $x_i^t$  denote the position of agent  $i$  at time  $t$ , and  $\mathbf{x}^t = \{x_i^t\}_i$  the positions of all agents at time  $t$ . The goal is to predict  $\mathbf{x}^t$  from the collective past observations  $\mathbf{x}^{0:t-1}$ .

### 3.1 Problem Formulation

We define the socio-temporal graph (STG) as a binary adjacency matrix to represent how current observations are related to the past in the form of the directed acyclic graph (DAG). For example, at time  $t=3$ , directed edges with value 1 in binary adjacency matrix  $\mathcal{G}^3$  (Fig. 2a) indicate a possible one-sided interaction between the agent represented by the source node and the agent represented by the destination node. Such interactions can span over time time  $\tau = \{0 : t-1\}$  to  $t$  to indicate the fact that one’s past behavior can influence another’s decision in the future. In other words, our



**Figure 2: Socio-temporal graph (STG).** (a) STG definition as the adjacency matrix between agents at the current step and all previous steps.  $G^3$  means at  $t=3$  all agents establish a possible connection with agents over  $t \leq 3$ . (b) Autoregressive model for trajectory prediction, where  $G^t$  is the embedding from an evolving sequence that constructs expanding  $G^3$ . (c)  $G^3$  has two components,  $G_u^3$  and  $G_v^3$ , which produce  $G^3$  through bilinear product.

hypothesis is that each observation  $\mathbf{x}^t$  is generated not only by  $\mathbf{x}^{0:t-1}$ , but also the underlying socio-temporal graph,  $G^t$ .

To learn the dynamics of  $G^t$  over time from past observations/training data as opposed to a static one

$$p(\mathbf{x}^t, G^t | \mathbf{x}^{0:t-1}, G^{0:t-1}) = p_\Phi(\mathbf{x}^t | \mathbf{x}^{0:t-1}, G^t) p_\Psi(G^t | G^{0:t-1}) \quad (1)$$

Specifically, inspired by the bilinear model for graph generation **STG dynamic module**. We model the STG prior model with a multivariate Gaussian model

$$p_\Psi(G^t | G^{0:t-1}) = \mathcal{N}(G^t; \mu_\Psi(G^{0:t-1}), I), \quad (2)$$

which is implemented by one Transformer decoder block **STG-aware trajectory module**. We first introduce the STG-aware attention block and then incorporate it into the trajectory module.

$G^t$  is obtained using Eqn. ?? to guide the evaluation of self-attention

Similarly, to construct  $p_\Phi(\mathbf{x}^t | \mathbf{x}^{0:t-1}, G^t)$ , we have

$$p_\Phi(\mathbf{x}^t | \mathbf{x}^{0:t-1}, G^t) = \mathcal{N}(\mathbf{x}^t; \mu_\Phi(\mathbf{x}^{0:t-1}, G^t), I). \quad (3)$$

with mean  $\mu^t$  and an identity matrix of a Multivariate Gaussian:

$$\mu_\Phi = \text{MLP}_x(\text{STG-aware-attn}^t) \quad (4)$$

We use another Transformer decoder, as well as an MLP to instantiate  $\mu_\Phi$ . The prediction of  $\mathbf{x}^t$  is then sampled according to the output of Eqn. 4, using our proposed STG-aware attention (STG-aware-attn).

## 3.2 Learning

**Variational Inference.** The goal of training is to determine  $\Psi$  and  $\Phi$  that maximize the log likelihood of the observed data  $p(x)$ . As obtaining the exact solution is intractable, we instead seek to evaluate its evidence lower bound (ELBO) with a posterior model for STG  $q_\Theta(\tilde{G}^t | \tilde{G}^{0:t-1}, \mathbf{x}^{0:t})$ . We denote the STG samples from the

posterior model as  $\tilde{G}^t$ , where those from the prior model as  $G^t$ .

$$\begin{aligned} \log p(\mathbf{x}^{0:T}) &\geq L(\Theta, \Phi, \Psi) = \mathbb{E}_{G^{0:T} \sim q_\Theta} \log \frac{p(\mathbf{x}^{0:T}, G^{0:T})}{q_\Theta(\tilde{G}^{0:T} | \mathbf{x}^{0:T})} \\ &= \mathbb{E}_{G^{0:T} \sim q_\Theta} \log \prod_{t=0}^T \left( \frac{p_\Phi(\mathbf{x}^t | \mathbf{x}^{0:t-1}, G^t) p_\Psi(G^t | G^{0:t-1})}{q_\Theta(\tilde{G}^t | \tilde{G}^{0:t-1}, \mathbf{x}^{0:t})} \right) \\ &= \sum_{t=0}^T \mathbb{E}_{G^{0:T} \sim q_\Theta} \log \left( \frac{p_\Phi(\mathbf{x}^t | \mathbf{x}^{0:t-1}, G^t) p_\Psi(G^t | G^{0:t-1})}{q_\Theta(\tilde{G}^t | \tilde{G}^{0:t-1}, \mathbf{x}^{0:t})} \right) \\ &= \sum_{t=0}^T \left[ \mathbb{E}_{G^t \sim q_\Theta} \log p_\Phi(\mathbf{x}^t | \mathbf{x}^{0:t-1}, G^t) \right. \\ &\quad \left. - D_{KL} \left( q_\Theta(\tilde{G}^t | \tilde{G}^{0:t-1}, \mathbf{x}^{0:t}) \| p_\Psi(G^t | G^{0:t-1}) \right) \right] \quad (5) \end{aligned}$$

where  $(\cdot)^{0:t-1}$  is an empty set when  $t=0$ .

**Training objective.** First, due to the normality of  $p_\Phi$ , the first term in Eqn. 3 reduces to the mean square error loss

$$L_{MSE} = \sum_{t=0}^T \mathbb{E}_{G^t \sim q_\Theta} (\|\mathbf{x}^t - \mu_\Phi(\mathbf{x}^{0:t-1}, G^t)\|^2), \quad (6)$$

where the expectation is evaluated numerically.

Then, we assume the posterior also takes a form of conditional Multivariate Gaussian:

$$q_\Theta(\tilde{G}^t | \tilde{G}^{0:t-1}, \mathbf{x}^{0:t}) = \mathcal{N}(\mu_\Theta(\mathbf{x}^{0:t}, \tilde{G}^{0:t-1}), I), \quad (7)$$

where  $\mu(\cdot; \Theta)$  is produced by a separate Transformer decoder

To further regularize the distribution of the resulting structure  $G^t$ , we also apply a  $L_0$ -norm penalty. In conclusion, we have the following training objective function

$$L = L_{MSE} + L_{KL} + \zeta \|G^{0:T}\|_0. \quad (8)$$

## 3.3 Inference

During the initialization, given the known trajectory  $\mathbf{x}^{0:t_0-1}$ , we can use the posterior model to generate  $\tilde{G}^{0:t_0-1}$ . We then feed  $\tilde{G}^{0:t_0-1}$  to the prior module to estimate  $G^{t_0}$  from Eq. 2 and ???. This design exploits the historical data to learn the reasonable  $G^t$  for predictions. To generate plausible future trajectories predictions, we sample from  $p_\Phi(\mathbf{x}^t | \mathbf{x}^{0:t-1}, G^t)$ , following Eqn. 3 and 4.

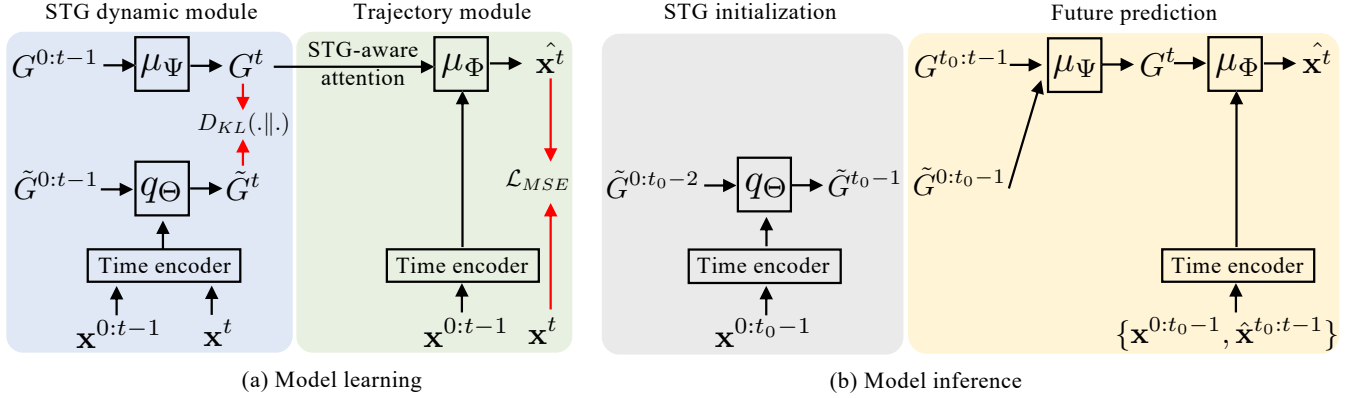
Again, we expect the distribution of our prediction maximizes the ELBO by minimizing the KL divergence in Eqn. 5 as the underlying hypothesis is that predictions over new observations and old samples used for training are drawn from a similar distribution being captured by our learned  $G^t$ .

## 4 EXPERIMENTS

We conducted extensive experiments to justify the effectiveness of our approach in terms of trajectory prediction. We also conduct experiments on the information captured by the STG by examining its existence, efficacy, and socio-temporal characteristics.

### 4.1 Experimental Setup

**Datasets.** The performance of our approach to trajectory forecasting was evaluated on two widely used datasets: ETH/UCY [15, 28] and the Stanford Drone Dataset (SDD) [31]. The ETH/UCY dataset contains five scenes, most of which include more than 700 different pedestrians. The human trajectories are captured in real-world scenes. The SDD dataset comprises long video sequences for 20



**Figure 3: Learning and inference for STGFormer model. (a) During training, we minimize the STG prediction from the prior and posterior modules and the reconstruction from the trajectory module. (b) During inference, we first initialize STGs with the posterior module and then predict future trajectories with the prior and trajectory modules.**

**Table 1: Results on the SDD dataset with best-of-20 ADE/FDE scores. For fair comparison, We rank methods with trajectory input only (T) 1st in bold and 2nd underlined. Also we showcase the results from those with extra image input (T+I).**

	Input	ADE/FDE ↓
Y-Net	T + I	7.85/11.85
V <sup>2</sup> -Net	T + I	7.12/11.39
NSP-SFM	T + I	6.52/10.61
Muse-VAE	T + I	6.36/11.10
Social-GAN	T	27.23/41.44
EvolveGraph	T	13.90/22.90
GroupNet	T	9.31/16.11
GP-Graph	T	9.10/13.80
LB-EBM	T	8.87/15.61
PCCSNET	T	8.62/16.16
NPSN	T	8.56/ <u>11.85</u>
MemoNet	T	8.56/12.66
AgentFormer	T	8.35/13.03
Expert	T	7.69/14.38
MID	T	<u>7.61</u> /14.30
<b>STGFormer (Ours)</b>	T	<b>7.35/11.39</b>

scenes captured using a drone in a top-down view around a university campus. It labels complete trajectories of different categorized moving objects (e.g., pedestrians, bicyclists, and vehicles) from entering the scene until they exit. Both datasets include highly dynamic scenarios with rich socio-temporal correlations, and as such, they are ideal for evaluating the performance of our approach.

Following state-of-the-art methods, [4, 9, 11, 12, 50], the trajectories of all datasets are sampled at 0.4 seconds intervals. The first 3.2 seconds (8 frames) are observed for each video, and the next 4.8 seconds (12 frames) are to be predicted. We adopt a leave-one-scene-out approach for our experiments on the ETH/UCY dataset, training and validating our model on videos in 4 scenes and testing on the 5th scene. We repeat this procedure for all five scenes. In addition, we apply the same training and testing procedure for all baseline methods to ensure a fair comparison.

**Evaluation metrics.** We follow the same evaluation metrics adopted by previous work [4, 8, 12, 48, 50], including the Best-of-20 average

displacement error (ADE) and final displacement error (FDE). In addition, our evaluation is performed for each agent compared to the respective ground truth, where ADE computes the mean square error (MSE) of 20 predictions and the ground truth, while FDE calculates the L2 distance between the final locations of 20 predictions and the ground truth.

## 4.2 Comparing with State-of-the-art Methods

**Methods in comparison.** We categorize methods as Trajectory-Only (T) and Trajectory-and-Image (T+I) methods, as the additional image information may be crucial in certain circumstances yet increases computation cost. Methods using T+I include Y-Net[23], V<sup>2</sup>-Net [41], NSP-SFM [51], Muse-VAE [14], and Social-BiGAT [12]. Methods using T-only include GP-Graph [2], MID [8], MemoNet [44], NPSN [3], GroupNet [43], Expert [52], PCCSNET [35], LB-EBM [26], Trajectron++ [33], Y-Net[23], AgentFormer [50], EvolveGraph [17], Social-GAN [9], Expert [52], CausalHTP [4], STAR [48], Trajectron++ [33], SG-net[39], and Social-GAN [9].

**Quantitative results on SDD dataset.** Tab. 1 shows that STGFormer outperforms other approaches with trajectories-only input consistently in ADE and FDE. Notably, our method improves the ADE by achieving a value of 7.35. In terms of FDE, our approach accomplishes a score of 11.39. These findings tip the balance towards our approach with respect to the other approaches.

**Quantitative results on ETH/UCY dataset.** In the second part of our benchmark experiments, our results on ETH/UCY datasets are summarized in Tab. 2. Experiment results demonstrate that our proposed method achieves the best average 0.18 ADE and 0.35 FDE over the trajectory-only approaches.


**Qualitative results.** Fig. 4 and Fig. 6 showcase the most-likely prediction from STGFormer and Expert [52] on SDD and ETH/UCY datasets. STGFormer generates more accurate predictions concerning the true future than Expert [52], especially for longer-term predictions.

## 4.3 Analysis of Learned STGs

What differentiates STGformer from other models is that past trajectories of copresent agents in the same scene are explicitly used by our model to predict any agent’s future trajectory. We validate the

**Table 2: The benchmark Best-of-20 ADE/FDE results (↓) on the ETH/UCY datasets. For fair comparison, We rank methods with trajectory input only (T) 1st in bold and 2nd underlined. Also we showcase the results from those with extra image input (T+I). † Up-to-date results from the official implementations are worse than the original ones due to an evaluation bug.**

	Input	ETH	Hotel	Univ.	Zara1	Zara2	AVG
Social-BiGAT	T + I	0.69/1.29	0.49/1.01	0.30/0.62	0.36/0.75	0.55/1.32	0.48/1.00
Y-Net	T + I	0.28/0.33	0.10/0.14	0.24/0.41	0.17/0.27	0.13/0.22	0.18/0.27
V <sup>2</sup> -Net	T + I	0.23/0.37	0.11/0.16	0.21/0.35	0.19/0.30	0.14/0.24	0.18/0.28
NSP-SFM	T + I	0.25/0.24	0.09/0.13	0.21/0.38	0.16/0.27	0.12/0.20	0.17/0.24
Social-GAN	T	0.81/1.52	0.72/1.61	0.60/1.26	0.34/0.69	0.42/0.84	0.58/1.18
CausalHTP	T	0.60/0.98	0.30/0.54	0.32/0.64	0.28/0.58	0.52/1.10	0.40/0.77
Trajectron++ †	T	0.67/1.18	0.18/0.28	0.30/0.54	0.25/0.41	0.18/0.32	0.32/0.55
SG-Net †	T	0.47/0.77	0.20/0.38	0.33/0.62	0.18/0.32	0.15/0.28	0.27/0.47
STAR	T	0.36/0.65	0.17/0.36	0.31/0.62	0.26/0.55	0.22/0.46	0.26/0.53
GroupNet †	T	0.46/0.73	0.15/0.25	0.26/0.49	0.21/0.39	0.17/0.33	0.25/0.44
GP-Graph	T	0.43/0.63	0.18/0.30	0.24/0.42	0.17/0.31	0.15/0.29	0.23/0.39
AgentFormer †	T	0.45/0.75	0.14/0.22	0.25/0.45	0.18/0.30	0.14/0.24	0.23/0.39
LB-EBM	T	0.30/0.52	0.13/0.20	0.27/0.52	0.20/0.37	0.15/0.29	0.21/0.38
PCCSNET	T	0.28/0.54	0.11/0.19	0.29/0.60	0.21/0.44	0.15/0.34	0.21/0.42
MID	T	0.39/0.66	0.13/0.22	0.22/0.45	0.17/0.30	0.13/0.27	0.21/0.38
NPSN	T	0.36/0.59	0.16/0.25	0.23/0.39	0.18/0.32	0.14/0.25	0.21/0.36
MemoNet	T	0.40/0.61	0.11/0.17	0.24/0.43	0.18/0.32	0.14/0.24	0.21/0.35
Expert	T	0.37/0.65	0.11/0.15	0.20/0.44	0.15/0.31	0.12/0.26	0.19/0.36
<b>STGFormer (Ours)</b>	T	<b>0.27/0.56</b>	<b>0.11/0.17</b>	0.22/0.45	0.16/0.31	0.14/0.24	<b>0.18/0.35</b>



**Figure 4: Qualitative results on SDD. Each example overlays the predictions from MID [8] (yellow dashed line) and ours STGFormer (cyan dashed line), historical observations (blue line), and ground truth future (red line)**

usefulness of this information by examining its distribution in space and time, and its socio-temporal dynamics. On SDD dataset, we compute the statistics of the edge weights of each learned instance of  $G^t$  and examine the relevant properties.

**Spatial analysis.** Fig. 5a depicts the likelihood of the existence of edges of a value 1 over space. The plot shows one agent is more likely to be correlated with her/his nearby agents; the further away two agents are located, the less likely they are to be correlated.

**Temporal analysis.** The samples are re-grouped in terms of the duration that each edge spans. The distribution presented in Fig. 5b indicates two facts:(1) there are social edges that are learned to be

non-trivial, indicating the socio-temporal notion pair-wise relationship between two agents can be captured through learning; (2) this notion becomes more apparent as the relationship is established in a shorter time-span. The plot shows that “Short-term” edges are more likely to be considered than their “long-term” counterparts for trajectory predictions.

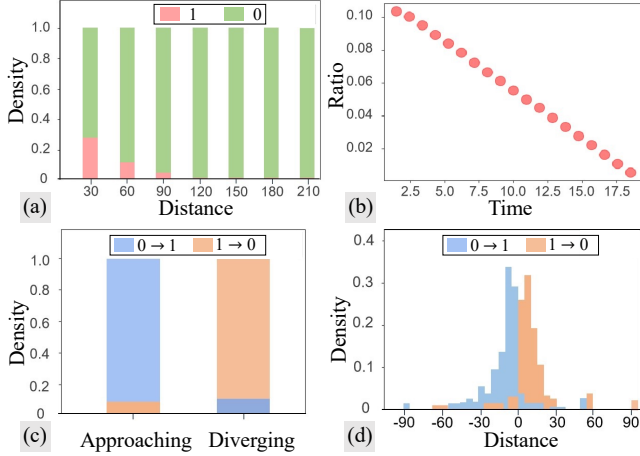
**Socio-temporal dynamics.** We also study how the pair-wise correlations change over time and space. To this end, we define two events: agent  $i$  and agent  $j$  are said to be “approaching” each other when the corresponding distance is decreasing:  $\|\hat{x}_i^t - \hat{x}_j^t\|_2 < \|\hat{x}_i^{t-1} - \hat{x}_j^{t-1}\|_2$ , or otherwise they are “diverging” from each other ( $\|\hat{x}_i^t - \hat{x}_j^t\|_2 > \|\hat{x}_i^{t-1} - \hat{x}_j^{t-1}\|_2$ ). The scenario is exemplified in Fig. 7

Fig. 5c shows that statistically, most edge values flip from 0 to 1 when two agents “approaching” each other, and vice-versa, most values flip from 1 to 0 when two agents are walking away from each other. Moreover, the symmetry and heavy tail distribution in Fig. 5d suggest that there exists a single underlying localized perimeter surrounding each agent, where when another agent enters or leaves it, the corresponding edge flips its value accordingly.

**Analysis of Learned STGs on ETH/UCY datasets.** We also conduct the same analysis of the learned STGs on each scene of the ETH/UCY datasets (Fig. 8A). The learned STGs on the ETH/UCY datasets have similar characteristics to that on the SDD dataset. Specifically, Fig. 8B reveals the likelihood of the existence of edges of a value 1 over space. The distribution presented in Fig. 8C indicates two facts:(1) the socio-temporal notion pair-wise relationship between two agents can be captured through learning STGs; (2) “Short-term” edges are more likely to be considered than their “long-term” counterparts for trajectory predictions. Fig. 8D illustrates

**Table 3: Ablation studies on the ETH/UCY and SDD datasets. We report the ADE / FDE scores over 20 predictions.**

	ETH	Hotel	Univ.	Zara1	Zara2	AVG	SDD
w/o social	0.79/1.12	0.61/0.74	0.53/0.85	0.69/0.92	0.62/1.07	0.65/0.94	14.71/21.07
w/o LP	0.62/1.07	0.49/0.55	0.45/0.69	0.42/0.61	0.47/0.66	0.49/0.72	12.93/17.42
Stationary- $G^t$	0.55/0.84	0.44/0.49	0.44/0.65	0.38/0.60	0.34/0.52	0.43/0.62	11.40/17.15
Short-term- $G^t$	0.45/0.69	0.31/0.42	0.35/0.63	0.30/0.52	0.31/0.38	0.34/0.45	9.04/14.29
w/o $G^t$	0.44/0.65	0.26/0.38	0.31/0.59	0.23/0.47	0.22/0.35	0.29/0.49	8.83/14.61
$\zeta = 0$	0.32/0.60	0.24/0.33	0.25/0.51	0.21/0.38	0.18/0.27	0.24/0.42	8.28/12.90
<b>STGFormer (Ours)</b>	<b>0.27/0.56</b>	<b>0.11/0.17</b>	<b>0.22/0.45</b>	<b>0.16/0.31</b>	<b>0.14/0.24</b>	<b>0.18/0.35</b>	<b>7.35/11.39</b>



**Figure 5: We compute the statistics of the learned edge values and study the socio-temporal notion of the pair-wise agent interactions these edges encode. (a) The sparse distribution of learned non-trivial STG edge values over the distance between the two agents the edges connect. Edges with a learned value of 1 mostly link two agents no more than of 120-pixel, meaning the behaviors of most agents outside this range are deemed less essential for trajectory predictions. (b) The distribution of edges with a value 1 over time. “Short-term” edges are more likely to be considered than their “long-term” counterparts for trajectory predictions. (c) The flip of edge from 0 to 1 is highly correlated with the event where two agents “approaching” each other, and vice-versa, the event that two agents “diverging” from each other are captured by the edge value flipping from 1 to 0. (d) Symmetry and spatial localization suggest that most events occur when one agent is entering or leaving another agent’s 30-pixel-perimeter.**

that statistically, most edge values flip from 0 to 1 when two agents “approaching” each other, and vice-versa, most values flip from 1 to 0 when two agents are walking away from each other. The symmetry and heavy tail distribution in Fig. 8E suggest that there exists a single underlying localized perimeter surrounding each agent, where when another agent enters or leaves it, the corresponding edge flips its value accordingly.

#### 4.4 Ablation Study

We conducted the following ablation studies to examine the effectiveness of the model design.

**Learning socio-temporal correlations over time: w/o  $G^t$ .** To highlight the advantages of learning the socio-temporal dependencies, we compared two models, one with  $G^t$  and one without it. A design with a predefined w/o  $G^t$  baseline is similar to a transformer decoder [30]. Again, a significant contribution of our work is to allow this entity to be learnable. We conduct ablation studies to show the benefit of learning data-driven STG over fixed ones for trajectory predictions (Tab. 3). This suggests that understanding socio-temporal correlations enables STGformer to generalize over various datasets over a more stable distribution.

**Long-term v.s. short-term correlations: Short-term- $G^t$ .** To verify that the information carried by non-trivial edges facilitates trajectory prediction, we also intentionally mask off this information in our ablation experiment. The result is reported in Tab. 3, annotated as “w/o social”. It can be easily seen that if relevant information learned by our model is intentionally ignored, the performance of our model drops, which means that the information uniquely learned by our model can be used to improve prediction accuracy.

Moreover, we divide all edges of  $G^t$  by the time duration they span into “short-term” and “long-term”, respectively. Tab. 3 reports the results if the samples only from the same group are used for prediction. The results show that the deeper we look back into history to consider potential correlations between past observations with present ones, the better prediction we make, and reasonably our performance gain becomes marginal as excessively old observations tend to carry obsolete information.

Specifically, we adjust the duration of the “look-back” window of  $G^t$  from  $\mathbb{R}^{n \times nt}$  to  $\mathbb{R}^{n \times n}$ . Such a design suggests connecting an individual with his/her cause persons at only 1 step before, respectively. Our approach drastically advances short-term  $G^t$  by improving from 9.04/14.29 to 7.35, 11.39 on SDD, and from 0.34/0.45 to 0.18, 0.35 on ETH/UCY. The results suggest that the long-term correlations provide valuable information for future predictions.

**Time varying property: Stationary- $G^t$ .** The performance of STGFormer is considerably better than its counterpart design with a learned-once-applied-to-all stationary- $G^t$  baseline, which assumes that the system is time-invariant. Under the time-invariant assumption,  $G^t$  at the first step is learned and applied to analyze all subsequent observations. The results suggest that the system governing human trajectories is time-varying.

**Generating  $G^t$ : w/o LP (learned prior).** We compare models with and without LP baseline. More specifically, the model without LP drops  $p_{\Psi}(G^t | G^{0:t-1})$  by setting  $G^t \sim \mathcal{N}(0, I)$ . Tab. 3 illustrates the advantage of STGFormer over the model without LP, indicating the

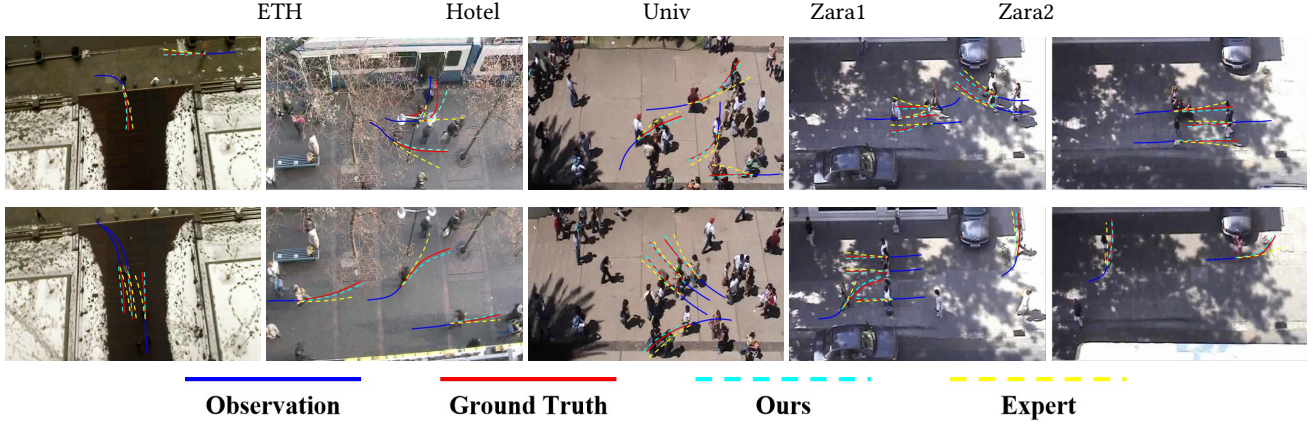


Figure 6: Qualitative results on ETH/UCY datasets. Each column overlays predictions for several individual agents. For each person, we show the path history (blue line), ground truth future (red line), and predictions from Expert [52] (yellow dashed line) and ours STGFormer (cyan dashed line).




Figure 7: Socio-temporal dynamics: the flip of edge values from 0 to 1 and 1 to 0 indicate the occurrence of “approaching” and “diverging”, respectively. For example, the initial edge value between two agents at  $t = 6$  is learned to be 0, indicating they are uncorrelated. Afterwards, at  $t = 12$ , when they are “approaching” each other close enough, the edge value flips from 0 to 1. At time  $t = 18$ , when they “diverge” from each other far enough, the corresponding edge value flips from 1 back to 0.

Table 4: Ablation studies on the architectures of the ADE / FDE scores over 20 predictions on the ETH and UCY datasets. The last row is our STGFormer results.

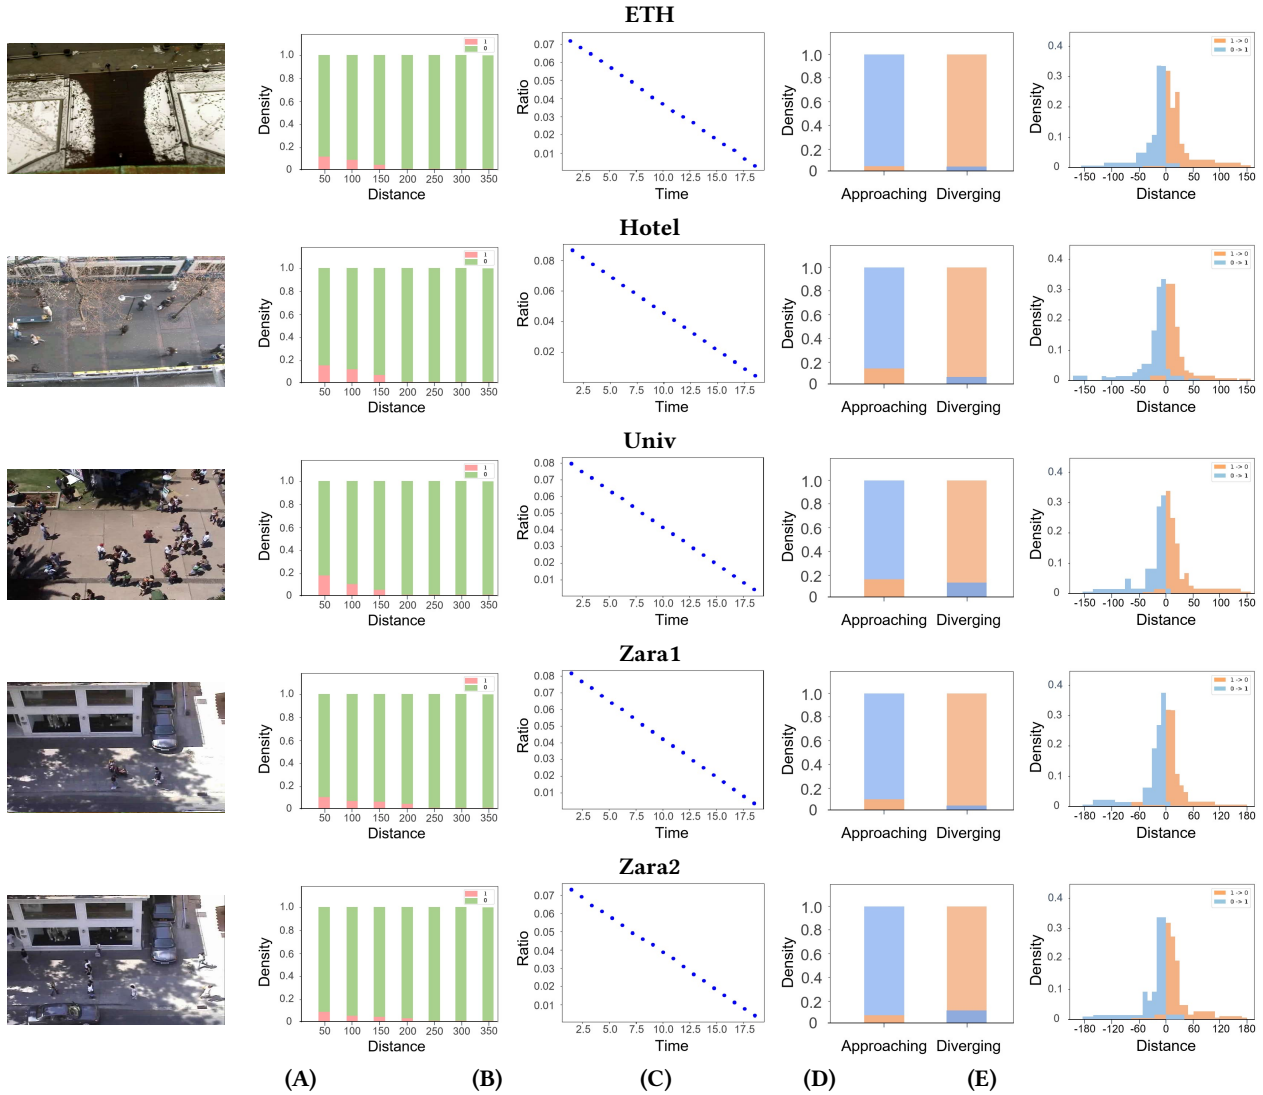
$G^t$	$x^t$	ETH	Hotel	Univ.	Zara1	Zara2	AVG	SDD
LSTM	GAT	0.55/0.76	0.49/0.64	0.50/0.65	0.38/0.54	0.42/0.57	0.47/0.63	12.71/19.84
LSTM	Transformer	0.40/0.62	0.31/0.38	0.36/0.54	0.32/0.50	0.30/0.42	0.34/0.45	9.55/14.10
Transformer	GAT	0.31/0.61	0.22/0.25	0.28/0.48	0.19/0.41	0.20/0.36	0.24/0.41	8.12/13.27
<b>Transformer</b>	<b>Transformer</b>	<b>0.27/0.56</b>	<b>0.11/0.17</b>	<b>0.22/0.45</b>	<b>0.16/0.31</b>	<b>0.14/0.24</b>	<b>0.18/0.35</b>	<b>7.35/11.39</b>

necessity of capturing the underlying temporal correlation of  $G^t$  over  $t$  upon learning from  $p_{\Psi}(G^t|G^{0:t-1})$ .

**STG sparsity term:**  $\zeta = 0$ . The sparsity term means  $\zeta \|\mathcal{G}^{0:T}\|_0$ . We favor a succinct socio-temporal interpretation for our observation. Setting  $\zeta = 0$  allows us to minimize the weight of edges encoding unnecessary correlations. We conducted an ablation study to verify this point that trains the model without  $\zeta \|\mathcal{G}^{0:T}\|_0$ . Tab. 3 shows that with the sparsity regularizer, the learner can generate better/more succinct interpretations, which lead to more accurate predictions by our design.

**Transformer Architecture.** In order to further verify the effectiveness of our architecture choices, we tested models with a different backbone, choosing LSTM [7] and GAT [38] as alternatives in contrast to the transformer. To ensure a fair comparison, we used a single LSTM layer with a hidden size of 256 to learn the prior  $p_{\Psi}(G^t|G^{0:t-1})$  and posterior  $q_{\Theta}(G^t|G^{0:t-1}, \mathbf{x}^{0:t})$ . The trajectory model, implemented with a Two-stacked GAT, is set to a size of 512. Both GAT layers use eight attention heads.





**Figure 8:** We compute the statistics of the learned edge values and study the socio-temporal notion of the pair-wise agent interactions these edges encode over the ETH/UCY dataset. (A) The different scenario settings in ETH/UCY dataset. (B) The sparse distribution of learned non-trivial STG edge values over the distance between the two agents the edges connect. Edges with a learned value of 1 mostly link two agents no more than 200 pixels, meaning the behaviors of most agents outside this range are deemed less essential for trajectory predictions. (C) The distribution of edges with a value of 1 over time. “Short-term” edges are more likely to be considered than their “long-term” counterparts for trajectory predictions. (D) The flip of edge from 0 to 1 is highly correlated with the event where two agents “approaching” each other, and vice-versa. The event that two agents “diverging” from each other is captured by the edge value flipping from 1 to 0. (E) Symmetry and spatial localization suggest that most events occur when one agent enters or leaves another agent around a 60-pixel-perimeter.

## 5 CONCLUSION

In this paper, we introduce and test a STGFormer for learning to better foresee human trajectories. The STGFormer models the joint distributions that are formulated in Eqn. 2. This formulation allows us to grasp the socio-temporal graph structures underlying representations of human trajectories. Experiment results demonstrate that our model delivers better performance in the task of path forecasting compared with other state-of-the-art trajectory-based

approaches. We believe that our STGFormer can pave the way for future improvements. One promising direction is to extend our framework to include context information, such as images, and trajectories, to discover the correlations between human agents and the environment in terms of socio-temporal interactions.

## REFERENCES

- [1] Alexandre Alahi, Krathar Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. 2016. Social LSTM: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 961–971.
- [2] Inhwan Bae, Jin-Hwi Park, and Hae-Gon Jeon. 2022. Learning Pedestrian Group Representations for Multi-modal Trajectory Prediction. In *European Conference on Computer Vision*. Springer, 270–289.
- [3] Inhwan Bae, Jin-Hwi Park, and Hae-Gon Jeon. 2022. Non-Probability Sampling Network for Stochastic Human Trajectory Prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6477–6487.
- [4] Guangyi Chen, Junlong Li, Jiwen Lu, and Jie Zhou. 2021. Human Trajectory Prediction via Counterfactual Analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 9824–9833.
- [5] Yilun Du and Igor Mordatch. 2019. Implicit Generation and Modeling with Energy Based Models. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc.
- [6] Francesco Giuliani, Irtiza Hasan, Marco Cristani, and Fabio Galasso. 2021. Transformer networks for trajectory forecasting. In *2020 25th international conference on pattern recognition (ICPR)*. IEEE, 10335–10342.
- [7] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press.
- [8] Tianpei Gu, Guangyi Chen, Junlong Li, Chunze Lin, Yongming Rao, Jie Zhou, and Jiwen Lu. 2022. Stochastic Trajectory Prediction via Motion Indeterminacy Diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [9] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. 2018. Social GAN: Socially Acceptable Trajectories With Generative Adversarial Networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [10] Fei Huang, Hao Zhou, Yang Liu, Hang Li, and Minlie Huang. 2022. Directed Acyclic Transformer for Non-Autoregressive Machine Translation. In *Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (Eds.). PMLR, 9410–9428.
- [11] Yingfan Huang, Huikun Bi, Zhaoxin Li, Tianlu Mao, and Zhaoyi Wang. 2019. STGAT: Modeling Spatial-Temporal Interactions for Human Trajectory Prediction. In *The IEEE International Conference on Computer Vision (ICCV)*.
- [12] Vineet Kosaraju, Amir Sadeghian, Roberto Martín-Martín, Ian Reid, Hamid Rezatofighi, and Silvio Savarese. 2019. Social-bigat: Multimodal trajectory forecasting using bicycle-gan and graph attention networks. *Advances in Neural Information Processing Systems* 32 (2019).
- [13] Parth Kothari, Brian Siffringer, and Alexandre Alahi. 2021. Interpretable Social Anchors for Human Trajectory Forecasting in Crowds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 15556–15566.
- [14] Mihee Lee, Samuel S Sohn, Seonghyeon Moon, Sejong Yoon, Mubbasir Kapadia, and Vladimir Pavlovic. 2022. Muse-VAE: multi-scale VAE for environment-aware long term trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2221–2230.
- [15] Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. 2007. Crowds by example. In *Computer graphics forum*, Vol. 26. Wiley Online Library, 655–664.
- [16] Jiachen Li, Fan Yang, Hengbo Ma, Srikanth Malla, Masayoshi Tomizuka, and Chiho Choi. 2021. Rain: Reinforced hybrid attention inference network for motion forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 16096–16106.
- [17] Jiachen Li, Fan Yang, Masayoshi Tomizuka, and Chiho Choi. 2020. Evolvegraph: Multi-agent trajectory prediction with dynamic relational reasoning. *Advances in neural information processing systems* 33 (2020), 19783–19794.
- [18] Lihuan Li, Maurice Pagnucco, and Yang Song. 2022. 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*. 2231–2241.
- [19] Yunzhu Li, Antonio Torralba, Anima Anandkumar, Dieter Fox, and Animesh Garg. 2020. Causal Discovery in Physical Systems from Videos. In *Advances in Neural Information Processing Systems*. 9180–9192.
- [20] Yuejiang Liu, Riccardo Cadei, Jonas Schweizer, Sherwin Bahmani, and Alexandre Alahi. 2022. Towards Robust and Adaptive Motion Forecasting: A Causal Representation Perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 17081–17092.
- [21] Yuejiang Liu, Qi Yan, and Alexandre Alahi. 2021. Social nce: Contrastive learning of socially-aware motion representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 15118–15129.
- [22] Lars Lorch, Jonas Rothfuss, Bernhard Schölkopf, and Andreas Krause. 2021. DiBS: Differentiable Bayesian Structure Learning. *Advances in Neural Information Processing Systems* 34 (2021).
- [23] Karttikeya Mangalam, Yang An, Harshayu Girase, and Jitendra Malik. 2021. From goals, waypoints & paths to long term human trajectory forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 15233–15242.
- [24] Karttikeya Mangalam, Harshayu Girase, Shreyas Agarwal, Kuan-Hui Lee, Ehsan Adeli, Jitendra Malik, and Adrien Gaidon. 2020. It is not the journey but the destination: Endpoint conditioned trajectory prediction. In *European Conference on Computer Vision*. Springer, 759–776.
- [25] Jacinto C Nascimento, Mário AT Figueiredo, and Jorge S Marques. 2009. Trajectory classification using switched dynamical hidden Markov models. *IEEE Transactions on Image Processing* 19, 5 (2009), 1338–1348.
- [26] Bo Pang, Tianyang Zhao, Xu Xie, and Ying Nian Wu. 2021. Trajectory prediction with latent belief energy-based model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11814–11824.
- [27] Judea Pearl, Madelyn Glymour, and Nicholas P Jewell. 2016. *Causal inference in statistics: A primer*. John Wiley & Sons.
- [28] S Pellegrini, A Ess, K Schindler, and L van Gool. 2009. You'll never walk alone: Modeling social behavior for multi-target tracking. In *2009 IEEE 12th International Conference on Computer Vision (ICCV)*.
- [29] Ruijie Quan, Linchao Zhu, Yu Wu, and Yi Yang. 2021. Holistic LSTM for pedestrian trajectory prediction. *IEEE transactions on image processing* 30 (2021), 3229–3239.
- [30] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1, 8 (2019), 9.
- [31] Alexandre Robicquet, Amir Sadeghian, Alexandre Alahi, and Silvio Savarese. 2016. Learning social etiquette: Human trajectory understanding in crowded scenes. In *European conference on computer vision*. Springer, 549–565.
- [32] Yulia Rubanova, Ricky T. Q. Chen, and David K Duvenaud. 2019. Latent Ordinary Differential Equations for Irregularly-Sampled Time Series. In *Advances in Neural Information Processing Systems 32 (NeurIPS 19)*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., 5320–5330.
- [33] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. 2020. Trajectron++: Multi-agent generative trajectory forecasting with heterogeneous data for control. In *The European Conference on Computer Vision (ECCV)*.
- [34] Yang Song and Stefano Ermon. 2019. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems* 32 (2019).
- [35] Jianhua Sun, Yuxuan Li, Hao-Shu Fang, and Cewu Lu. 2021. Three steps to multimodal trajectory prediction: Modality clustering, classification and synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 13250–13259.
- [36] Li-Wu Tsao, Yan-Kai Wang, Hao-Siang Lin, Hong-Han Shuai, Lai-Kuan Wong, and Wen-Huang Cheng. 2022. Social-SSL: Self-supervised Cross-Sequence Representation Learning Based on Transformers for Multi-agent Trajectory Prediction. In *European Conference on Computer Vision*. Springer, 234–250.
- [37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.
- [38] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *International Conference on Learning Representations (ICLR)*.
- [39] Chuhua Wang, Yuchen Wang, Mingze Xu, and David J Crandall. 2022. Stepwise goal-driven networks for trajectory prediction. *IEEE Robotics and Automation Letters* 7, 2 (2022), 2716–2723.
- [40] Xinhao Weng, Boris Ivanovic, Kris Kitani, and Marco Pavone. 2022. Whose Track Is It Anyway? Improving Robustness to Tracking Errors With Affinity-Based Trajectory Prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 6573–6582.
- [41] Conghao Wong, Beihao Xia, Ziming Hong, Qimu Peng, and Xinge You. 2021. View Vertically: A hierarchical network for trajectory prediction via fourier spectrums. *arXiv preprint arXiv:2110.07288* (2021).
- [42] Minye Wu, Haibin Ling, Ning Bi, Shenghua Gao, Qiang Hu, Hao Sheng, and Jingyi Yu. 2020. Visual tracking with multiview trajectory prediction. *IEEE Transactions on Image Processing* 29 (2020), 8355–8367.
- [43] Chenxin Xu, Maosen Li, Zhenyang Ni, Ya Zhang, and Siheng Chen. 2022. GroupNet: Multiscale Hypergraph Neural Networks for Trajectory Prediction with Relational Reasoning. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [44] Chenxin Xu, Weibo Mao, Wenjun Zhang, and Siheng Chen. 2022. Remember Intentions: Retrospective-Memory-based Trajectory Prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6488–6497.
- [45] Mengyue Yang, Furui Liu, Zhitang Chen, Xinwei Shen, Jianye Hao, and Jun Wang. 2021. CausalVAE: Disentangled representation learning via neural structural causal models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9593–9602.
- [46] Shuai Yi, Hongsheng Li, and Xiaogang Wang. 2016. Pedestrian behavior modeling from stationary crowds with applications to intelligent surveillance. *IEEE*

- transactions on image processing* 25, 9 (2016), 4354–4368.
- [47] Ziyi Yin, Ruijin Liu, Zhiliang Xiong, and Zejian Yuan. 2021. Multimodal Transformer Networks for Pedestrian Trajectory Prediction.. In *IJCAL*. 1259–1265.
  - [48] Cunjun Yu, Xiao Ma, Jiawei Ren, Haiyu Zhao, and Shuai Yi. 2020. Spatio-temporal graph transformer networks for pedestrian trajectory prediction. In *European Conference on Computer Vision*. Springer, 507–523.
  - [49] Cunjun Yu, Xiao Ma, Jiawei Ren, Haiyu Zhao, and Shuai Yi. 2020. Spatio-temporal graph transformer networks for pedestrian trajectory prediction. In *European Conference on Computer Vision*. Springer, 507–523.
  - [50] Ye Yuan, Xinshuo Weng, Yanglan Ou, and Kris M Kitani. 2021. AgentFormer: Agent-Aware Transformers for Socio-Temporal Multi-Agent Forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 9813–9823.
  - [51] Jiangbei Yue, Dinesh Manocha, and He Wang. 2022. Human trajectory prediction via neural social physics. *arXiv preprint arXiv:2207.10435* (2022).
  - [52] He Zhao and Richard P Wildes. 2021. Where are you heading? dynamic trajectory prediction with expert goal examples. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7629–7638.