This WACV 2021 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

GyF

GraphTCN: Spatio-Temporal Interaction Modeling for Human Trajectory Prediction

Chengxin Wang* Shaofeng Cai* Gary Tan National University of Singapore

{wangcx, shaofeng, gtan}@comp.nus.edu.sg

Abstract

Predicting the future paths of an agent's neighbors accurately and in a timely manner is central to the autonomous applications for collision avoidance. Conventional approaches, e.g., LSTM-based models, take considerable computational costs in the prediction, especially for the long sequence prediction. To support more efficient and accurate trajectory predictions, we propose a novel CNN-based spatial-temporal graph framework GraphTCN, which models the spatial interactions as social graphs and captures the spatio-temporal interactions with a modified temporal convolutional network. In contrast to conventional models, both the spatial and temporal modeling of our model are computed within each local time window. Therefore, it can be executed in parallel for much higher efficiency, and meanwhile with accuracy comparable to best-performing approaches. Experimental results confirm that our model achieves better performance in terms of both efficiency and accuracy as compared with state-of-the-art models on various trajectory prediction benchmark datasets.

1. Introduction

Trajectory prediction is a fundamental and challenging task, which needs to forecast the future path of the agents in autonomous applications, such as autonomous vehicles, socially compliant robots, agents in simulators, to navigate in a shared environment. With multi-agent interaction in these applications, the agents are required to respond timely and precisely to the environment for collision avoidance. Therefore, the ability of the agents to predict the future paths of their neighbors in an efficient and accurate manner is thus much needed. Although recent works [25, 33, 17, 28] have achieved great improvement in modeling complex social interactions among agents to generate accurate future paths, trajectory prediction is still a challenging task, where the deployment of the prediction models in real-world applica-



Figure 1: The illustration of trajectory prediction in a crowd. The solid blue lines are the observed trajectories, and the dash blue lines indicate the plausible future path. The influence levels α between two agents are different based on their relative movement trends, e.g., α_{ij} and α_{ji} between agent *i*, and agent *j* are different.

tions is mostly restricted by its high computational cost and long inference time. For example, some small robots are only equipped with limited computing devices that can not afford the high inference cost with existing solutions.

In particular, the trajectory prediction is typically modeled in two dimensions, i.e., the temporal dimension and the spatial dimension, which is illustrated in Fig. 1. The temporal dimension models the historical movement dynamics for each agent. Most of the state-of-the-art approaches [1, 12, 25, 17, 19, 34] have focused on Recurrent Neural Networks (RNNs), to capture such sequence dynamics since RNNs are designed for sequence modeling. However, besides the training difficulties of gradient vanishing and exploding [30] in modeling long sequential data, both training and inference of RNN models are notoriously slow compared with their feed-forward counterparts, e.g., Convolutional Neural Networks (CNNs). This is largely due to the fact that each hidden state of RNNs is dependent on the previous inputs and hidden states. As a consequence, the prediction of RNNs is produced sequentially, and thus not parallelizable.

The spatial dimension models the human-human interaction, i.e., interactions between the agent and its neighbors.

^{*}These authors contributed equally to this work.

There are mainly three categories of methods proposed to capture the spatial interaction, including the distancebased [1, 12, 25], attention-based [33, 7, 42, 19] and graphbased [17, 21, 51, 28] approaches. Distance-based approaches introduce a social pooling layer to summarize the crowd interactions, while the attention-based approaches instead dynamically generate the importance of neighbors using soft attention. The graph-based approaches model the agents' representation with a graph and utilize graph neural network, e.g., GCN [20, 51, 28] or GAT [41], to capture the spatial interaction features of agents, which is empirically more intuitive and effective in modeling complex social interactions. However, existing graph-based approaches are mostly based on the simple aggregation of neighbor features or their absolute geometric distance, which neglects the relative relation between agents in the spatial modeling.

To improve effectiveness and efficiency, we propose a novel Graph-based Temporal Convolutional Network (GraphTCN), to capture the spatial and temporal interaction for trajectory prediction. In the temporal dimension, different from RNN-based methods, we adopt a modified gated convolutional network (TCN) to capture the temporal dynamics for each agent. The gated highway mechanism introduced to CNNs dynamically regulate the information flow by focusing on more salient features, and the feedforward nature of CNN makes it more tractable in training and parallelizable for much higher efficiency both in training and inference. In the spatial dimension, we propose an edge feature based graph attention network (EFGAT) with skip connections and the gate mechanism for each time step to model the spatial interaction between the agents. Specifically, nodes in the graph represent agents, and edges between agents denote their relative spatial relation. EFGAT learns the adjacency matrix, i.e., the spatial interaction, of the graph adaptively. Together, the spatial and temporal modules of GraphTCN support more effective and efficient modeling of the interactions within each time step between agents and across the time steps for each and every agent. We summarize our main contributions as follows:

- We propose an edge feature based graph attention network (EFGAT), which introduces the relative spatial relation as prior knowledge, to capture the spatial interaction adaptively with attention.
- We propose to model the temporal interactions with a gated convolutional network (TCN), which empirically proves to be more efficient and effective.
- Our spatial-temporal framework achieves better performance compared with best-performing approaches. Specifically, we reduce the average displacement error by 19.4% and final displacement error by 13.6% with 5 times less predicted paths, and achieves up to 5.22x wall-clock time speedup over existing solutions.

We organize this paper as follows: in Section 2, we introduce the background and discuss related works in detail. Our GraphTCN framework is introduced in Section 3. Then in Section 4, results of GraphTCN measured in both accuracy and efficiency are compared with state-of-the-art approaches. Finally, Section 5 concludes the paper.

2. Related Work

Human-Human Interactions. Research in the crowd interaction model can be traced back to the Social Force model [14], which adopts the nonlinearly coupled Langevin equations to represent the attractive and repulsive forces for human movement in the crowed scenarios. Similar handcrafted approaches[39, 2, 44] have proved successful in crowd simulation [16, 32], crowd behavior detection [27], and trajectory prediction [47]. Recent works instead investigate deep learning techniques to capture the interaction between agents and neighbors. The distance-based approaches [1, 12, 26] either adopt the grid-based pooling or symmetric function to aggregate the hidden states from neighbors or encode the geometric relation between the agents. Different from the distance-based methods, attention-based approaches [33, 42, 7, 51] provide better crowd modeling since they differentiate the importance of neighbors by soft attention or gating mechanisms. More recent works [17, 21, 28] adopts graph-based networks to learn the social interaction by aggregating neighborhood features adaptively with the adjacency matrix, which provides an effective way to represent the pedestrian's topology in a shared space. Social-STGCNN [28] captures the spatial relation by introducing a kernel function on the weighted adjacency matrix; STGAT [17, 21] adopts GAT directly on the LSTM hidden states to capture the spatial interaction between pedestrians. However, Social-STGCNN only focuses on distance features between agents, and STGAT simple aggregates neighbor features. EGNN [9] incorporates the edge feature into the graph attention mechanism to exploit richer graph information. However, EGNN neglects the relative relation between pedestrians. We propose to model the pedestrian interactions with a novel edge feature based graph network, which integrates the relative distance feature into graph attention to learn an adaptive adjacency matrix for the most salient interaction information.

Pattern-based Sequence Prediction. Sequence prediction refers to the problem of predicting the future sequence using historical information. Recently, pattern-based methods prevails for many sequence prediction tasks, e.g., speed recognition [40, 4, 11], activity recognition [6, 18], and natural language processing [3, 37, 8]. In particular, trajectory prediction can be formulated as a sequence prediction task, which uses historical movement patterns of the agent to predict the future path. Most trajectory prediction methods adopt recurrent neural networks (RNNs), e.g.,

Long Short-Term Memory (LSTM) networks [15], to capture the temporal movement. However, RNN-based models suffer from gradient vanishing and exploding in training and focus more on recent inputs during prediction, especially for long input sequences. Many sequence prediction works [40, 45] instead adopt convolutional neural networks (CNNs). The convolutional networks can effectively capture long-term dependency and greatly improve prediction efficiency. The superiority of CNN-based methods can be largely attributed to the convolutional operation, which is independent of preceding time-steps and thus can process in parallel. The recent work [29] proposes a compact CNN model to capture the temporal information, and the results confirm that the CNN-based model can yield competitive performance in trajectory prediction. However, it fails to model the spatial interaction between pedestrians. In this work, we propose to capture the spatial interaction with EF-GAT and introduce gated convolutional networks to better capture the temporal dynamics.

Graph Networks for Trajectory Prediction. Many studies adopt spatial-temporal graph neural networks (STGNNs) for the sequence prediction task, such as action recognition [48, 35], taxi demand prediction [50], and traffic prediction [49]. Specifically, the sequence can be formulated as a sequence of graphs of nodes and edges, where nodes correspond to agents and edges denote their interactions. The sequence can then be effectively modeled with the spatial-temporal graph network. Likewise, the trajectory prediction task can be modeled with the spatial-temporal graph network [42, 43, 17, 24, 28]. In particular, the prediction task needs to be modeled in two dimensions, i.e., the spatial dimension and the temporal dimension. The spatial dimension models the interaction between the agent and its neighbors, and the temporal dimension models the historical trajectory for each agent. Specifically, each node in the graph represents one pedestrian of a scene, and each edge between two nodes captures the interaction between the two corresponding pedestrians. For example, social attention [42] models each node with the location of the agent, and edge with the distance between pedestrians, where the spatial relation is modeled with an attention module and then the temporal with RNNs. Similarly, [43] constructs the STGNN with Edge RNN and Node RNN based on the location. STGAT [17] adopts GAT to capture the spatial interaction by assigning different importance to neighbors and adopts extra LSTMs to capture the temporal information. The major limitation of these methods lies in capturing the spatial interaction along the temporal dimension. Notably, the future path of an agent is not only dependent on the current position but meanwhile the neighbors'. However, the information of such spatial interaction may be lost during the aggregation of the node features along the temporal dimension using RNN-based models. Different from RNN-based methods, Social-STGCNN [28] and Graph WaveNet [45] adopts CNNs to alleviate parameter inefficiency and demonstrate the capability of CNNs in the temporal modeling of long sequences. In this paper, we propose an enhanced temporal convolutional network to integrate both the temporal dynamics of the agent and the spatial interactions.

3. GraphTCN

The goal of trajectory prediction is to predict the future paths of all agents that are present in a scene. Naturally, the future path of an agent depends on its historical trajectory, i.e., the temporal interaction, and is influenced by the trajectories of neighboring agents, i.e., the spatial interaction. Therefore, the trajectory prediction model needs to take into account both features when modeling the spatial and temporal interactions for the prediction.

Problem Formulation Formally, the trajectory prediction can be defined as follows: given N pedestrians observed in a scene with T_{obs} steps, the position of a single pedestrian $i \in \{1, \ldots, N\}$ at the time step $t \in \{1, \ldots, T_{obs}\}$ is denoted as $X_i^t = (x_i^t, y_i^t)$. Therefore, the observation positions of the pedestrian X_i can be represented as $X_i^{1:T_{obs}} = X_i^1, X_i^2, ..., X_i^{T_{obs}}$. The goal of trajectory prediction is then to predict all the future positions \hat{Y}_i^t ($t \in \{T_{obs+1}, \ldots, T_{pred}\}$) concurrently.

3.1. Overall Framework

As illustrated in Fig. 2(a), GraphTCN comprises three key modules, including the edge feature graph attention (EFGAT) module, temporal convolutional (TCN) module, and a decoder. First, we embed the absolute positions and relative positions of each pedestrian into a fixed-length hidden space and feed these trajectories features into the EF-GAT module. The residual learning mechanism and skip connections [13] are incorporated into the network to facilitate the gradient backpropagation and encourage the intermediate feature reusage. The TCN module is a feed-forward one-dimensional convolutional network with a gating activation unit [40] for capturing the most salient features. Finally, the decoder module produces future trajectories of all pedestrians. We elaborate on the details of each module of GraphTCN in the following sections.

3.2. EFGAT Module for Spatial Interaction

The EFGAT module shown in Fig. 2(b) is designed to encode the spatial interaction between pedestrians with graph attentional layers and graph residual connections. Formally, pedestrians within the same time step can be formulated as a directed graph $\mathcal{G} = (\mathcal{V}^t, \mathcal{E}^t)$, where each node $v_i^t \in \mathcal{V}^t, i \in \{1, \ldots, N\}$ corresponds to the *i*-th pedestrian, and the weighted edge $(v_i^t, v_j^t) \in \mathcal{E}^t$ represents the human-human interaction from pedestrian *i* to *j*. The adjacency



Figure 2: (a) The overview of GraphTCN, where EFGAT captures the spatial interaction between agents for each time step and is based on the historical trajectory embedding. TCN further captures the temporal interaction across time steps. The decoder module then produces multiple socially acceptable trajectories for all the agents simultaneously. (b) EFGAT captures the spatial salient information with graph attentional layers (GAL) and skip connections.

matrix $A^t \in \mathbb{R}^{N \times N}$ of \mathcal{G} thus represents the spatial relationships between pedestrians.



Figure 3: An illustration of the graph attentional layer with 5 nodes employed by our EFGAT module. The attention α_{ij}^t between node *i* and its neighbors is learned from their embedding features and relative spatial relation \hat{A}_{ij}^t .

We represent the spatial relation of nodes as an asymmetric, non-negative matrix in this task since the influence between the agents should be different based on their relative movement behavior. Therefore, instead of constructing graphs with undirected spatial distance, we introduce relative spatial location as prior edge feature knowledge of the adjacency matrix:

$$\hat{A}_{ij}^{t} = \phi_s \left(x_i^t - x_j^t, y_i^t - y_j^t; W_s \right)$$
(1)

where $\phi_s(\cdot)$ embeds the relative distance features to a higher dimension F_1 with a linear transformation, and \mathbf{W}_s is the embedding weight. We feed the learnable edge weights and node features into graph attentional layers illustrated in Fig. 3 to capture the spatial interaction:

$$\alpha_{ij}^{t} = \frac{\exp\left(\sigma\left(\mathbf{W}_{1}\hat{h}_{i}^{t} + \mathbf{W}_{2}\hat{h}_{j}^{t} + \hat{A}_{ij}^{t}\right)\right)}{\sum_{k \in \mathcal{N}_{i}}\exp\left(\sigma\left(\mathbf{W}_{1}\hat{h}_{i}^{t} + \mathbf{W}_{2}\hat{h}_{k}^{t} + \hat{A}_{ij}^{t}\right)\right)} \quad (2)$$

where $\hat{h}_i^t \in \mathbb{R}^{F_1}$ is the node input feature for pedestrian i at time step t, F_1 is the dimension of node features, \mathcal{N}_i is the set of neighbors for node i in the graph, $\sigma(\cdot)$ is the LeakyReLU activation, and $\mathbf{W}_1, \mathbf{W}_2$ are learnable weights. In this way, α_{ij}^t gives the importance weight of neighbor j to pedestrian i dynamically via the self-attention mechanism. The gating function empirically proves to be effective in controlling the signal bypassing [40, 46, 5]. We therefore adopt the gated activation unit to dynamically regulate the information flow and select salient features:

$$g_i^t = g\left(\mathbf{W}_h \hat{h}_i^t + b_h\right) \odot \left(\mathbf{W}_h \hat{h}_i^t + b_h\right)$$
(3)

where $g(\cdot)$ is the *tanh* activation function, \mathbf{W}_h is the affine transformation parameters, b_h is the bias, and \odot denotes the element-wise multiplication. This can be understood as a multiplicative skip connection which facilitates gradients flow through layers [5]. To stabilize the self-attention process [41, 45], we adopt the multi-head attention mechanism:

$$h_i^t = \prod_{k=1}^K \sigma\left(\sum_{j \in \mathcal{N}_i} \alpha_{kij}^t g_j^t\right) + \mathbf{R}(\hat{h}_i^t; \mathbf{W}_{\mathbf{r}}) \qquad (4)$$

where \mathbf{W}_r is the learnable parameters, || denotes concatenation, and K is the number of attention heads. $\mathbf{R}(\cdot)$ denotes the graph residual term [40, 41, 45]. We name the proposed multi-head graph attention layers as GAL, which can be stacked multiple times for better modeling spatial relation (e.g., twice as in Fig. 2(b)). Subsequently, we can obtain final node representations of $\mathbf{h} = \{h_1, h_2, \ldots, h_N\}$, where $h_i \in \mathbb{R}^{T_{obs} \times (K \cdot F_1)}$ captures the aggregated spatial interaction between pedestrian *i* and all the neighbors at each time step. Therefore, the EFGAT module can learn a self-adaptive adjacency matrix that captures the relative importance for different pedestrians.

3.3. TCN for Spatial and Temporal Interaction Modeling

The movement pattern of a pedestrian is greatly influenced by the historical trajectory and the moving patterns of neighboring pedestrians. We therefore propose to capture the spatial and temporal interaction between pedestrians using a modified temporal convolution network (TCN), which is illustrated in Fig. 4.



Figure 4: (a) An illustration of TCN with a stack of 3 convolution layers of kernel size 3. The input \mathbf{h}_{in} (i.e., $\mathbf{h}^{(0)}$) contains the spatial information captured by the preceding EFGAT modules. The output of TCN \mathbf{h} is collected by concatenating \mathbf{h}_{out} (i.e., $\mathbf{h}^{(L)}$) across time. (b) The gating function in each of the TCN layers to control the bypass signals.

The network shown in Fig. 4(a) can be regarded as a short-term and long-term encoder, where lower convolution layers focus on local short-term interactions, while in higher layers, long-term interactions are captured with a larger receptive field. For example, if the kernel size of the TCN is k, the receptive field size in the *l*-th layer is $(k-1)\cdot l+1$, which increases linearly ascending layers. Therefore, the top layer of TCN captures interactions within a longer time span. Since the order of the input is important in the sequence prediction task, we therefore adopt the left padding of size k - 1 instead of symmetric padding for the convolution, where each convolution output convolves over the input of the corresponding time step and the preceding k - 1 time steps. The output size of each convolution then remains the

same as the input. In each layer of TCN 4 in Fig.(b), the gated activation unit utilizes two non-linear functions to dynamically regulate the information flow formed as:

$$\mathbf{h}^{(l+1)} = g\left(\mathbf{W}_{g}^{(l)} * \mathbf{h}^{(l)}\right) \odot \sigma\left(\mathbf{W}_{f}^{(l)} * \mathbf{h}^{(l)}\right) \quad (5)$$

where $\mathbf{h}^{(0)}$ is the output \mathbf{h} from the EFGAT modules, $\mathbf{h}^{(l)} \in \mathbb{R}^{N \times T_{obs} \times F_2}$, \mathbf{W}_g and \mathbf{W}_f are the learnable 1Dconvolution parameters, and $\sigma(\cdot)$ denotes the *sigmoid* function. The final output of the TCN module can be denoted as $\mathbf{\vec{h}} \in \mathbb{R}^{N \times T_{obs} \times F_2}$. In this way, the embedding vector \vec{h}_i captures all the spatial-temporal interaction between the *i*th pedestrian and its neighbors. We note that TCN can handle much longer input sequences with the dilated convolution [40], which is more efficient than RNN-based methods.

3.4. Future Trajectory Prediction

In real-world applications, given the historical trajectory, there are multiple plausible paths of future movements. We also model such uncertainty of the final movement in our decoder module for the trajectory prediction.

Following STGAT [17], the decoder module produces multiple socially acceptable trajectories by introducing a shared random noise $z \in \mathbb{R}^{T_{obs} \times F_3}$, which is concatenated with the spatial-temporal embedding $\mathbf{\vec{h}}$, as part of the decoder input. Specifically, the input of the decoder can be denoted as $\mathbf{\tilde{h}} \in \mathbb{R}^{N \times T_{obs} \times (F_2 + F_3)}$. We adopt a canonical MLP layer to generate the relative future locations $\Delta \hat{Y} \in \mathbb{R}^{N \times T_{pred} \times 2}$ and denote the architecture with such an MLP decoder as GraphTCN.

The predicted relative location $\Delta \hat{Y}$ is the relative position to the origin for all the pedestrians. We then convert relative positions to absolute positions \hat{Y} and adopt the variety loss as the loss function for training, which computes the minimum ADE loss among the M plausible trajectories:

$$\mathcal{L}_{\text{ADE}}(\hat{Y}) = \frac{\sum_{i=1}^{N} \sum_{t=1}^{T_{pred}} \left\| \hat{Y}_{i}^{t} - Y_{i}^{t} \right\|_{2}}{NT_{pred}} \tag{6}$$

$$\mathcal{L}_{\text{variety}} = \min_{m} \left(\mathcal{L}_{ADE} \left(\hat{Y}^{(m)} \right) \right) \tag{7}$$

where Y is the ground truth, $\hat{Y}^{(1)}, \ldots, \hat{Y}^{(M)}$ are the M plausible trajectories predicted. Although this loss function may lead to a diluted probability density function [38], we empirically find that it facilitates better predictions of multiple future trajectories.

We further integrate the deep generative strategy [36, 19, 26] adopted widely in multimodal prediction to enhance the decoder of our GraphTCN. Specifically, during training, we concatenate $\vec{\mathbf{h}}$ with the ground-truth future trajectories encoded by an MLP layer, and then further encode the two features with an MLP to produce μ and σ for the

noise distribution $\hat{z} = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma}), \hat{z} \in \mathbb{R}^{N \times F_4}$ following CVAE [36, 26]. Note that \hat{z} is randomly sampled from distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$ during inference. For the final relative location prediction, we again concatenate \hat{z} with $\vec{\mathbf{h}}$ and feed them into an MLP layer to produce $\Delta \hat{Y}$. We further introduce the KL divergence regularization term [10, 22] to stabilize the training process:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\text{variety}} + \lambda_2 D_{KL}(\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma}) \| \mathcal{N}(\mathbf{0}, \mathbf{I}))$$
(8)

GraphTCN with such a decoder is denoted as GraphTCN-G in the following experiments.

4. Experiments

In this section, we evaluate our GraphTCN on two world coordinates trajectory prediction datasets, i.e., ETH [31] and UCY [23], and compare the performance of GraphTCN with state-of-the-art approaches.

4.1. Datasets and Evaluation Metrics

ETH and UCY datasets comprise five outdoor environments that are recorded from a fixed top-view. The ETH dataset includes ETH and Hotel, and the UCY dataset consists of UNIV, ZARA1, and ZARA2. In these datasets, pedestrians exhibit complex behaviors, including nonlinear trajectories, moving from different directions, walking together, walking unpredictably, avoiding collisions, standing, etc. The crowd density of a single scene in each environment varies from 0 to 51 pedestrians per frame. All datasets are records at 25 frames per second (FPS), and the pedestrian trajectory is extracted at every 2.5 FPS.

We use two metrics to evaluate model performance: Average Displacement Error (ADE) defined in Equation 6, which is the average Euclidean distance between the predicted trajectory and the ground truth over all prediction time steps, and Final Displacement Error (FDE), which is the Euclidean distance between the predict position and the ground truth position at the final time step T_{pred} .

The model is trained with the leave-one-out policy [1, 12, 42, 17]. We produced 4 samples for the next 4.8 seconds (12 timesteps) based on 3.2 seconds (8 timesteps) observations.

4.2. Implementation Details

We train with Adam optimizer in 50 epochs with a learning rate of 0.0001. The node feature embedding size is set to 64. The EFGAT module comprises two graph attention layers with attention heads K = 2, 1, and a output dimension F1 = 16, 32 for the first and second GAL respectively. F_2 is also set to 16, and the noise z has a dimension of $F_3 = 4$. For the GraphTCN-G decoder, the ground truth trajectory during training is encoded into a dimension of 64 for F_4 . M is set to 4 and 20 for predicting 4 and 20 sample paths. All the LeakyReLU in our model has a negative slope of 0.2. λ_1 is set to 1, and λ_2 is set to 0.5 for the first 15 epochs and 0.2 for the rest epochs for GraphTCN-G.

4.3. Baselines

We compare our framework with the baselines and four state-of-the-art approaches: LSTM adopts the vanilla LSTM encoder-decoder model to predict the sequence of every single pedestrian. Social LSTM [1] builds on top of LSTM and introduces a social pooling layer to capture the spatial interaction between pedestrians. CNN [29] adopts the CNNs to predict the sequence. SR-LSTM [51] obtains the spatial influence by iteratively refining the LSTM hidden states through the gate and attention mechanism. Social GAN [12] improves over Social LSTM with socially generative GAN to generate multiple plausible trajectories. Trajectron [19] utilizes LSTM to capture the spatial and temporal relations and incorporates CVAE [36] to generate the distributions of future paths. SoPhie [33] introduces the social and physical attention mechanisms to an LSTM based GAN model. Social-STGCNN [28] is one of the SOTA approaches that utilizes CNNs to extract spatio-temporal features. STGAT [17] is one of the SOTA approaches which adopts vanilla GAT to model the spatial interactions and LSTMs to capture temporal interaction.

4.4. Quantitative Results

The results in Table 1 show that GraphTCN achieves consistently better performance compared with existing models on these benchmark datasets. Our model generates multiple trajectories at once for future trajectories. We can notice that GraphTCN achieves better prediction performance than other baselines with only four predictions instead of 20 as in most baselines, e.g., STGAT [17], with an ADE of 0.36 and an FDE of 0.72 on average. These results confirm that our GraphTCN yields competitive results even with less generated paths compared with previous approaches in terms of prediction accuracy, especially on the more complex dataset UNIV, ZARA1, and ZARA2.

Ablation Study. We evaluate each module of GraphTCN through ablation studies in Table 2. *w/o EGNN* refers to the model without the spatial module. *vanilla GAT* refers to the model with GAT as the spatial module while ignoring the relative relation between pedestrians in spatial modeling. *GraphTCN-G* refers to the model which integrates VAE for multi-modal future path prediction. The result demonstrates that introducing graph neural networks (GNN) into the framework can reduce ADE and FDE, and adding edge relative relation to GNN leads to further improvement. However, these spatial interactions can only improve performance mildly. We further scrutinize the dataset and attribute these findings to the fact that pedestrians seldom change their path suddenly to avoid their neighbors.

Method	ETH	HOTEL	UNIV	ZARA1	ZARA2	AVG
LSTM* [1]	1.09 / 2.41	0.86 / 1.91	0.61 / 1.31	0.41 / 0.88	0.52 / 1.11	0.70 / 1.52
Social-LSTM* [1]	1.09 / 2.35	0.79 / 1.76	0.67 / 1.40	0.47 / 1.00	0.56 / 1.17	0.72 / 1.54
CNN* [29]	1.04 / 2.07	0.59 / 1.27	0.57 / 1.21	0.43 / 0.90	0.34 / 0.75	0.59 / 1.22
SR-LSTM* [51]	0.63 / 1.25	0.37 / 0.74	0.51 / 1.10	0.41 / 0.90	0.32 / 0.70	0.45 / 0.94
Social-GAN [12]	0.81 / 1.52	0.72 / 1.61	0.60 / 1.26	0.34 / 0.69	0.42 / 0.84	0.58 / 1.18
Trajectron [19]	0.59 / 1.14	0.35 / 0.66	0.54 / 1.13	0.43 / 0.83	0.43 / 0.85	0.56 / 1.14
SoPhie [33]	0.70/1.43	0.76 / 1.67	0.54 / 1.24	0.30 / 0.63	0.38 / 0.78	0.54 / 1.15
Social-STGCNN [28]	0.64 / 1.11	0.49 / 0.85	0.44 / 0.79	0.34 / 0.53	0.30/0.48	0.44 / 0.75
STGAT [17]	0.65 / 1.12	0.35 / 0.66	0.52 / 1.10	0.34 / 0.69	0.29 / 0.60	0.43 / 0.83
GraphTCN $(M = 4)$	0.59 / 1.12	0.27 / 0.52	0.42 / 0.87	0.30 / 0.62	0.23 / 0.48	0.36 / 0.72
GraphTCN ($M = 20$)	0.39 / 0.71	0.21 / 0.44	0.33 / 0.66	0.21 / 0.42	0.17 / 0.43	0.26 / 0.51
GraphTCN-G $(M = 4)$	0.60 / 1.21	0.27 / 0.52	0.41 / 0.84	0.28 / 0.58	0.22 / 0.47	0.36 / 0.72
GraphTCN-G ($M = 20$)	0.39 / 0.75	0.18/0.33	0.30 / 0.60	0.20 / 0.39	0.16 / 0.32	0.25 / 0.48

Table 1: Quantitative results of GraphTCN compared with baseline approaches. Evaluation metrics are reported in ADE / FDE in meters (the lower numerical result is better). The * mark denotes the deterministic model, and the rest of the baseline approaches are stochastic models with M = 20 prediction samples.

Method	M = 4	M = 20
w/o EGNN	0.38 / 0.78	0.28 / 0.54
vanilla GAT	0.37 / 0.74	0.27 / 0.54
GraphTCN	0.36/0.72	0.26 / 0.51
GraphTCN-G	0.36 / 0.72	0.25 / 0.48

Table 2: Ablation studies of GraphTCN.

As a consequence, the temporal features already contain part of the spatial interactions for the prediction. Therefore, spatial information is less critical in the prediction. Meanwhile, compared with RNN-based approaches, GraphTCN can model the whole observed sequence better without losing important temporal information.

Inference Speed. We compare the inference speed of GraphTCN with state-of-the-art methods, including Social GAN [12], SR-LSTM [51], Social-STGCNN [28] and STGAT [17]. Table 3¹ reports the model inference time and the speedup factor compared with the Social GAN in wall-clock second. As can be observed from the results, GraphTCN achieves much faster inference compared with these baseline approaches. In particular, GraphTCN takes 0.00067 second inference time to generate 4 samples, which is 42.82 times and 5.22 times faster than Social-GAN and the most similar prior approach STGAT respectively.

4.5. Qualitative Analysis

We investigate the prediction results of our GraphTCN by visualizing and comparing the predicted trajectories

	Inference Time	Speed-up
Social-GAN [12]	0.02869	$1 \times$
Social-STGCNN [28]	0.00861	$3.33 \times$
STGAT [17]	0.00350	$8.20 \times$
Trajectron [19]	0.00081	$35.42 \times$
GraphTCN (M=4)	0.00066	43.47 ×
GraphTCN-G (M=4)	0.00067	$42.82 \times$
GraphTCN-G (M=20)	0.00075	$38.25 \times$

Table 3: The inference time and speedup of GraphTCN compared with baseline methods. The inference time is the average of the total inference steps per pedestrians. The results are reported on an Intel Core i9-9880H Processor.

with the best-performing approach STGAT in Fig. 5. We choose three different scenarios in which the complex interactions take place. The complex interactions include pedestrian standing, pedestrian merging, pedestrian following and pedestrian avoidance.

In Fig. 5, we can observe that GraphTCN achieves better performance on: firstly, *direction and speed*, from Fig. 5(a)(b), we find that trajectories generated by GraphTCN follow the same direction as the ground truth, while predictions of STGAT deviate from the path notice-ably. In Fig. 5(a), one pedestrian moves in an unexpected direction, and GraphTCN generates an acceptable prediction accordingly. Besides, GraphTCN generates plausible short trajectories to the stationary pedestrian and the pedestrian who moves slowly. Secondly, *collision-free future paths*, Fig. 5(b)(c) show that STGAT may fail to make satisfactory predictions when pedestrians come from different groups, while GraphTCN generates better prediction in scenarios where one pedestrian meets another group. In Fig. 5(b),

¹For a fair comparison, the reported time includes the data processing time since some approaches require extra time to construct the graph during inference. Note that we use the corresponding official implementations and settings for each model, and the batch size is one in the evaluation.



Figure 5: Comparison of our GraphTCN (M=4) and STGAT predictions with ground truth trajectories. To better illustrate the results, only part of the pedestrian trajectories are presented. The solid red line, solid blue line, and dashed yellow line denote the observed trajectory, ground truth future trajectory, and predicted trajectory, respectively.

GraphTCN can successfully produce predictions avoiding future collisions when the pedestrian moves in the same direction from an angle. Further, GraphTCN produces socially acceptable predictions even in the more complex scenario in Fig. 5(c) when the pedestrian departs for the opposite directions or walks towards the same direction.



Figure 6: Illustration of EFGAT attention weights. The solid green/red line is the trajectory, and the arrow indicates the trajectory direction. The circle color shows the attention at each time step, and the circle size corresponds to the attention weight. The green trajectory without circles denotes the target pedestrian.

Social attention. In Fig. 6, we illustrate the learned attention weights by the EGAT module. The results show that our model can capture the relative importance of the target's neighbors, and the attention weights between two pedestrians vary along the path. Further, the attention weight from pedestrian i to pedestrian j and pedestrian j to pedestrian i is different, which is not considered in existing approaches [17, 21, 51, 28]. Less attention weight:

in Fig. 6(a), the stationary pedestrian has less impact on its moving neighbors, and the model assigns small importance to the pedestrians far away from the target. *More significant influence*: our model assigns a higher attention weight to the pedestrian moving toward the target in Fig. 6(a), moving ahead or moving in the rear while having a higher velocity in Fig. 6(b), and moving from the opposite directions before meeting the target Fig. 6(c). These cases demonstrate that reasonable attention weights are successfully assigned to the target pedestrian's neighbors according to all pedestrian movement patterns in the scene.



Figure 7: Visualizations of diverse predicted trajectories. (a) and (b) show four trajectories produced by GraphTCN and GraphTCN-G, and (c) shows the 20 trajectories generated by STGAT.

Diverse trajectory predictions. Fig. 7 is the visualization of diverse predictions. The result shows that GraphTCN can generate the prediction closer to the ground truth even with a smaller number of samples and can make good predictions for the pedestrian, who have relatively unexpected behaviors. In this scenario, one pedestrian has the intention to change its direction from the observation, and GraphTCN can generate both the normal and unexpected predictions for it. And for other pedestrians who have a more consistent observation, the model produces future paths with normal behaviors. Further, from Fig. 7(b) and (c), the prediction area of GraphTCN is much smaller and precise than STGAT with 20 predicted trajectories.

5. Conclusion

In this paper, we proposed GraphTCN for trajectory prediction, which captures the spatial and temporal interaction between pedestrians effectively by integrating EFGAT to model their spatial interactions, and TCN to model both the spatial and temporal interactions. The proposed GraphTCN is completely based on feed-forward networks, which is more tractable during training, and achieves better prediction accuracy and higher inference speed compared with existing RNN-based solutions. Experimental results confirm that our GraphTCN outperforms state-of-the-art approaches on various benchmark datasets.

Acknowledgement This research is supported by National Research Foundation (NRF) Singapore (Award NRF2018AU-SG01).

References

- [1] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–971, 2016.
- [2] Gianluca Antonini, Michel Bierlaire, and Mats Weber. Discrete choice models of pedestrian walking behavior. *Transportation Research Part B: Methodological*, 40(8):667–687, 2006.
- [3] Kyunghyun Cho, Bart van Merrienboer, Çaglar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL, pages 1724–1734. ACL, 2014.
- [4] Jan Chorowski, Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. End-to-end continuous speech recognition using attention-based recurrent nn: First results. arXiv preprint arXiv:1412.1602, 2014.
- [5] Yann N Dauphin, Angela Fan, Michael Auli, and David Grangier. Language modeling with gated convolutional networks. In *Proceedings of the 34th International Conference* on Machine Learning-Volume 70, pages 933–941. JMLR. org, 2017.
- [6] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2625–2634, 2015.
- [7] Tharindu Fernando, Simon Denman, Sridha Sridharan, and Clinton Fookes. Soft+ hardwired attention: An lstm framework for human trajectory prediction and abnormal event detection. *Neural networks*, 108:466–478, 2018.
- [8] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. Convolutional sequence to sequence learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1243–1252. JMLR. org, 2017.
- [9] Liyu Gong and Qiang Cheng. Exploiting edge features for graph neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9211–9219, 2019.
- [10] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. http://www. deeplearningbook.org.
- [11] Alex Graves and Navdeep Jaitly. Towards end-to-end speech recognition with recurrent neural networks. In *International conference on machine learning*, pages 1764–1772, 2014.
- [12] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social GAN: socially acceptable trajectories with generative adversarial networks. In 2018 IEEE Conference on Computer Vision and Pattern Recognition,

CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018, pages 2255–2264, 2018.

- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceed-ings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [14] Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. *Physical review E*, 51(5):4282, 1995.
- [15] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [16] Lei Hou, Jian-Guo Liu, Xue Pan, and Bing-Hong Wang. A social force evacuation model with the leadership effect. *Physica A: Statistical Mechanics and its Applications*, 400:93–99, 2014.
- [17] Yingfan Huang, HuiKun Bi, Zhaoxin Li, Tianlu Mao, and Zhaoqi Wang. Stgat: Modeling spatial-temporal interactions for human trajectory prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6272– 6281, 2019.
- [18] Mostafa S Ibrahim, Srikanth Muralidharan, Zhiwei Deng, Arash Vahdat, and Greg Mori. A hierarchical deep temporal model for group activity recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1971–1980, 2016.
- [19] Boris Ivanovic and Marco Pavone. The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2375–2384, 2019.
- [20] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.
- [21] Vineet Kosaraju, Amir Sadeghian, Roberto Martín-Martín, Ian Reid, Hamid Rezatofighi, and Silvio Savarese. Socialbigat: Multimodal trajectory forecasting using bicycle-gan and graph attention networks. In Advances in Neural Information Processing Systems, pages 137–146, 2019.
- [22] Namhoon Lee, Wongun Choi, Paul Vernaza, Christopher B Choy, Philip HS Torr, and Manmohan Chandraker. Desire: Distant future prediction in dynamic scenes with interacting agents. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 336–345, 2017.
- [23] Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. Crowds by example. In *Computer graphics forum*, pages 655–664. Wiley Online Library, 2007.
- [24] Junwei Liang, Lu Jiang, Kevin Murphy, Ting Yu, and Alexander Hauptmann. The garden of forking paths: Towards multi-future trajectory prediction. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 10508–10518, 2020.
- [25] Junwei Liang, Lu Jiang, Juan Carlos Niebles, Alexander G Hauptmann, and Li Fei-Fei. Peeking into the future: Predicting future person activities and locations in videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5725–5734, 2019.
- [26] Karttikeya Mangalam, Harshayu Girase, Shreyas Agarwal, Kuan-Hui Lee, Ehsan Adeli, Jitendra Malik, and Adrien

Gaidon. It is not the journey but the destination: Endpoint conditioned trajectory prediction. *arXiv preprint arXiv:2004.02025*, 2020.

- [27] Ramin Mehran, Alexis Oyama, and Mubarak Shah. Abnormal crowd behavior detection using social force model. In 2009 IEEE Conference on Computer Vision and Pattern Recognition, pages 935–942. IEEE, 2009.
- [28] Abduallah Mohamed, Kun Qian, Mohamed Elhoseiny, and Christian Claudel. Social-stgcnn: A social spatio-temporal graph convolutional neural network for human trajectory prediction. arXiv preprint arXiv:2002.11927, 2020.
- [29] Nishant Nikhil and Brendan Tran Morris. Convolutional neural network for trajectory prediction. In *Proceedings* of the European Conference on Computer Vision (ECCV), pages 0–0, 2018.
- [30] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013, volume 28 of JMLR Workshop and Conference Proceedings, pages 1310–1318. JMLR.org, 2013.
- [31] Stefano Pellegrini, Andreas Ess, and Luc Van Gool. Improving data association by joint modeling of pedestrian trajectories and groupings. In *European conference on computer vision*, pages 452–465. Springer, 2010.
- [32] Priscila Saboia and Siome Goldenstein. Crowd simulation: applying mobile grids to the social force model. *The Visual Computer*, 28(10):1039–1048, 2012.
- [33] Amir Sadeghian, Vineet Kosaraju, Ali Sadeghian, Noriaki Hirose, Hamid Rezatofighi, and Silvio Savarese. Sophie: An attentive gan for predicting paths compliant to social and physical constraints. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 1349– 1358, 2019.
- [34] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Multi-agent generative trajectory forecasting with heterogeneous data for control. arXiv preprint arXiv:2001.03093, 2020.
- [35] Chenyang Si, Wentao Chen, Wei Wang, Liang Wang, and Tieniu Tan. An attention enhanced graph convolutional lstm network for skeleton-based action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1227–1236, 2019.
- [36] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *Advances in neural information processing systems*, pages 3483–3491, 2015.
- [37] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.
- [38] Luca Anthony Thiede and Pratik Prabhanjan Brahma. Analyzing the variety loss in the context of probabilistic trajectory prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9954–9963, 2019.
- [39] Adrien Treuille, Seth Cooper, and Zoran Popović. Continuum crowds. ACM Transactions on Graphics (TOG), 25(3):1160–1168, 2006.

- [40] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W. Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. In *The 9th ISCA Speech Synthesis Workshop, Sunnyvale, CA, USA, 13-15 September* 2016, page 125. ISCA, 2016.
- [41] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph Attention Networks. *International Conference on Learning Representations*, 2018.
- [42] Anirudh Vemula, Katharina Muelling, and Jean Oh. Social attention: Modeling attention in human crowds. In 2018 IEEE International Conference on Robotics and Automation (ICRA), pages 1–7. IEEE, 2018.
- [43] Allan Wang, Zirui Wang, and Wentao Yuan. Pedestrian trajectory prediction with graph neural networks. *Semantic Scholar*, 2019.
- [44] Jack M Wang, David J Fleet, and Aaron Hertzmann. Gaussian process dynamical models for human motion. *IEEE transactions on pattern analysis and machine intelligence*, 30(2):283–298, 2007.
- [45] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, and Chengqi Zhang. Graph wavenet for deep spatial-temporal graph modeling. In Sarit Kraus, editor, *Proceedings of* the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019, pages 1907–1913. ijcai.org, 2019.
- [46] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In Proceedings of the European Conference on Computer Vision (ECCV), pages 305–321, 2018.
- [47] Kota Yamaguchi, Alexander C Berg, Luis E Ortiz, and Tamara L Berg. Who are you with and where are you going? In *CVPR 2011*, pages 1345–1352. IEEE, 2011.
- [48] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [49] Huaxiu Yao, Xianfeng Tang, Hua Wei, Guanjie Zheng, Yanwei Yu, and Zhenhui Li. Modeling spatial-temporal dynamics for traffic prediction. *arXiv preprint arXiv:1803.01254*, 2018.
- [50] Huaxiu Yao, Fei Wu, Jintao Ke, Xianfeng Tang, Yitian Jia, Siyu Lu, Pinghua Gong, Jieping Ye, and Zhenhui Li. Deep multi-view spatial-temporal network for taxi demand prediction. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [51] Pu Zhang, Wanli Ouyang, Pengfei Zhang, Jianru Xue, and Nanning Zheng. Sr-lstm: State refinement for lstm towards pedestrian trajectory prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12085–12094, 2019.