

S2TNet: Spatio-Temporal Transformer Networks for Trajectory Prediction in Autonomous Driving

Wei Huang Chen

Fangfang Wang

Hongbin Sun (corresponding author)

Xi'an Jiaotong University, Xi'an, China

CHENWEIHUANG@STU.XJTU.EDU.CN

FFWANG123@MAIL.XJTU.EDU.CN

HSUN@MAIL.XJTU.EDU.CN

Editors: Vineeth N Balasubramanian and Ivor Tsang

Abstract

To safely and rationally participate in dense and heterogeneous traffic, autonomous vehicles require to sufficiently analyze the motion patterns of surrounding traffic-agents and accurately predict their future trajectories. This is challenging because the trajectories of traffic-agents are not only influenced by the traffic-agents themselves but also by spatial interaction with each other. Previous methods usually rely on the sequential step-by-step processing of Long Short-Term Memory networks (LSTMs) and merely extract the interactions between spatial neighbors for single type traffic-agents. We propose the Spatio-Temporal Transformer Networks (S2TNet), which models the spatio-temporal interactions by spatio-temporal Transformer and deals with the temporal sequences by temporal Transformer. We input additional category, shape and heading information into our networks to handle the heterogeneity of traffic-agents. The proposed methods outperforms state-of-the-art methods on ApolloScape Trajectory dataset by more than 7% on both the weighted sum of Average and Final Displacement Error.

Keywords: Trajectory prediction, Transformer, Autonomous Driving.

1. Introduction

Autonomous driving is an innovative and advanced research field that can reduce the number of road fatalities, increase traffic efficiency, decrease environmental pollution and give mobility to handicapped members of our society (Milakis et al. (2017)). In order to achieve desired goals and avoid collisions of other agents, autonomous vehicles need to have the ability to perceive the environment and make intelligent decisions. As a part of perception, trajectory prediction can well reflect the future behaviors of surrounding agents and build a bridge between perception and decision-making. However, complex temporal prediction is inevitably accompanied by spatial agent-agent interactions at the same time, especially in the dense and highly dynamic traffic composed of heterogeneous *traffic-agents*, including pedestrians, cyclists, human drivers. The heterogeneity means that these traffic-agents have diverse shapes, sizes, dynamics and behaviors. Moreover, a variety of potentially reasonable spatial interactions between traffic-agents may occur, e.g. human drivers may overtake another vehicle or slow down to follow other vehicles (Lefèvre et al. (2014)). Consequently, trajectory prediction is a challenging task that plays an important role in autonomous driving.

Classical methods treat traffic-agents as individual entities without any spatial interactions and abstract their motion as kinematic and dynamic models (Brännström et al. (2010)), Gaussian Processes (Rasmussen (2003)) and etc., making it difficult to comprehend complex scenarios or accomplish long-term predictions. With the success in deep neural networks, recent trajectory prediction methods mainly focus on using these networks to extract features on spatial and temporal dimensions (Alahi et al. (2016); Huang et al. (2019); Ivanovic and Pavone (2019); Mohamed et al. (2020);). Long Short-Term Memory networks (LSTMs) are widely used for modeling temporal features. The LSTMs are based on consecutively processing sequences and storing the latent states to represent knowledge about the motion of traffic-agents (Giuliari et al. (2021)). However, LSTM-based methods remember the history with a single vector with limited memory and regularly have difficulty in handling complex temporal dependencies (Vaswani et al. (2017)). After that, pooling mechanism (Deo and Trivedi (2018)), attention mechanisms (Ivanovic and Pavone (2019)) and graph convolution mechanisms (Li et al. (2019); Yu et al. (2020)) are used to model the spatial interactions. The limitation of these methods is that they only model the interactions of spatially proximal traffic-agents and ignore the influence by traffic-agents beyond the given spatial limits. This assumption may work well when the speed of traffic-agents is low, but lose efficacy with speed increasing. Besides, the majority of trajectory prediction algorithms are developed for homogeneous traffic-agents in a single scene, which corresponding to human pedestrians in crowds (Alahi et al. (2016)) or moving vehicles on a highway (Deo and Trivedi (2018)). These methods may have great limitation on dealing with dense urban environments where heterogeneous traffic-agents coexist and interact with each other.

In this paper, we address all these limitations by employing Spatio-Temporal Transformer Networks (S2TNet) for heterogeneous traffic-agents trajectory prediction. S2TNet is proposed based on the vanilla Transformer architecture, which discards the sequential nature of data and models features with only the effective self-attention mechanism. For the spatial dimension, we propose spatial self-attention mechanism to capture the interactions between all traffic-agents in the road network, not limited to the interactions between spatial neighbors. For the temporal dimension, temporal convolution network (TCN) is adopted to extract temporal dependencies of consecutive frame and combined with spatial self-attention to form the spatio-temporal Transformer where a set of new spatio-temporal features are obtained. Based on temporal self-attention mechanism, temporal Transformer could refine the temporal features for each traffic-agent independently and produce the future trajectories auto-regressively. In addition to history trajectories, we input additional shape, heading, category features into our networks to handle the heterogeneity of traffic-agents. Main contributions of this paper are summarized as follows:

- We put forward an innovative approach for heterogeneous traffic-agents trajectory prediction, employing Transformer-based networks to accurately extract interaction information both on the spatial and temporal dimensions.
- Spatio-temporal Transformer is designed creatively to merge spatial and temporal information from the history features of traffic-agents. After that, temporal Transformer is utilized to enhance capturing temporal dependencies and output future trajectories with specified length.

- S2TNet outperforms prior methods on ApolloScape Trajectory dataset by 7.2% on the weighted sum of Average Displacement Error (WSFDE) and 7.7% on the weighted sum of Final Displacement Error (WSFDE)

2. Background

2.1. Problem Formulation

Trajectory prediction aims to accurately predict the future long-term trajectories of traffic-agents, given their history trajectories and other information such as shapes and categories.

The input of S2TNet is

$$\mathbf{X} = [\mathbf{x}^1, \dots, \mathbf{x}^{t_{obs}}] \quad (1)$$

where,

$$\mathbf{x}^i = \{(x_0^i, y_0^i, l_0^i, w_0^i, \theta_0^i, \tau_0^i, \dots, x_n^i, y_n^i, l_n^i, w_n^i, \theta_n^i, \tau_n^i) \mid i \in (1 : t_{obs})\} \quad (2)$$

are the history feature vectors (including global coordinates x and y , lengths l , widths w , headings θ and categories τ) of n traffic-agents being predicted in a road network. The subscript n in (2) refers to all agents in general and varies with different scenes. We currently take into account five types of traffic-agents $c \in (1, 2, 3, 4, 5)$, representing small vehicles, big vehicles, pedestrian, cyclist and others sequentially. We hold that additional features if available to each traffic-agent could handle the heterogeneity of traffic-agents and improve trajectory accuracy.

The output of S2TNet is

$$\mathbf{Y} = [\mathbf{y}^{t_{obs}}, \dots, \mathbf{y}^{t_{fut}}] \quad (3)$$

where,

$$\mathbf{y}^i = \{(x_0^i, y_0^i, \dots, x_n^i, y_n^i) \mid i \in (t_{obs}+1 : t_{fut})\} \quad (4)$$

are the future feature vectors including global coordinates x and y . It is noted that S2TNet outputs future positions of all observed traffic-agents simultaneously other than merely predicting the location of one specific traffic-agent.

With the objective to hierarchically represent the trajectory sequences, we construct a spatio-temporal graph $G = (V, E)$ on a trajectory sequence with N traffic-agents and T frames featuring both intra-frame and inter-frame connection. In this graph, the node set $V = \{\mathbf{x}_i^t \mid t \in (1, T), i \in (1, N)\}$ includes all the feature vectors of traffic-agents, and E represents the set of edges connected between nodes. We utilize node and traffic-agent equally in the following description. The edge set E consists of two subsets. The first subset depicts the virtual spatial connection between traffic-agents in the same frame, denoted as $E_S = \{(\mathbf{x}_i^t, \mathbf{x}_j^t) \mid i, j \in (1, N), t \in (1, T)\}$. The second subset contains the temporal edges which connects the same traffic-agent in consecutive frames as $E_T = \{(\mathbf{x}_i^t, \mathbf{x}_i^{t_1}) \mid i \in (1, N), t, t_1 \in (1, T)\}$.

2.2. Trajectory prediction networks overview

Trajectory prediction using RNNs Recurrent Neural Networks (RNNs) and their variant structures such as LSTMs and Gated Recurrent Units (GRU) have made great progress in trajectory prediction. As one of the earliest RNNs using in trajectory prediction,

Social-LSTM (Alahi et al. (2016)) addresses interaction between neighborhood by defining a spatial grid based pooling scheme to aggregate the recurrent outputs of all the agents around the agent being predicted. However, this hand-crafted solution is inefficient and fails to capture global context since cells in one grid are treated equally. CS-LSTM (Gupta et al. (2018)) combines a novel pooling mechanism with generative adversarial networks to tackle intrinsically multimodality of pedestrian trajectory. TrafficPredict (Ma et al. (2019)) utilizes LSTM to refine the similarities of motion pattern of instances into category features for heterogeneous agent prediction. SR-LSTM (Zhang et al. (2019)) introduces a message passing and selecting mechanism to capture the crucial current intention of the neighbors.

Trajectory prediction using hybrid networks Recently, the approaches of trajectory prediction have been extended to hybrid networks by combining RNNs with Convolution Networks (CNNs), Generative Adversarial Network (GAN) or Graph Neural Networks (GNNs). Taphic (Chandra et al. (2019)) introduces CNNs into pooling mechanism for maneuver based trajectory prediction. Sophie (Sadeghian et al. (2019)) concatenates the outputs of social and physical attention with the scene features extracted by a CNNs and takes advantage of GAN to generate more realistic samples for the path prediction of multiple interacting agents. Social-BiGAT (Kosaraju et al. (2019)) captures the social interactions information between pedestrians on the basis of graph attention networks. GRIP (Li et al. (2019)) directly models traffic-agents history trajectories as a spatio-temporal graph and forecasts future trajectories based on Spatio-Temporal Graph Convolutional Networks (ST-GCNs).

Trajectory prediction using Transformers On account of Transformer’s (Vaswani et al. (2017)) unique attention mechanism and superior performance in NLP, there is an emerging interest in applying Transformer architectures to prediction tasks. Without considering any complicated interaction information, (Giuliari et al. (2021)) utilizes vanilla Transformer for pedestrian trajectory forecasting and achieves plausible results. STAR (Yu et al. (2020)) interleaves a variant of the graph convolution, named TGConv, and the original Temporal Transformer for spatio-temporal interactions modeling. Inspired by the parallel version of transformer used in (Carion et al. (2020)), mmTransformer (Liu et al. (2021)) uses stacked transformers to aggregate multiple information and achieves multi-modal prediction.

2.3. Self-attention in Transformer

The core of Transformer networks is their unique self-attention mechanism which is used in parallel, while LSTMs only serially combine current word with the embedding of previous words which have been processed. The first step in calculating the self-attention of Transformer is to learn three vectors separately, i.e. query $q_i \in \mathbb{R}^{d_q}$, key $k_i \in \mathbb{R}^{d_k}$ and value $v_i \in \mathbb{R}^{d_v}$ through trainable linear mapping from each embedding $w_i \in \mathbb{R}^e, i \in (1, n)$, where n is the number of words being considered. After that, a score is calculated by the dot product of a query and key, $\alpha_{ij} = q_i \cdot k_j^T, \forall i, j \in (1, n)$, where superscript T is the transpose of vector. Through the softmax function, all scores belonging to the same node are normalized. Finally, the i th self-attention is obtained by multiplying each v_j by normalized scores α_{ij} and summing the weighted results.

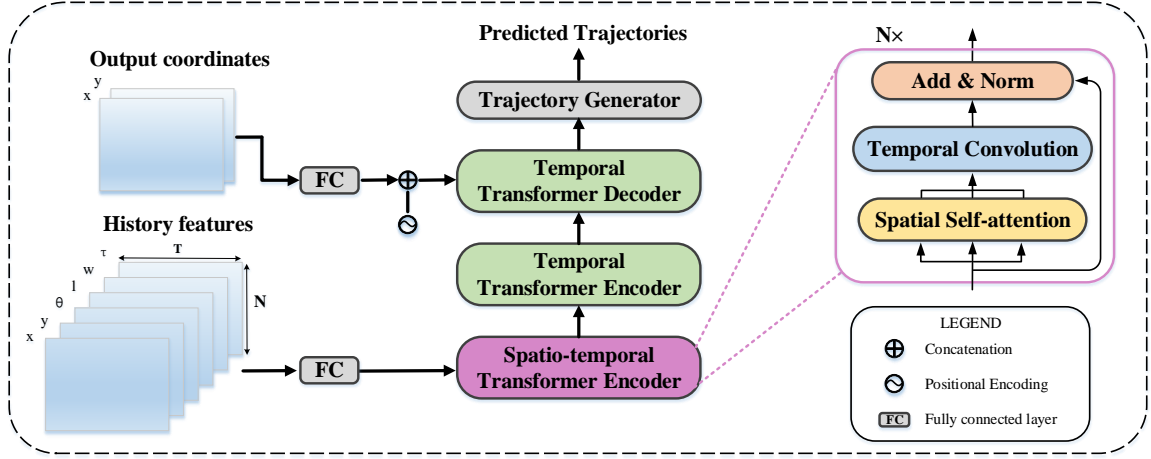


Figure 1: Overview of S2TNet. S2TNet leverages the encoders representation of history features, i.e. x and y coordinates, length l , width w , heading θ and category τ , of all N traffic-agents in T frames, and the decoder to obtain the refined output spatio-temporal features, and further generates future trajectories by passing them to trajectory generator. The two encoders and one decoder contains a stack of $N = 6$ identical layers respectively. The detailed Temporal Transformer can be found in appendix A.

In practice, the attention function is computed on a set of words simultaneously. Three vectors (q, k, v) of all words are individually packed together into three matrices (Q, K, V) . The output of this process, named as *scaled dot-product attention*, can be written as:

$$Attention(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = softmax\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}}\right)\mathbf{V} \quad (5)$$

where d_k is the dimension of each query. The division by $\sqrt{d_k}$ is used to increase gradients stability.

By adding *multi-head attention mechanism*, we can further improve the performance of self-attention. It gives multiple representation sub-spaces for self-attention and enables the model to jointly deal with information from varied sub-spaces at separate positions.

3. Proposed S2TNet Model

3.1. Overview of S2TNet Model

As illustrated in Fig. 1, the whole model can be viewed as an encoder-decoder architecture in which spatio-temporal Transformer encoder, temporal Transformer encoder and temporal Transformer decoder are aggregated hierarchically. For the sake of acquiring abundant motion information, the history feature vectors of each traffic-agent are embedded onto a higher dimensional space by means of a fully connected layer. Then, the spatial interactions in intra-frame are captured by spatial self-attention and the temporal features of inter-frame are obtained by TCN. Our model emphasizes the coupled spatio-temporal modeling by interleaving the spatial self-attention and TCN in a single spatio-temporal Transformer

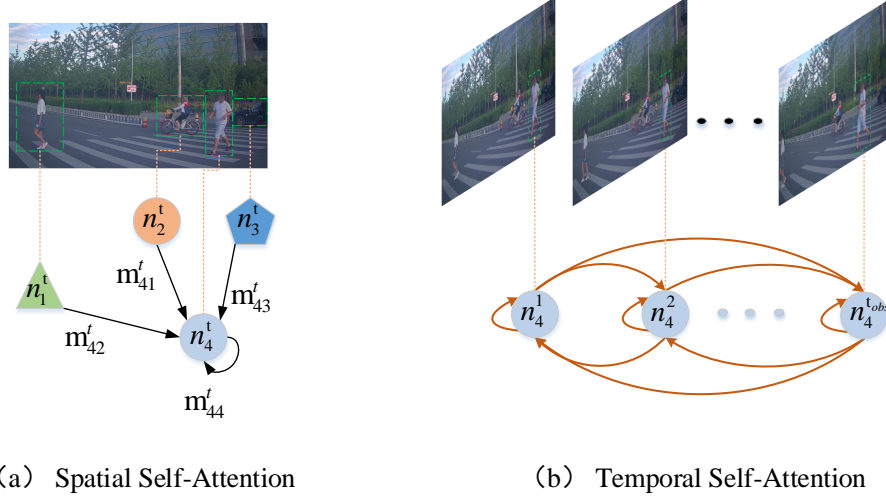


Figure 2: Spatial and Temporal Self-Attention. (a) The spatial interactions of node 4 in frame t is modeled. n_i^t ($i = 1, 2, 3, 4$) is the embeddings of node i . m_{4j}^t ($j = 1, 2, 3, 4$) is the message passing from node j to 4. (b) The temporal correlations between inter-frame are computed in temporal Transformer where the nodes are independent of each other.

layer. In order to further capture the temporal dependencies on all history frames, we perform post-processing of the input embeddings with the second temporal Transformer encoder. Temporal Transformer decoder refines the output embeddings based on the spatio-temporal features provided by encoders and the previously predicted output embeddings produced by previously output coordinates. Finally, the trajectory generator outputs all the traffic-agents future trajectories $\mathbf{Y}^{(t_{obs}+1, t_{fut})}$ simultaneously by decoding the output embeddings.

3.2. Spatio-temporal Transformer

In order to handle the spatial interactions coupled with temporal continuity, we creatively design a spatio-temporal Transformer encoder that captures spatial information through a spatial self-attention sub-layer and extracts dependencies along the temporal dimension through a TCN sub-layer. We interleave two sub-layers to merge the spatio-temporal features.

Spatial Self-attention sub-layer From a different perspective of Transformer, the spatial attention could be regard as spatial-edge in the spatio-temporal graph. We adopt message passing mechanism on the spatial-edge to preform the suitable processing. For each node i in the scene at time t , query $q_i^t \in \mathbb{R}^{d_q}$, key $k_i^t \in \mathbb{R}^{d_k}$ and value $v_i^t \in \mathbb{R}^{d_v}$ is computed by the linear projection from input embeddings $h_i^t \in \mathbb{R}^C$:

$$q_i^t = W_q \cdot h_i^t, k_i^t = W_k \cdot h_i^t, v_i^t = W_v \cdot h_i^t \quad (6)$$

where $W_q \in \mathbb{R}^{C \times d_q}, W_k \in \mathbb{R}^{C \times d_k}, W_v \in \mathbb{R}^{C \times d_v}$. Attention score between node i and $j \in (1, V)$ is then obtained by applying scaled dot-product between q_i^t and k_j^t , representing the

spatial-edge message m_{ij}^t send from j to i , as depicted in Fig. 2(a).

$$m_{ij}^t = q_i^t \cdot k_j^{tT} \quad (7)$$

The messages sent from all j to i is normalized over the weights of spatial-edges and summed to get a single attention head of node i , as in the following:

$$head_i^t = \sum_j softmax\left(\frac{m_{ij}^t}{\sqrt{d_k}}\right)v_j \quad (8)$$

By repeating this embedding extraction process h times, multi-head attention are concatenated and projected to output embeddings with an fully connected layer:

$$MultiHead_i^t = \mathbf{W}_o \cdot concat([head_{i_0}^t, \dots, head_{i_h}^t]) \quad (9)$$

Temporal Convolution sub-layer After spatial information is obtained, we impose temporal convolution operation on the temporal-edge in the spatio-temporal graph to model temporal dynamics within trajectory sequence. Given the input graph of shape (T, N, C) , where T is history frame, N is node number and C is the embeddings, we use a standard 2D convolution with the kernel size $(K \times 1)$ to force on processing along the temporal dimension, as expressed in the following:

$$output_i = Conv_{K \times 1}(MultiHead_i^t) \quad (10)$$

As a Transformer structure, we regularly imply layer normalization (Ba et al. (2016)) after skip connection in the end of TCN sub-layer. That is, the output of sub-layer is $LayerNorm(x + Sublayer(x))$. In this way, we have a well-defined operation on the constructed spatio-temporal graph.

3.3. Temporal Transformer

Temporal Transformer consists of an encoder and decoder. The capability of temporal Transformer encoder is performed to better study the dynamics of each node independently along the temporal dimension. The temporal Transformer decoder is used to refine the output embeddings by encoder outputs and the previously predicted embeddings.

Encoder Temporal Transformer encoder layer is composed of temporal self-attention sub-layer and separable convolution sub-layer. Each encoder layer has two sub-layers. The first sub-layer, called temporal self-attention, uses a multi-head self-attention mechanism similar to spatial self-attention sub-layer in Spatial Transformer with the difference that correlations along the temporal dimension are computed independently for each node. As shown in Fig. 2(b), the temporal self-attention for node i represented as:

$$MultiHead_i = \mathbf{W}_u \cdot concat([head_{i_0}, \dots, head_{i_h}]) \quad (11)$$

$$where, head_i = softmax\left(\frac{Q_i \cdot K_i^T}{\sqrt{d_k}}\right)V_i \quad (12)$$

Where Q^i , K^i and V^i are query, key and value matrix learned from the embeddings of input node i .

Instead of fully connected network used in vanilla Transformer, the second sub-layer is the separable convolution (Chollet (2016)) in order to achieve higher accuracy.

Decoder To inject the relative position information of previous output trajectories to decoder, we add the positional encodings to output embeddings:

$$PE_{(pos,2i)} = \sin(pos/10000^{2i/d_{model}}) \quad (13)$$

$$PE_{(pos,2i)} = \cos(pos/10000^{2i/d_{model}}) \quad (14)$$

where pos is the position, i is the dimension and d_{model} the total dimensions of the output embeddings.

Compared with temporal self-attention in encoder, decoder employs a masked temporal self-attention sub-layer to ensure that the predictions for time t can only depend on the known outputs at times less than t . Besides masked temporal self-attention and separable convolution, a third sub-layer is inserted into decoder layer which performs multi-head attention over the output of the temporal Transformer encoder.

3.4. Implementation Details

The scheme is implemented using PyTorch. The dimensions of embedding features is set to 32. We apply dropout to the output of each sub-layer before the skip connection step and the output of positional encodings in the decoder stacks. The dropout ratio is $P_{drop} = 0.1$.

An L2-loss is adopted

$$Loss = \sum_{t=t_{obs}+1}^T |Y_{pred}^t - Y_{GT}^t|^2 \quad (15)$$

where Y_{pred}^t and Y_{GT}^t are predicted positions and ground truth positions respectively. We use Adam Kingma and Ba (2014) as the optimizer and impose a learning rate variation strategy as follows:

$$learning_rate = d_{model}^{-0.5} \cdot \min(step_num^{-0.5}, step_num \cdot warmup_steps^{-1.5}) \quad (16)$$

where $warmup_step$ is set to 5000. Random rotation is implemented for data augmentation in the training.

4. Experiments

4.1. Dataset and Evaluation Metrics

Our model is evaluated on ApolloScape Trajectory dataset (Ma et al. (2019)) which is collected by Apollo autonomous vehicles. The ApolloScape Trajectory dataset contains images, point clouds, and manually annotated trajectories. It is gathered under various lighting conditions and traffic densities in Beijing, China. More specifically, it comprises vastly complex traffic flows mixed with vehicles, riders, and pedestrians. The dataset includes 53 minute training sequences and 50 minute S2TNet sequences captured at 2 frames per second. We need to predict six future frames based on six history frames. Due to the

testset of ApolloScape Trajectory dataset is not public, we obtain the results of our model and other baselines by uploading to the ApolloScape Trajectory Leaderboard ¹.

Two metrics are used to evaluate model performance: the Average Displacement Error (ADE) (Pellegrini et al. (2009)) and the Final Displacement Error (FDE). ADE is the mean Euclidean distance over all predicted positions and ground truth positions during the prediction time, and FDE is the last item of ADE. Obviously, ADE shows the average prediction performance, while the FDE reflects just the prediction accuracy at the end points. Because the trajectories of heterogeneous traffic-agents are diverse in scales, we use the following weighted sum of ADE (WSADE) and weighted sum of FDE (WSFDE) as metrics:

$$WSADE = D_v \cdot ADE_v + D_p \cdot ADE_p + D_b \cdot ADE_b \quad (17)$$

$$WSFDE = D_v \cdot FDE_v + D_p \cdot FDE_p + D_b \cdot FDE_b \quad (18)$$

where $D_v = 0.20$, $D_p = 0.58$, and $D_b = 0.22$ are relevant with reciprocals of the average velocity of vehicles, pedestrian and cyclist in the dataset.

4.2. Baselines

To evaluate the performance of S2TNet, we compare S2TNet with a wide range of baselines, including:

- *Constant Velocity (CV)*: We use the average velocity of history trajectories as the constant velocity during the future to predict trajectories.
- *TrafficPredict*: A LSTM-based method using a hierarchical architecture by (Ma et al. (2019)).
- *StarNet*: (Zhu et al. (2019)) builds a star topology to consider the collective influence among all pedestrians.
- *Social LSTM (S-LSTM)*: (Alahi et al. (2016)) uses LSTM to extract single pedestrian feature and devises a social pooling mechanism to capture neighbor information.
- *Social GAN (S-GAN)*: (Gupta et al. (2018)) predicts socially plausible futures by a conditional GAN.
- *Transformer*: (Giuliani et al. (2021)) uses vanilla temporal Transformer to model pedestrian separately without any complex human-human interactions nor scene interaction terms.
- *STAR*: (Yu et al. (2020)) interleaves spatial and temporal Transformer to capture the social intersection between pedestrians.
- *TPNet*: (Fang et al. (2020)) first generates a candidate set of future trajectories, then gets the final predictions by classifying and refining the candidates.
- *GRIP++*: (Li et al. (2019)) is the SOTA trajectory predictor which uses an enhanced graph to represent the interactions of close objects, and applies ST-GCNS to extract spatio-temporal features.

1. http://apolloscape.auto/leader_board.html

4.3. Quantitative Results and Analyses

We compare S2TNet with the state-of-the-art approaches as mentioned in Section 4.2. All methods are compared by the results released on ApolloScape Trajectory Leaderboard. The main results are presented in Table 1.

From Table 1 we observe that the performance of S2TNet is superior to the baseline methods of all traffic-agent types by a large margin. More specifically, our method reduces the ADE of vehicles, pedestrians, and cyclists over GRIP++ by 11.28%, 4.31% and 10.24% respectively. Meanwhile, our method reduces the FDE of vehicles, pedestrians, and cyclists over GRIP++ by 12.21%, 4.98% and 5.87% sequentially. It is notice worthy that the degree of improvement in vehicles and cyclists is better than pedestrians. We believe it is because that the motion pattern of pedestrians are more flexible than vehicles and bikes with non-holonomic constraint. Another remarkable finding is that simple model CV which just makes use of average velocity of history trajectories outperforms many deep learning methods including the SOTA model, STAR. This suggests that homogeneous methods may not handle dense urban scenes efficiently. On the contrary, our approach performs better in heterogeneous and dense urban environments. We will further demonstrate this in Section 4.4 with visualized results.

Table 1: Comparison with baselines models on ApolloScape Trajectory dataset.

Method	WSADE	ADE _v	ADE _p	ADE _b	WSFDE	FDE _v	FDE _p	FDE _b
TrafficPredict	8.5881	7.9467	7.1811	12.8805	24.2262	12.7757	11.121	22.7912
S-LSTM	1.8922	2.9456	1.2856	2.5337	3.4024	5.2802	2.3240	4.5384
S-GAN	1.5829	3.0430	0.9836	1.8354	2.7796	5.0913	1.7264	3.4547
STAR	1.5400	2.5644	0.9473	2.1714	2.8602	4.6324	1.8029	4.0366
CV	1.4762	2.6454	0.8547	2.0519	2.7601	4.7944	1.6428	3.8564
StraNet	1.3425	2.3860	0.7854	1.8628	2.4984	4.2857	1.5156	3.4645
Transformer	1.2803	2.2322	0.7398	1.8398	2.4024	4.0317	1.4309	3.4826
TPNet	1.2800	2.2100	0.7400	1.8500	2.3400	3.8600	1.4100	3.4000
GRIP++	1.2588	2.2400	0.7142	1.8024	2.3631	4.0762	1.3732	3.4155
S2TNet	1.1679	1.9874	0.6834	1.7000	2.1798	3.5783	1.3048	3.2151

4.4. Qualitative Results and Analyses

In Fig. 3, we visualize several prediction results of ApolloScape Trajectory dataset. We separately present the trajectory of single traffic-agent with different type selected from complex scenes and show the complete prediction results of a scene.

- *S2TNet has the ability to forecast long horizon trajectories for different categories of traffic-agents.* After observing 6 frames (3s) of history trajectories, S2TNet could accurately predict the trajectories over 3 seconds horizon. Moreover, S2TNet does well in the case of sharp turns for the vehicle, e.g. Fig. 3(a) and (b). With the increase of prediction length, the prediction results of S2TNet are more realistic and the cumulative error of S2TNet is better than GRIP++, e.g. Fig. 3(c) and (d).

- *S2TNet is able to model spatio-temporal interaction accurately.* In the top right portion of Fig. 3(e) and (f), a vehicle runs in opposite directions to an unknown traffic-agent. While the predicted trajectories of GRIP++ deviates from ground truth, S2TNet precisely captures the interactive routes.
- *S2TNet successfully identify the stationary traffic-agent.* In the lower-left of Fig. 3(e) and (f), two vehicles decelerate to near standstill. Compared with GRIP++, S2TNet successfully predicts the corresponding stationary trajectories.

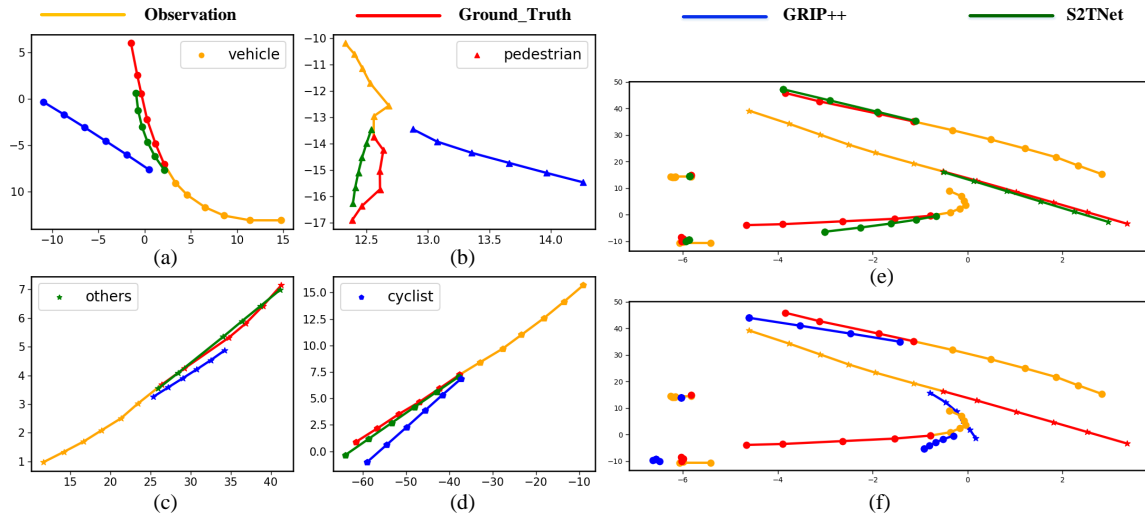


Figure 3: Visualized Prediction Results in heterogeneous and dense traffic. S2TNet successfully captures spatio-temporal information and outperforms the SOTA model, GRIP++. (a, b, c, d) Comparison the future trajectories of different types of traffic-agents between two methods. (e, f) The prediction results of GRIP++ and S2TNet in a complete traffic scene.

4.5. Ablation Studies

In this section, we conduct extensive ablation studies and focus on the effect of the proposed components. The results are presented in Table 2.

- *The spatio-temporal Transformer could sufficiently extract information both in spatial and temporal dimensions.* In (1), (2) and (3), we remove one or two sub-layers in spatio-temporal Transformer. Compared (1) to (2), the model contains TCN sub-layer outperforms solely temporal Transformer. On the contrast to outperforming in our validation set, (3) which contains the spatial self-attention sub-layer and temporal Transformer is worse than (1) in final test set. We hold that merely stacking attention on the spatial dimension without merging temporal information results in overfitting.
- *The temporal Transformer encoder enhance capturing temporal dependencies.* In (4), we remove the temporal Transformer encoder and gain a lower performance compared

with (8). This indicates that temporal self-attention mechanism could effectively improve the ability to extract temporal information.

- *The separable convolution outperforms full connected feed-forward network in temporal Transformer.* In (5), we replace separable convolution sub-layer in temporal Transformer with full connected feed-forward network and the performance slightly descends.
- *More features, higher accuracy.* Instead of feeding all features into S2TNet, we input only history trajectories in (6). We find that rich information helps the network to understand the heterogeneity of traffic-agents.
- *The spatial self attention of the whole scene is better than that of the given spatial limits* We use a masked attention mechanism in (7) to ignore the influence out of the given spatial limits (15m) as (Li et al. (2019)) does. We find that the traffic-agents in the whole scene have a great influence on the accuracy of trajectory prediction.

Table 2: Ablation study. **SS** denotes spatial self-attention sub-layer in spatio-temporal Transformer. **TE** denotes temporal Transformer encoder layer. **SC** denotes separable convolution. **FC** denotes full connected layer. **TD** denotes temporal Transformer decoder layer. **HF** denotes history features. **A** denotes history features including global coordinates, category, length, width and heading. **C** denotes only global coordinates. **LM** denotes spatial limits used in spatial self-attention sub-layer. **W** denotes the spatial self-attention without spatial limits. **N** denotes the spatial self-attention of neighbors (15m).

	Components						Performance
	SS	TCN	TE	TD	HF	LM	(WSADE/WSFDE)
(1)	×	×	SC	SC	A	W	1.2300/2.2949
(2)	×	✓	SC	SC	A	W	1.2189/2.2570
(3)	✓	×	SC	SC	A	W	1.2500/2.3561
(4)	✓	✓	×	SC	A	W	1.2674/2.4086
(5)	✓	✓	FC	FC	A	W	1.1945/2.2613
(6)	✓	✓	SC	SC	C	W	1.2170/2.3036
(7)	✓	✓	SC	SC	A	N	1.2686/2.3548
(8)	✓	✓	SC	SC	A	W	1.1679/2.1798

5. Conclusion

In this paper, we propose S2TNet, a Transformer-based framework to predict the trajectories of heterogeneous traffic-agents around autonomous driving cars. Spatio-temporal Transformer is designed to capture spatio-temporal interactions between all traffic-agents, not limited to spatial neighbor. The temporal Transformer is utilized to enhance modeling

temporal dependencies and output future trajectories auto-regressively. The experimental results from ApolloScape Trajectory dataset show that the proposed method achieves the state-of-the-art performance and substantially improves the accuracy of the predicted trajectories. In the future work, we intend to integrate additional map information on S2TNet framework and implement real time prediction on autonomous driving platform by S2TNet.

Acknowledgments

This research is supported by National Natural Science Foundation of China (No. 61790563).

References

- 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.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization, 2016.
- Mattias Brännström, Erik Coelingh, and Jonas Sjöberg. Model-based threat assessment for avoiding arbitrary vehicle collisions. *IEEE Transactions on Intelligent Transportation Systems*, 11(3):658–669, 2010.
- Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229. Springer, 2020.
- Rohan Chandra, Uttaran Bhattacharya, Aniket Bera, and Dinesh Manocha. Ttraffic: Trajectory prediction in dense and heterogeneous traffic using weighted interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8483–8492, 2019.
- François Chollet. Xception: Deep learning with depthwise separable convolutions. *CoRR*, abs/1610.02357, 2016. URL <http://arxiv.org/abs/1610.02357>.
- Nachiket Deo and Mohan M Trivedi. Convolutional social pooling for vehicle trajectory prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1468–1476, 2018.
- Liangji Fang, Qinhong Jiang, Jianping Shi, and Bolei Zhou. TpNet: Trajectory proposal network for motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2020.
- Francesco Giuliari, Irtiza Hasan, Marco Cristani, and Fabio Galasso. Transformer networks for trajectory forecasting. In *2020 25th International Conference on Pattern Recognition*, pages 10335–10342. IEEE, 2021.

- Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social GAN: socially acceptable trajectories with generative adversarial networks. *CoRR*, abs/1803.10892, 2018. URL <http://arxiv.org/abs/1803.10892>.
- 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/CVF International Conference on Computer Vision*, pages 6272–6281, 2019.
- Boris Ivanovic and Marco Pavone. The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2375–2384, 2019.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Vineet Kosaraju, Amir Sadeghian, Roberto Martín-Martín, Ian Reid, S Hamid Rezatofighi, and Silvio Savarese. Social-bigat: Multimodal trajectory forecasting using bicycle-gan and graph attention networks. *arXiv preprint arXiv:1907.03395*, 2019.
- Stéphanie Lefèvre, Dizan Vasquez, and Christian Laugier. A survey on motion prediction and risk assessment for intelligent vehicles. *ROBOMECH journal*, 1(1):1–14, 2014.
- Xin Li, Xiaowen Ying, and Mooi Choo Chuah. GRIP: graph-based interaction-aware trajectory prediction. *CoRR*, abs/1907.07792, 2019. URL <http://arxiv.org/abs/1907.07792>.
- Yicheng Liu, Jinghuai Zhang, Liangji Fang, Qinhong Jiang, and Bolei Zhou. Multimodal motion prediction with stacked transformers. *arXiv preprint arXiv:2103.11624*, 2021.
- Yuexin Ma, Xinge Zhu, Sibozhang, Ruigang Yang, Wenping Wang, and Dinesh Manocha. Trafficpredict: Trajectory prediction for heterogeneous traffic-agents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6120–6127, 2019.
- Dimitris Milakis, Bart Van Arem, and Bert Van Wee. Policy and society related implications of automated driving: A review of literature and directions for future research. *Journal of Intelligent Transportation Systems*, 21(4):324–348, 2017.
- Abduallah Mohamed, Kun Qian, Mohamed Elhoseiny, and Christian Claudel. Social-stgcnn: A social spatio-temporal graph convolutional neural network for human trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14424–14432, 2020.
- S. Pellegrini, A. Ess, K. Schindler, and L. van Gool. You’ll never walk alone: Modeling social behavior for multi-target tracking. In *2009 IEEE 12th International Conference on Computer Vision*, pages 261–268, 2009. doi: 10.1109/ICCV.2009.5459260.
- Carl Edward Rasmussen. Gaussian processes in machine learning. In *Summer school on machine learning*, pages 63–71. Springer, 2003.

- Amir Sadeghian, Vineet Kosaraju, Ali Sadeghian, Noriaki Hirose, Hamid Reza Tofighi, and Silvio Savarese. Sophie: An attentive gan for predicting paths compliant to social and physical constraints. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1349–1358, 2019.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017.
- Cunjun Yu, Xiao Ma, Jiawei Ren, Haiyu Zhao, and Shuai Yi. Spatio-temporal graph transformer networks for pedestrian trajectory prediction. *CoRR*, abs/2005.08514, 2020. URL <https://arxiv.org/abs/2005.08514>.
- 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 12085–12094, 2019.
- Yanliang Zhu, Deheng Qian, Dongchun Ren, and Huaxia Xia. Starnet: Pedestrian trajectory prediction using deep neural network in star topology. *CoRR*, abs/1906.01797, 2019. URL <http://arxiv.org/abs/1906.01797>.

Appendix A. Temporal Transformer encoder and decoder

The detailed temporal Transformer architecture used in S2TNet is visualized in Fig. 4. Input embeddings is passed to the temporal Transformer encoder to enhance capturing the temporal features of observed traffic-agents. Then, the temporal Transformer decoder receives the previously output embeddings and produced the refined output embeddings through masked temporal self-attention, decoder-encoder attention and separable convolution layers.




Figure 4: Temporal Transformer encoder and decoder