

UST: Unifying Spatio-Temporal Context for Trajectory Prediction in Autonomous Driving

Hao He¹, Hengchen Dai¹, Naiyan Wang¹

Abstract—Trajectory prediction has always been a challenging problem for autonomous driving, since it needs to infer the latent intention from the behaviors and interactions from traffic participants. This problem is intrinsically hard, because each participant may behave differently under different environments and interactions. This key is to effectively model the interlaced influence from both spatial context and temporal context. Existing work usually encodes these two types of context separately, which would lead to inferior modeling of the scenarios. In this paper, we first propose a unified approach to treat time and space dimensions equally for modeling spatio-temporal context. The proposed module is simple and easy to implement within several lines of codes. In contrast to existing methods which heavily rely on recurrent neural network for temporal context and hand-crafted structure for spatial context, our method could automatically partition the spatio-temporal space to adapt to the data. Lastly, we test our proposed framework on two recently proposed trajectory prediction dataset ApolloScape and Argoverse. We show that the proposed method substantially outperforms the previous state-of-the-art methods while maintaining its simplicity. These encouraging results further validate the superiority of our approach.

I. INTRODUCTION

With the great development of deep learning techniques in recent years, the perception system equipped in autonomous driving system has been significantly advanced. However, the difficulty of another equally important task, predicting the future trajectories for traffic participants in real-world scenarios is still underestimated, since trajectory prediction task requires understanding the latent intention from the behaviors and interactions of participants.

The key to this challenging task is to model the *complicated spatio-temporal social context* and tolerate the *imperfect output from perception systems*. For the first task, the difficulty lies in that we can only infer the intention from indirect observations. In a certain circumstance, different drivers or pedestrians may make distinct decisions. On the other hand, even for a certain driver or pedestrian, their behaviors may be easily affected by diverse interactions at different places and different time (see Fig. 1(a)). For the second challenge, we cannot expect an oracle perception system due to occlusion and limited range of sensors. Common mistakes include trajectory interruption, unstable speed estimation, etc. Existing works always make the strong assumption that each neighbor has a fix length trajectory and highly depend on accurate speed feature. Some works even eliminate those agents who have incomplete history[1], [2].

¹Hao He, Hengchen Dai and Naiyan Wang are with TuSimple, Beijing, China. {hugohaohe, dhcisyx, winsty}@gmail.com

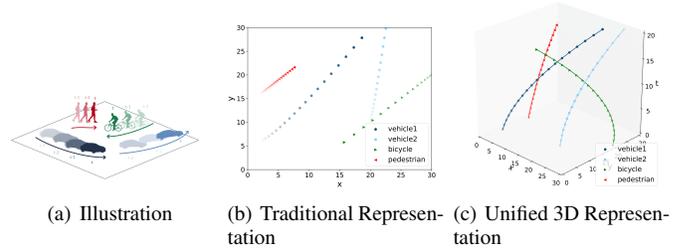


Fig. 1. Illustration and representations of the trajectory prediction task. Blue, green, red colors show trajectories for vehicles, bicycles, pedestrians respectively. (b) shows the common representation, which represents the surrounding agents as sequences of positions in 2D spatial space. (c) shows our proposed trajectory representation in a unified spatio-temporal space.

However, these are just the common scenarios we meet in real autonomous driving systems.

Most learning-based trajectory prediction methods can be categorized into the following encoder-decoder framework: they usually consist of *spatio-temporal context encoder* and *future trajectory decoder* [3], [1]. In the first module, we want to utilize all the available history information to model the intention of agents. The information we can use include the history trajectory of one single agent (temporal context) and the interaction between agents in one single step (spatial context). Other meta-data such as maps or traffic rules can also be incorporated. While in the second module, given the encoded context, we need to generate the future trajectory of the agent. The future trajectory can be represented either by a deterministic path [4], [5], a path with uncertainty estimation [3], [6] or several sampling paths [1], [7].

Most existing spatio-temporal context encoders encode spatial context and temporal context separately. They typically extract each agent’s temporal context with a recurrent neural network (RNN). After that, a spatial context extractor (e.g., convolution [2], attention [8], pooling [1], graph network [9]) will be applied to further aggregate the spatial context. The spatial information between agents is considered at the last time step, thus the interaction among agents in previous time steps are all abandoned. To improve upon these methods, recent works [6], [10] have proposed to leverage spatio-temporal graph. The basic idea of these methods is to extract spatial context at each time step, and then fuse it into individual temporal features for next time aggregation. However, they still extract spatio-temporal context in this cascade style (first temporal, then spatial) essentially. On one hand, extracting spatial context at each time step can be time-consuming (reported in [1] as 16x slower); on the other hand,

it cannot model complex spatial context across different time step. There is also an intrinsic problem involved with RNN-based temporal context extractor. When existing missing data, either eliminating the whole history or padding with interpolations is necessary. All these drawbacks call for a unified spatio-temporal modeling method for trajectory prediction.

In this paper, we propose a novel method to address the above challenges by Unifying the Spatial and Temporal context (UST) into one single representation. The core idea of UST is to jointly represent the spatio-temporal context in a higher dimensional space, which does not distinguish these two concepts explicitly. Then the status of a certain agent at a certain time is represented as a point in this space. (See Fig. 1(c) for illustration.) Other meta-data can also be easily incorporated. As a result, the spatio-temporal context becomes the distribution of these unordered points in the new space. Next, inspired by the classical work PointNet [11], we devise a novel encoder structure for joint context modeling. Based on the extracted context, various off-the-shelf decoders can be used to incorporate different purposes from subsequent modules. To summarize, our main contributions are as follows:

- We first propose to unify 2D location and discrete time space into one single 3D space. We treat them equally.
- Based on this representation, we devise a simple and effective network to encode the spatio-temporal context. The method can be easily implemented within ten lines of codes.
- Experiments on two major trajectory prediction datasets demonstrates the effectiveness of UST.

II. RELATED WORK

In this section, we will give a brief review of the history and recent researches about trajectory prediction problem.

The idea of predicting 2D future trajectories using spatio-temporal context could date back to the work of Helbing and Molnar [12], who proposed a model with attractive and repulsive *Social Force*. It has achieved remarkable successes in robotics[13], [14] and activity understanding[15], [16]. Recently, many researchers extended these interaction-aware trajectory prediction methods to a variety of areas with advanced deep learning techniques[3], [1], [4], [7], including pedestrians trajectory prediction[3], vehicles trajectory prediction[2], intention prediction[17] and heterogeneous prediction[18]. Despite they are dealing with various tasks with different approaches, we summarize most deep learning based methods into the *spatio-temporal context encoder* and *future trajectory decoder* framework.

a) *Spatio-temporal Context Encoder*: The raw data of trajectory prediction task consists of both spatial and temporal information and various metadata. They cannot be directly processed by off-the-shelf neural networks. Popular representation forms include sequential points-based [1], [6], [10], occupy grid-based [19], [2], displacement volume [20], [21] and rasterized image [22], [23]. Directly representing trajectory as a sequence of 2D location is the most natural

way, however its unstructured characteristics make context extraction hard. Although all other representations preserve the structural information, their structures are hand-crafted thus need deliberately tuning, which makes the generalization across different scenarios infeasible. Moreover, these hand-crafted structures are sensitive to the quality of perception output. If the quality of the output of upstream modules changes, the structure may need to be redesigned. Consequently, how to adaptively build the structure from data still remains an open problem.

Based on these representations, researchers have devised various networks to extract the spatio-temporal context. However, all existing work extract the temporal context and spatial context individually, and aggregate them in a cascade way. For the temporal context, Recurrent Neural Network (RNN) is the most widely used. There are also some other algorithms that use Convolution Neural Network (CNN) to encode temporal context[24], [5]. For the spatial context (a.k.a “social context”), its extraction remains to a hot research topic in recent years. Various methods including convolution[2], pooling[3], [1], attention network[6], graph neural network[10], relation network[25] have been used in past few years. Instead of manually building the spatial partition and then encoding the temporal context and spatial context separately, our proposed UST unifies the spatio-temporal space and learn to partition this joint space end-to-end.

b) *Future Trajectory Generation*: The task of the generator or decoder is to generate the future position of the target agent based on both the ego information and the encoded context information. The most common one is to use a recurrent neural network (RNN) to regress the future trajectory directly[10], [4]. There are also other variants based on RNN to accommodate other demands from downstream modules. To name a few, Gupta *et al.*[1] incorporated noise to the RNN decoder and train it with a discriminator and variety loss to generate diverse socially acceptable trajectories. Deo *et al.* [2] tried to classify maneuvers firstly then used the classification result to construct a Gaussian mixture model for multi-modal trajectory prediction. Chai *et al.* [26] classified trajectories to the pre-clustered anchors instead of predefined maneuvers. In this work, the decoder is not our focus. Thus we directly utilize the off-the-shelf decoders in our proposed method.

III. METHOD

This section gives details of our proposed method. Subsection III-A formalizes the trajectory prediction problem mathematically. Subsection III-B details how to encode the spatial-temporal context in a unified framework. Subsection III-C elaborates how we use the encoded spatio-temporal context representation to generate future trajectories.

A. Problem Definition

Let’s assume that there are N agents and T time steps in total, and $\mathbf{x}_n^t \in \mathbb{R}^2$ denotes the raw 2D spatial location of agent n at time step t with respect to a predefined reference

frame. We can further formulate the history of one single agent n as $\mathbf{X}_n \doteq \{\mathbf{H}_n, \mathbf{F}_n\}$, where $\mathbf{H}_n \doteq \{x_n^1, x_n^2, \dots, x_n^t\}$ refers to history locations of agent n starting from time 1 to time t , while $\mathbf{F}_n \doteq \{x_n^{t+1}, x_n^{t+2}, \dots, x_n^T\}$ refers to future locations of agent n starting from time $t+1$ to time T . \mathcal{I} is the additional information such as traffic-agent type, map information, taillight, heading, etc. The goal of trajectory prediction is to utilize all $\{\mathbf{H}_n\}_{n=1}^N$ and \mathcal{I} to accurately predict $\{\mathbf{F}_n\}_{n=1}^N$.

B. Spatio-temporal Context Encoder

In this subsection, we elaborate details of our spatio-temporal context encoder. We first introduce the input representation of our method, then followed by the encoder structure.

1) *Input Representation*: For the sake of simplicity and flexibility, we design spatio-temporal point sets to represent the raw input (sequence of locations of multiple agents) in the following form:

$$\begin{aligned} \mathbf{p}_n^t &= \{x_n^t, v_n^t, \phi_n, t, \text{id}\}, \\ \mathcal{P} &= \{p_n^t \mid \forall t \in \{1, \dots, T\}, \forall n \in \{1, \dots, N\}\} \end{aligned} \quad (1)$$

where x_n^t, v_n^t are 2D locations and 2D velocity of agent n in current reference frame, ϕ_n denotes some intrinsic properties of agent n such as type. $t \in \{1, \dots, T\}$ refers to the time step. And id is a binary number, 0 is for neighbor agents, and 1 for the target agent to be predicted. We don't distinguish different neighborhood agents, and treat them equally.

After applying the above input representation module, we can treat the snapshot of status of agent n at time step t as a single point with metadata in a 3D space spanned by 2D location and time. An intuitive illustration is depicted in Fig. 1(c). Note that all the information we need to model the social interaction of the target agent n is presented in \mathcal{P} . One desired property of this representation is that it is invariant to the order of each \mathbf{p}_n^t in it and robust to missing data. By this uniform representation, we unify space and time into one representation, which eases the subsequent context modeling task.

2) *Unified Spatio-Temporal Context Extraction*: To deal with such unordered and variable length data, the structure and operations of this feature extractor should be deliberately designed to fit the nature of the data. Fortunately, the same challenge has been met in the area of point cloud processing. Inspired by PointNet [11], we propose the following two key components for context extraction.

a) *Embedding*: The goal of this embedding network ϕ is to map p_n^t into a hidden representation e_n^t , in which the spatial context and temporal context are unified:

$$e_n^t = \phi(p_n^t; W_e). \quad (2)$$

In our implementation, ϕ is implemented by a multiple layer perceptron (MLP). W_e is the embedding weight. Batch normalization is applied over all layers with ReLU activation functions. Although there are several works that also apply MLP on time series prediction[27], they treat the whole time series as a fixed-length chronological data instead of a set of permutation invariant snapshots.

b) *Permutation Invariant Aggregator*: After embedding each \mathbf{p}_n^t , we need a permutation invariant aggregator to form the global context feature. [11] has shown that simple pooling operator is capable for the task. It is the simplest symmetric function that enjoys this property. By default, we use max pooling as the aggregator.

Each dimension in the global feature actually corresponds to a partition in the feature space. Each of them represents one configuration of spatio-temporal context. And this partition is learnt end-to-end from the data. To intuitively understand these operation, we visualize several examples in Fig. 2 with NGSIM [28] dataset.

We randomly choose two dimensions from the global context feature, and find the most influencing positions in this 3D space. We further pickup some typical cases and present their activation of this dimension in these figures. For the case in the Fig. 2(a), most points spread in the left front and right front area. Observing the scenarios in Fig. 2(b) - Fig. 2(e), we can find that the scenario with more agents appearing in left front and right front has larger value in the corresponding dimension of the global feature. Fig. 2(f) - Fig. 2(j) demonstrates another case that the agent should focus on past status of the front vehicle. Examining the high activation of the cases, we can find it actually corresponds to the decelerating of front vehicle. It is clearly seen that the learned partition of spatio-temporal space by our proposed method coincides with human interpretation.

c) *Recursive Refinement*: Although this aggregator is capable to aggregate the overall environment around the target agent, the pooling operator is still a first-order aggregator. It fails to model the second-order information such as interactions between agents. As a simple solution, we concatenate the global context feature to every individual feature, and recursively apply the aforementioned steps. In the second step, the embedding is aware of the status of individual agent and all the global context, thus could capture the interactions.

We summarize the pseudocodes of our encoder structure in Alg. 1. It can be easily implemented with several lines of code.

C. Future Trajectory Decoder

The decoder should be designed to satisfy various demands from downstream modules like planning. We believe that when the spatio-temporal context is well modeled, off-the-shelf decoder can still achieve superior results. As a default option, we feed the encoded spatio-temporal feature into a standard LSTM. At each time step, we minimize the l_2 loss between predicted 2D positions and the ground-truth. If desired, it is easy to extend to a stochastic decoder. We could simply inject a zero-mean Gaussian noise to the encoded feature C in each decoding process, and using variety loss as supervision as in[1], [7], [10].

In section IV, extensive experiments show that our method can achieve state-of-the-art performance in both deterministic and stochastic prediction tasks despite the simplicity of our decoder structure.

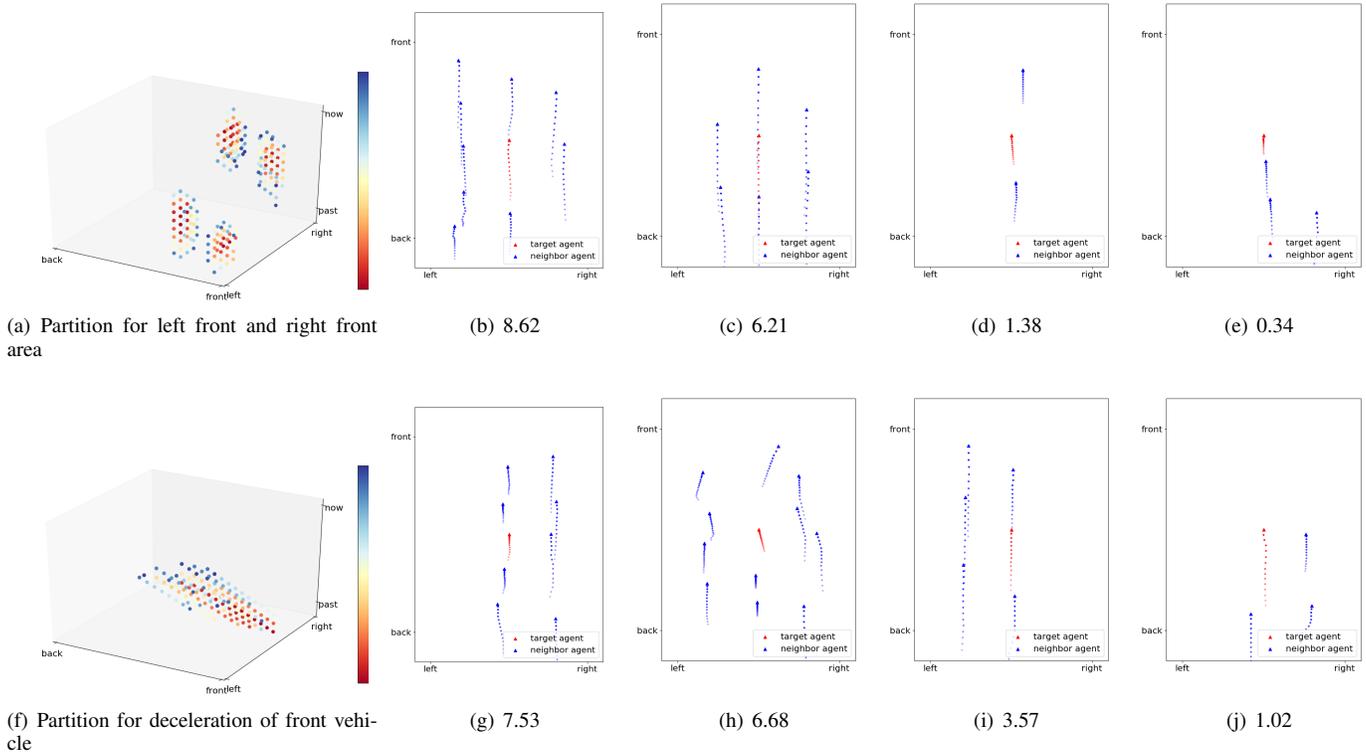


Fig. 2. (a) and (f) show activation patterns of two typical neurons in the pooled spatio-temporal features. The number in other subfigures indicates the value of activation of this neuron of the case.

Algorithm 1 Pseudocode of Spatio-temporal pooling in a PyTorch-like style.

```
def embedding(x, n_layers):
    for layer in range(n_layers):
        x = Linear(x)
        x = BN(x)
        x = ReLU(x)
    return x

def STPooling(x):
    '''spatial temporal context extractor
    Args:
        x: [N, K, H].
            N samples with K nbrs,
            snapshot's dimension H,
    Return:
        context: [N, O]
    '''
    e1 = embedding(x, n_layers=2)
    context1 = pooling(e1, dim=1, keep_dim=True)
    context1 = context1.repeat(1, x.shape[1], 1)
    context2 = cat((context1, x), dim=2)
    e2 = embedding(context2, n_layers=2)
    context2 = pooling(e2, dim=1)
    return context2
```

pooling: max pooling; cat: concatenation; BN: batch normalization.

D. Discussion

In this section, we highlight several key advantages of our proposed method, and compare with other state-of-the-art method.

a) RNN free: One benefit is that we do not rely on a fixed-length history trajectory, thus we can better handle the missing data due to the false negative of detection or tracker interruption. In contrast, current RNN based methods [1],

[2] have to eliminate the incomplete history of an agent because RNN can not differentiate the exact time step of the remaining trajectory, which reduces the amount of effective context one can utilize when modeling the behavior of the target agent. Second, RNN-based encoder for varying length trajectories is hard to be batch parallelled. Padding dummy values can be a solution, but it will introduce noise and affect performance. Besides, RNN-based encoder does not work when the time gap of the history sequence is not fixed, which is usually the case in real world autonomous driving scenarios because of the latency of the upstream fusion and perception modules. On the contrary, the time dimension in our input representation can be a dynamic floating point, which is much more flexible.

b) Comparisons with other social pooling methods: Although there are several works that also use pooling module to model social interaction [3], their methods and motivation are different from us. Their method needs to partition the space manually, and compute additional occupancy grids for each agent, while ours has already embedded and partitioned both space and time adaptively thanks to the unified spatio-temporal representation. Therefore, a simple global pooling is enough.

c) The importance of modeling cross time step social interaction: People need the reaction time to handle the sudden circumstance on the road. For example, there exists a delay before one can step on brakes when the front car slows down. In other words, an agent's behavior may be

affected by the actions of other agents in previous time steps. Thus, extracting the cross-time step interactions among agents becomes necessary.

IV. EXPERIMENT

We compare our proposed method with other state-of-the-art methods in two recently published trajectory datasets: ApolloScape [29], [18] and Argoverse [30]. Both of these datasets are collected from the first-person perspective by sensor-equipped acquisition cars in real-world driving scenarios, but they focus on different problems in the trajectory prediction field. We use the official metrics for these datasets and compare our proposed UST with the state-of-the-arts. In addition, we construct variants of our method for ablation study and qualitative analysis to show how UST can model complex spatio-temporal context in the high dimensional space.

A. Dataset and Experiment Setting

In this section, we detail the datasets as well as the corresponding experiment settings.

a) ApolloScape [29]: ApolloScape dataset is recorded in urban streets by various sensors (e.g. LiDAR, radar, camera). This dataset provides 3 seconds as trajectory history and aims to predict the next 3 seconds at 0.5s interval. The scene consists of heterogeneous traffic agents, including pedestrians, vehicles and bicycles. So the spatio-temporal context is much more complex to model. To give a fair comparison for this heterogeneous traffic-agent dataset, we follow the same experiment setting in ApolloScape challenge. We submit our result to the challenge leaderboard, and compare it with other submitted methods. To measure the performance of algorithms, we report the prediction error of each type of traffic-agents. The main metrics used for this dataset are Average Distance Error (ADE) and Final Distance Error (FDE).

- ADE: mean Euclidean distance between predicted coordinates and the ground truth over all time steps.

$$\text{ADE} = \frac{\sum_{n=1}^N \sum_{t=1}^T \|\tilde{x}_n^t - x_n^t\|_2}{N * T}$$

- FDE: Euclidean distance between the predicted coordinates and the ground truth at the final prediction timestep.

$$\text{FDE} = \frac{\sum_{n=1}^N \|\tilde{x}_n^T - x_n^T\|_2}{N}$$

Based on these metrics, we report a weighted sum of ADE (WSADE) and weighted sum of FDE (WSFDE) by assigning coefficient 0.20, 0.58, 0.22 to vehicles, pedestrians, bicycles respectively as in the ApolloScape Challenge.

$$\text{WSADE} = 0.2 \cdot \text{ADE}_v + 0.58 \cdot \text{ADE}_p + 0.22 \cdot \text{ADE}_b$$

$$\text{WSFDE} = 0.2 \cdot \text{FDE}_v + 0.58 \cdot \text{FDE}_p + 0.22 \cdot \text{FDE}_b$$

WSADE is also the metric by which the ApolloScape challenge is ranked. Because the difference of intrinsic behavior between pedestrians and vehicles, we separate the model for pedestrians by increasing the pedestrians' weight in the loss function.

b) Argoverse [30]: Argoverse is a large-scale autonomous driving dataset, containing 320 hours of data as well as rich map information. The given vector map is a semantic graph that provides detailed lane information that an agent might follow. Therefore we can get the multi-modal future trajectories explicitly with the help of multiple candidate lane centerlines. The collected trajectories in this dataset are individual 5 seconds trajectory segments. The first 2 seconds are used as history and predict spatial locations of the vehicles for up to 5 seconds. The traffic-agent type of Argoverse is not provided, so it is not included in our input representation.

We follow the official metric to benchmark multiple predictions on this dataset. The metric we choose is Minimum over N (MoN) metric as in previous works [1], [32]. It computes the error between the ground truth and the closest trajectory provided in the N output predictions. Specifically, we evaluate the top-N ADE and FDE with $N = 6$. To generate multi-modal future trajectory, we use the basic implementation offered by Argoverse baseline. In training, we choose a *2D curvilinear coordinate system* with axes tangential and perpendicular to the most possible centerline of the trajectory. At inference, we generate diverse future trajectories by using a different centerline as the 2-d curvilinear coordinate system. With multiple candidate centerlines, we can define various origins and reference frames to predict diverse futures.

c) NGSIM: Besides real-world autonomous driving dataset, we also evaluate our method on a widely used trajectory dataset Next Generation Simulation (NGSIM) dataset [28]. It consists of 45 minutes of highway driving trajectories at 10Hz for each roadway and contains various traffic conditions and diverse interactions among different traffic-agents. Note that trajectories in this dataset are recorded from fixed bird-eye view cameras. It can observe traffic agents without any coverage or missing data, which is inconsistent with the real autonomous driving scenarios. Thus we only list the results for reference. We adopt the same experiment setting as [2] to use 3s as history and predict 5s for future trajectory. The evaluation metric is root mean square error (RMSE) in meters over all future timesteps.

B. Implementation Details

We set the filter number of the fully connected layer to 128. Also the dimension of the hidden state for decoder LSTM keeps the same too. We iteratively train the network with a batch size of 128 for 50 epochs using Adam with an initial learning rate of 0.0003. We set the weight decay to 0.0001. The basic input of our model is 2D location, velocity, discrete time, and ID. We also add the type of agents on the ApolloScape dataset since in urban scenarios usually mixed types of agents are presented.

C. Results

We compare our proposed method with several state-of-the-art algorithms and baseline models in this section.

Methods	(a) ADE (3s)				(b) FDE (3s)			
	vehicles	pedestrians	bicycles	weighted	vehicles	pedestrians	bicycles	weighted
StarNet[31]	2.38	0.78	1.86	1.34	4.28	1.51	3.46	2.49
TrafficPredict[18]	7.94	7.18	12.88	8.58	12.77	11.12	22.79	24.22
LSTM	2.88	0.94	2.09	1.58	5.25	1.84	3.87	2.97
CV	2.59	0.81	2.17	1.47	4.64	1.58	4.02	2.73
UST	2.10	0.75	1.77	1.24	3.65	1.44	3.14	2.25

TABLE I

PERFORMANCE ON APOLLOSCAPE TRAJECTORY DATASET. ALL RESULTS ARE RECORDED FROM APOLLOSCAPE PUBLIC LEADERBOARD. OUR METHOD RANK FIRST IN ALL PUBLIC METHODS CURRENTLY. (A) SHOWS THE AVERAGE DISTANCE ERROR (ADE) OF DIFFERENT ALGORITHMS ON DIVERSE TRAFFIC-AGENTS. (B) SHOWS THE FINAL DISTANCE ERROR (FDE) IN 3 SECOND.

a) *ApolloScape*: State-of-the-art algorithms on ApolloScape and baseline models are shown as follows. The results are excerpted from ApolloScape leaderboard to make sure a fair comparison.

- StarNet[31]: The champion of trajectory prediction challenge in “CVPR 2019 Workshop on Autonomous Driving — Beyond Single Frame Perception”, which utilizes a centralized hub network to model spatio-temporal contexts with low time complexity.
- TrafficPredict[18]: An LSTM-based algorithm which focuses on predicting trajectories for heterogeneous traffic-agents in urban environment.
- LSTM: A simple LSTM encoder-decoder which only models each agent’s temporal context independently.
- CV(constant velocity): A baseline which uses a constant velocity Kalman filter to predict trajectories in the future.

In Table I, the results show that our proposed method achieves best performance in every metric and every agent type. Especially, compared with the current state-of-the-art method StarNet which is an LSTM-based attention network, we observe a **0.63m** improvement in FDE for the vehicle prediction task. Although they have leveraged a hub network to model agents’ interactions, the encoded context is still worse than ours. Our method improves considerably compared with all other methods of all three kinds of traffic agents.

Methods	3s ADE	3s FDE
NN + map	2.28	4.80
LSTM + map	2.25	4.67
LSTM + social + map	2.46	4.67
UST	1.47	2.94

TABLE II

EVALUATIONS OF UST AND MULTIPLE BASELINES ON THE ARGOVERSE DATASET. ERROR IS MINIMUM OVER SIX SAMPLES.

b) *Argoverse*: We compare our method with several baselines from Argoverse.

- Nearest Neighbor with map (NN + map): Weighted Nearest Neighbor regression trajectories where trajectories are queried by coordinate in the curvilinear coordinate system.

- LSTM + map: standard LSTM encoder-decoder structure, which uses map information to define multiple curvilinear coordinate systems. Therefore, it can generate multiple possible future trajectories.
- LSTM + social + map : Similar to LSTM + map, but with additional hand-crafted social features. Social features include minimum distance to the vehicle in front and in back, and number of neighbors.

Table II shows that UST outperforms all baselines on both 3s ADE and 3s FDE metrics. We also provide qualitative results to visualize how UST can predict multiple possible future trajectories according to maps (as shown in Fig. 3). Once the spatio-temporal context is well modeled, we can easily generate multiple multi-modal trajectories with additional map information under different complex conditions.

c) *NGSIM*: We report the performance of UST and several state-of-the-art algorithms on the NGSIM dataset (see Table. III). Only deterministic version result is reported. In addition, we construct a variant of our method denoted as *UST-180* which doubles the context range by enlarging the longitudinal direction range from ± 90 to ± 180 .

Without any doubt, UST outperforms any other algorithm in the same setting. In particular, the results of ± 180 improves dramatically. It outperforms the standard version, especially in the result of 5s. We believe that it is owe to the richer spatio-temporal context provided, and our unified modeling method could digest these information seamlessly.

D. Ablation Study

In this section, we design several variants of UST to investigate how different setting influences the final performance. All methods are evaluated on Argoverse validation set. The metric in ablation study is top-N ADE and FDE with N=1 to control all possible factors except the encoder.

a) *Input Representation*: We study the feature used in the input representation of UST. Table. IV summarizes the results. Note that the velocity comes from the differential of the provided position. As can be noticed, velocity only brings in marginal improvement. We owe this observation to that UST can model the relationship between different agents across different time steps, thus speed feature has already been implicitly learned. This property is especially useful when the observations from upstream are noisy, in which the results of explicit differential are not reliable.

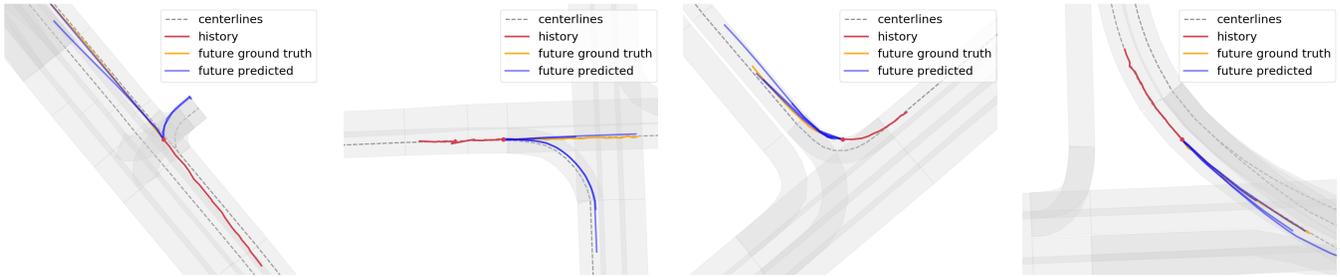


Fig. 3. Visualization of diverse trajectory predictions on Argoverse dataset. With basic implementation, we can still generate multiple diverse trajectory using the well modeled spatio-temporal context.

Time	CV	CV-GMM[33]	GAIL-GRU[34]	LSTM	MATF [7]	CS-LSTM [2]	S-LSTM [3]	UST	UST-180
1s	0.73	0.66	0.69	0.68	0.67	0.61	0.65	0.58	0.56
2s	1.78	1.56	1.56	1.65	1.51	1.27	1.31	1.20	1.15
3s	3.13	2.75	2.75	2.91	2.51	2.09	2.16	1.96	1.82
4s	4.78	4.24	4.24	4.46	3.71	3.10	3.25	2.92	2.58
5s	6.68	5.99	5.99	6.27	5.12	4.37	4.55	4.12	3.45

TABLE III

COMPARISON WITH DIFFERENT METHODS. ROOT MEAN SQUARE ERRORS (RMSE) IN METERS FROM 1s TO 5s ARE REPORTED.

Position	Time	Velocity	3s ADE	3s FDE
✓			3.46	7.60
✓	✓		2.71	6.01
✓	✓	✓	2.68	5.97

TABLE IV

RESULTS OF ABLATION STUDIES ON THE ARGOVERSE VALIDATION SET.

b) *Number of Times of Iterative Refinement*: We also conduct experiments to investigate the time of recursive refinement. The results in Fig. 4 show that when the time of refinement exceeds two, the performance becomes saturated. This may suggest that the gain of modeling even higher-order interactions more than two is marginal.

V. CONCLUSION

In this paper, we have presented UST for trajectory prediction problem, which integrates 2D locations and discrete time space into one unified 3D space, then learn the spatio-temporal context end-to-end. Despite the simplicity, it still shows state-of-the-art performance on various trajectory prediction datasets. We hope our method could be a strong baseline to trajectory prediction field and these encouraging results could inspire more advanced methods on the spatio-temporal context representation.

REFERENCES

[1] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social GAN: Socially acceptable trajectories with generative adversarial networks," in *CVPR*, 2018, pp. 2255–2264.
 [2] N. Deo and M. M. Trivedi, "Convolutional social pooling for vehicle trajectory prediction," in *CVPRW*, 2018, pp. 1468–1476.

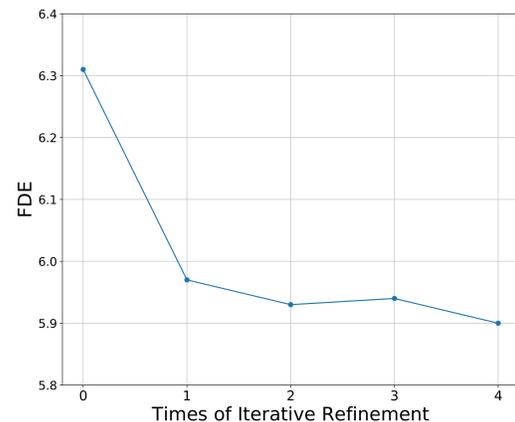


Fig. 4. Results on the Argoverse validation set under the different time of iterative refinement.

[3] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social LSTM: Human trajectory prediction in crowded spaces," in *CVPR*, 2016, pp. 961–971.
 [4] P. Zhang, W. Ouyang, P. Zhang, J. Xue, and N. Zheng, "SR-LSTM: State refinement for LSTM towards pedestrian trajectory prediction," in *CVPR*, 2019, pp. 12 085–12 094.
 [5] X. Li, X. Ying, and M. C. Chuah, "GRIP: Graph-based interaction-aware trajectory prediction," *arXiv preprint arXiv:1907.07792*, 2019.
 [6] A. Vemula, K. Muelling, and J. Oh, "Social attention: Modeling attention in human crowds," in *ICRA*, 2018, pp. 1–7.
 [7] T. Zhao, Y. Xu, M. Monfort, W. Choi, C. Baker, Y. Zhao, Y. Wang, and Y. N. Wu, "Multi-agent tensor fusion for contextual trajectory prediction," in *CVPR*, 2019, pp. 12 126–12 134.
 [8] K. Messaoud, I. Yahiaoui, A. Verroust-Blondet, and F. Nashashibi, "Non-local social pooling for vehicle trajectory prediction," 2019.
 [9] V. Kosaraju, A. Sadeghian, R. Martín-Martín, I. Reid, S. H. Rezatofighi, and S. Savarese, "Social-BiGAT: Multimodal trajectory forecasting using bicycle-gan and graph attention networks," *arXiv*

preprint *arXiv:1907.03395*, 2019.

- [10] Y. Huang, H. Bi, Z. Li, T. Mao, and Z. Wang, “STGAT: Modeling spatial-temporal interactions for human trajectory prediction,” in *ICCV*, 2019, pp. 6272–6281.
- [11] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, “Pointnet: Deep learning on point sets for 3D classification and segmentation,” in *CVPR*, 2017, pp. 652–660.
- [12] D. Helbing and P. Molnar, “Social force model for pedestrian dynamics,” *Physical Review E*, vol. 51, no. 5, p. 4282, 1995.
- [13] M. Luber, J. A. Stork, G. D. Tipaldi, and K. O. Arras, “People tracking with human motion predictions from social forces,” in *ICRA*, 2010, pp. 464–469.
- [14] G. Ferrer, A. Garrell, and A. Sanfeliu, “Robot companion: A social-force based approach with human awareness-navigation in crowded environments,” in *IROS*, 2013, pp. 1688–1694.
- [15] K. Yamaguchi, A. C. Berg, L. E. Ortiz, and T. L. Berg, “Who are you with and where are you going?” in *CVPR*, 2011, pp. 1345–1352.
- [16] W. Choi and S. Savarese, “Understanding collective activities of people from videos,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 6, pp. 1242–1257, 2013.
- [17] W. Ding, J. Chen, and S. Shen, “Predicting vehicle behaviors over an extended horizon using behavior interaction network,” *arXiv preprint arXiv:1903.00848*, 2019.
- [18] Y. Ma, X. Zhu, S. Zhang, R. Yang, W. Wang, and D. Manocha, “Trafficpredict: Trajectory prediction for heterogeneous traffic-agents,” in *AAAI*, 2019, pp. 6120–6127.
- [19] B. Kim, C. M. Kang, J. Kim, S. H. Lee, C. C. Chung, and J. W. Choi, “Probabilistic vehicle trajectory prediction over occupancy grid map via recurrent neural network,” in *ITSC*, 2017, pp. 399–404.
- [20] S. Yi, H. Li, and X. Wang, “Pedestrian behavior understanding and prediction with deep neural networks,” in *CVPR*, 2016, pp. 263–279.
- [21] Y. Li, “Which way are you going? Imitative decision learning for path forecasting in dynamic scenes,” in *CVPR*, 2019, pp. 294–303.
- [22] H. Cui, V. Radosavljevic, F.-C. Chou, T.-H. Lin, T. Nguyen, T.-K. Huang, J. Schneider, and N. Djuric, “Multimodal trajectory predictions for autonomous driving using deep convolutional networks,” in *ICRA*, 2019, pp. 2090–2096.
- [23] N. Djuric, V. Radosavljevic, H. Cui, T. Nguyen, F.-C. Chou, T.-H. Lin, and J. Schneider, “Motion prediction of traffic actors for autonomous driving using deep convolutional networks,” *arXiv preprint arXiv:1808.05819*, 2018.
- [24] N. Nikhil and B. Tran Morris, “Convolutional neural network for trajectory prediction,” in *ECCV*, 2018.
- [25] C. Choi and B. Dariush, “Looking to relations for future trajectory forecast,” *arXiv preprint arXiv:1905.08855*, 2019.
- [26] Y. Chai, B. Sapp, M. Bansal, and D. Anguelov, “Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction,” *arXiv preprint arXiv:1910.05449*, 2019.
- [27] T. Koskela, M. Lehtokangas, J. Saarinen, and K. Kaski, “Time series prediction with multilayer perception, fir and elman neural networks,” 1996. [Online]. Available: <https://academic.microsoft.com/paper/1558675754>
- [28] J. Colyar and J. Halkias, “Us highway 101 dataset,” *Federal Highway Administration (FHWA), Tech. Rep. FHWA-HRT-07-030*, 2007.
- [29] X. Huang, X. Cheng, Q. Geng, B. Cao, D. Zhou, P. Wang, Y. Lin, and R. Yang, “The ApolloScape dataset for autonomous driving,” in *CVPRW*, 2018, pp. 954–960.
- [30] M.-F. Chang, J. Lambert, P. Sangkloy, J. Singh, S. Bak, A. Hartnett, D. Wang, P. Carr, S. Lucey, D. Ramanan, *et al.*, “Argoverse: 3D tracking and forecasting with rich maps,” in *CVPR*, 2019, pp. 8748–8757.
- [31] Y. Zhu, D. Qian, D. Ren, and H. Xia, “StarNet: Pedestrian trajectory prediction using deep neural network in star topology,” *arXiv preprint arXiv:1906.01797*, 2019.
- [32] N. Lee, W. Choi, P. Vernaza, C. B. Choy, P. H. Torr, and M. Chandraker, “Desire: Distant future prediction in dynamic scenes with interacting agents,” in *CVPR*, 2017, pp. 336–345.
- [33] N. Deo, A. Rangesh, and M. M. Trivedi, “How would surround vehicles move? A unified framework for maneuver classification and motion prediction,” *IEEE Transactions on Intelligent Vehicles*, vol. 3, no. 2, pp. 129–140, 2018.
- [34] A. Kuefler, J. Morton, T. Wheeler, and M. Kochenderfer, “Imitating driver behavior with generative adversarial networks,” in *IV*. IEEE, 2017, pp. 204–211.