View Vertically: A Hierarchical Network for Trajectory Prediction via Fourier Spectrums

Conghao Wong^{*1}, Beihao Xia^{*1}, Ziming Hong¹, Qinmu Peng¹, Wei Yuan¹, Qiong Cao², Yibo Yang², and Xinge You¹

¹ Huazhong University of Science and Technology ² JD Explore Academy {conghao_wong, xbh_hust}@hust.edu.cn

Abstract. Understanding and forecasting future trajectories of agents are critical for behavior analysis, robot navigation, autonomous cars, and other related applications. Previous methods mostly treat trajectory prediction as time sequence generation. Different from them, this work studies agents' trajectories in a "vertical" view, i.e., modeling and forecasting trajectories from the spectral domain. Different frequency bands in the trajectory spectrums could hierarchically reflect agents' motion preferences at different scales. The low-frequency and high-frequency portions could represent their coarse motion trends and fine motion variations, respectively. Accordingly, we propose a hierarchical network V^2 -Net, which contains two sub-networks, to hierarchically model and predict agents' trajectories with trajectory spectrums. The coarse-level keypoints estimation sub-network first predicts the "minimal" spectrums of agents' trajectories on several "key" frequency portions. Then the finelevel spectrum interpolation sub-network interpolates the spectrums to reconstruct the final predictions. Experimental results display the competitiveness and superiority of V²-Net on both ETH-UCY benchmark and the Stanford Drone Dataset.

Keywords: Hierarchical Trajectory Prediction, Fourier Spectrums

1 Introduction

Trajectory prediction aims at inferring agents' possible future trajectories considering potential influencing factors. It has been an essential but challenging task, which can be widely applied to behavior analysis [2,5,39], robot navigation [48], autonomous driving [8,9,19,22], tracking [38,43], detection [10,11,34], and many other computer vision tasks [36,53]. Researchers have widely studied interactive factors, including the agent-agent interaction [1,14] and the agent-scene interaction [21,42]. Another line of researchers have explored creative ways to model trajectories. Neural networks like Long-Short Term Memory Networks (LSTMs) [16,30,55], Graph Convolution Networks (GCNs) [24,46,56], and Transformers

^{*} Equal contribution.

Codes are available at https://github.com/cocoon2wong/Vertical.



Fig. 1. V^2 -Net Motivation. (a) We show the trajectory in two views, the time sequences view and the Fourier spectrums view. The x-axis trajectory (green dots) and the y-axis trajectory (yellow dots) have similar "shapes", while they are quite different in the spectrum views (shown with amplitudes and phases), which motivates us to view trajectories "vertically" with spectrums rather than time sequences. (b) shows a trajectory with interactions. We utilize the first n/N low-frequency portions to reconstruct the trajectory, and the results show that different frequency portions have different contributions to their global plannings and interactive preferences.

[13,57] are employed to encode agents' trajectories. Some researchers have also studied agents' multi-modal characteristics [5,17] and the distributions of their goals [7,12,40], thus forecasting multiple future choices.

However, most previous methods treat trajectory prediction as time sequence generation and predict trajectories recurrently, which could be challenging to reflect agents' motion preferences at different scales hierarchically. In other words, researchers usually tend to explore agents' behaviors and their changing trends dynamically but lack the overall analyses at different temporal scales. In fact, pedestrians always plan their activities at different levels simultaneously. For example, they may first determine their coarse motion trends and then make fine decisions to interactive variations. Although some methods [57,13,58] employ neural networks with attention mechanisms (like Transformers) as the backbone to model agents' status, they may still be difficult to directly represent agents' detailed motion differences at different temporal scales.

The Fourier transform (FT) and its variations have significantly succeeded in the signal processing community. Recently, researchers have also introduced FTs to some computer vision tasks, such as image de-noising [20], edge extraction [18], and image super-resolution [6]. FTs decompose sequential inputs into a series of sinusoids with different amplitudes and phases on different frequencies. Furthermore, these sinusoids could reflect the differentiated frequency response characteristics at different frequency scales, which could be difficult to obtain directly in the original signal. They provide a "vertical" view for processing and analyzing sequences, thus presenting elusive features in the original signals.

Some researchers have applied FTs in tasks similar to trajectory prediction. For example, Mao *et al.* [32,31] employ the discrete cosine transform (DCT) to

3

help predict human motions. Cao *et al.* [3] use the graph Fourier transform to attempt to model interaction series in trajectory prediction. Unfortunately, FTs have not been employed to *model trajectories* in the trajectory prediction task directly. Inspired by the successful use of FTs, we try to employ the Discrete Fourier Transform (DFT) to obtain trajectory spectrums to capture agents' detailed motion preferences at different frequency scales. Fig.1(a) demonstrates an agent's two-dimensional trajectory in the time series and spectrum view, respectively. Although the two projection trajectories have similar shapes, they show considerable differences in the spectrum view. It means that the trajectory spectrums (including amplitudes and phases) obtained through FTs could reflect subtle differences that are difficult to discern within the time series.

Furthermore, some works have divided trajectory prediction into a two-stage pipeline, which we call the *hierarchical* prediction strategy. [47,29] divide this two-stage process into destination prediction and destination-controlled prediction. [28] introduces several "waypoints" to help predict agents' potential intentions rather than the only destination point. FTs decompose time series into the combination of different frequency portions. Inspired by these hierarchical approaches, a natural thought is to hierarchically predict agents' trajectories on different *frequency scales*, including: (a) *Global Plannings*, *i.e.*, agents' coarse motion trends. The low-frequency portions (slow-changing portions) in the trajectory spectrums could reflect their plannings globally. (b) *Interactive Preferences*, *i.e.*, agents' detailed interactions. The high-frequency portions (fast-changing portions) in the spectrums will directly show the rapidly changing movement rules, thus further describing their personalized interactive preferences.

Fig.1(b) demonstrates the trajectory reconstructed by different number of frequency portions. Agent's overall plannings could be reconstructed through a few low-frequency portions of the trajectory spectrum. By continually adding new high-frequency portions, the reconstructed trajectory would be able to reflect finer motion details and interactive preferences. Accordingly, we introduce a "coarse-fine" strategy to hierarchically model global plannings and interactive preferences by trajectory spectrums at different levels correspondingly.

In this paper, we propose the V^2 -Net to hierarchically forecast trajectories in the spectrum view. It contains two sub-networks. The coarse-level keypoints estimation sub-network first predicts the "minimal" spectrums of agents' trajectories on several "key" frequency portions, and then the fine-level spectrum interpolation sub-network interpolates the spectrums to reconstruct the final predictions. Meanwhile, agents' social and scene interactions will also be concerned to make the network available to give predictions that conform to social rules and physical constraints. Our contributions are list as follows:

- We introduce the Fourier Transform to model and predict trajectories with spectrums to better capture agents' behaviors from a different perspective.
- A Transformer-based V²-Net containing coarse-level keypoints estimation and fine-level spectrum interpolation sub-networks is proposed to model and predict trajectories "coarse-to-fine" hierarchically.

- 4 C. Wong, B. Xia, et al.
- Experiments demonstrate that V²-Net achieves competitive performance on both ETH-UCY benchmark and the Stanford Drone Dataset.

2 Related Work

Trajectory Prediction. Recently, researchers have studied how interactive factors affect agents' trajectory plans, like social interaction [1,15,35,59] and scene interaction [30,55]. The modeling of agents' trajectories also matters how the trajectory prediction networks perform. Alahi et al. [1] treat this task as a sequence generation problem, and they employ LSTMs to model and predict pedestrians' positions in the next time step recurrently. A series of works have also introduced Graph Neural Networks (GNNs), e.g. Graph Attention Networks (GATs) [15,21,26] and GCNs [4,35], to handle interactive factors when forecasting. Moreover, the attention mechanism has been employed [50,54,59] to focus on the most valuable interactive targets to give reasonable predictions. With the success of Transformers [49] in sequence processing such as natural language processing, researchers [13,57] have employed Transformers to obtain better feature representations. Some researchers address agents' uncertainty and randomness by introducing generative models. Generative Adversarial Networks (GANs) are employed in [14, 42, 21] to generate multiple stochastic trajectories to suit agents' various future choices. Some works like [17,44] use the Conditional Variation AutoEncoder (CVAE) to achieve the similar goal.

Applications of Fourier Transforms. Many approaches achieve better performance in different computer vision tasks by introducing FTs. Cheng et al. [6] present a Fast Fourier Transform-based algorithm, which brings high computational efficiency and reliability for the multichannel interpolation in image superresolution. Kaur et al. [18] propose a Fractional Fourier Transform based Riesz fractional derivative approach for edge detection and apply it to enhance images. Komatsu et al. [20] construct the 3-D mean-separation-type short-time DFT and apply it to denoise moving images. It is worth noting that FTs have been widely applied in handling time-series forecasting problems. Mao et al. [31,32] employ DCT to help predict human skeleton-graphs in motion forecasting. Cao et al. [4,3] propose spectral temporal graph to model interaction series (not trajectories) in trajectory forecasting. Forecasting trajectories could also be regarded as one of the time-series forecasting tasks [1]. The successful use of FTs motivates us to model agents' trajectories in the Fourier domain, therefore trying to obtain better representations. Unfortunately, there seems to be no method that directly uses FTs to describe agents' trajectories in the field of trajectory prediction. This paper attempts to model and forecast *agents' trajectories* in the spectrum view for the first time.

Hierarchical Trajectory Prediction. More and more researchers [7,12,40] have treated trajectory prediction as a "two-stage" problem and predict trajectories *hierarchically*. Some researchers divide this process into "destination



Fig. 2. V^2 -Net Overview. The coarse-level keypoints estimation sub-network aims at forecasting the N_{key} -point keypoint spectrums $[A^{key}, \Phi^{key}]$ from the observed trajectory spectrums $[A, \Phi]$. The other fine-level spectrum interpolation sub-network aims at predicting the complete spectrums $[\hat{A}, \hat{\Phi}]$ on each future moment based on the keypoint spectrums and the interaction representations.

prediction" and "destination-controlled prediction". Mangalam *et al.* [29] infer trajectory endpoints to assist in long-range multi-modal trajectory prediction. Tran *et al.* [47] attempt to obtain multi-modal goal proposals from the additional goal channel to generate multiple predictions. Wong *et al.* [51] use a set of generators to give destination proposals with different styles, and then interpolate to forecast multiple future predictions. Others like [28] also introduce several "waypoints" to help better predict agents' potential future intentions rather than the only destination points. These hierarchical approaches concern more on factor "destination" (or "waypoints") and "predictions". Considering the Fourier transforms could decompose time series into the combination of different frequency portions, a natural thought is to predict agents' future trajectories hierarchically on different *frequency scales*. Therefore, these two-stage factors may become the "low-frequency-portions" and "high-frequency-portions" for trajectory spectrums. It motivates us to forecast trajectories from coarse global plannings to fine interactive preferences with trajectory spectrums hierarchically.

3 Method

As shown in Fig.2, V²-Net has two main sub-networks, the coarse-level keypoints estimation sub-network and the fine-level spectrum interpolation sub-network.

Formulations. Let $p_t = (x_t, y_t) \in \mathbb{R}^2$ denote the two-dimensional coordinates of one agent at step t. Given a video clip $\{I\}$ that contains M agents' observed trajectories $\{X_i\}_{i=1}^M (X_i = (p_1, p_2, ..., p_{t_h})^T$ represents *i*-th agent's observed trajectory) during the observation period, trajectory prediction aims to forecast their possible future coordinates $\{\hat{Y}_i\}_{i=1}^M (\hat{Y}_i = (\hat{p}_{t_h+1}, \hat{p}_{t_h+2}, ..., \hat{p}_{t_h+t_f})^T$ denotes one of the predictions) during the corresponding future t_f steps considering their observations and the interactive context.

Table 1. Architecture details of the proposed V^2 -Net.

Layers	Network Architecture
MLP_t	$(a_x, a_y, \phi_x, \phi_y) \to \text{fc}(64, \text{ReLU}) \to \text{fc}(64, \tanh) \to f_t$
MLP_c	$C \to \operatorname{MaxPool}(5 \times 5) \to \operatorname{Flatten} \to \operatorname{fc}(64t_h, \operatorname{tanh}) \to \operatorname{Reshape}(t_h, 64) \to f_c$
T_k	$f_e \to \text{TransformerEncoder} (128) \to f'_k;$
	$f'_k, (a_x, a_y, \phi_x, \phi_y) \to \text{TransformerDecoder} (128) \to f''_k$
MLP_e	$f_k'' \to \text{fc}(128, \tanh) \to \text{fc}(128) \to f$
MLP_d	$f \to \text{fc}(128, \text{ReLU}) \to \text{fc}(128, \text{ReLU}) \to \text{fc}(4N_{key}) \to$
	Reshape $(N_{key}, 4) \rightarrow (a_x^{key}, a_y^{key}, \phi_x^{key}, \phi_y^{key})$
T_i	$[f_t^{key}, f_c] \to \text{TransformerEncoder} (128) \to f'_i;$
	$f'_i, (a^{key}_x, a^{key}_y, \phi^{key}_x, \phi^{key}_y) \to \text{TransformerDecoder} (4) \to (\hat{a}_x, \hat{a}_y, \hat{\phi}_x, \hat{\phi}_y)$

Keypoints Estimation Sub-network. The coarse-level keypoints estimation sub-network aims to forecast agents' keypoint spectrums with *a lower spatio-temporal resolution*. An encoder-decoder structure is designed to generate multiple random predictions to adapt to agents' uncertainty and randomness.

We first use Discrete Fourier Transform (DFT) on agents' observed trajectories to obtain the t_h -point trajectory spectrums. We apply 1D-DFT on *each dimension* in the observed trajectory $X = \{(x_t, y_t)\}_{t=1}^{t_h}$ to obtain their spectrums, including the amplitudes $A = \{a_x, a_y\}$ and the phases $\Phi = \{\phi_x, \phi_y\}$:

$$\mathcal{X} = \mathrm{DFT}[(x_1, x_2, ..., x_{t_h})] \in \mathbb{C}^{t_h}, \quad \mathcal{Y} = \mathrm{DFT}[(y_1, y_2, ..., y_{t_h})] \in \mathbb{C}^{t_h},$$

$$A = \{a_x, a_y\} = \{\|\mathcal{X}\|, \|\mathcal{Y}\|\}, \quad \Phi = \{\phi_x, \phi_y\} = \{\arg \mathcal{X}, \arg \mathcal{Y}\}.$$
(1)

We employ an embedding MLP (the MLP_t in Table 1) to embed agents' observed trajectory spectrums $(a_x, a_y, \phi_x, \phi_y) \in \mathbb{R}^{t_h \times 4}$ into the high-dimensional f_t . Similar to previous works like S-GAN [14], we model agents' multimodality by sampling the random noise vector $z \sim \mathcal{N}(0, I)$ and then concatenating the corresponding random representations f_i to the f_t . The encoder for these noise vectors (we call the MLP_i) has the same structure like the MLP_t. We combine the above representations to obtain the embedded vector f_e . Formally,

$$f_t = \mathrm{MLP}_t((a_x, a_y, \phi_x, \phi_y)) \in \mathbb{R}^{t_h \times 64},$$

$$f_i = \mathrm{MLP}_i(z) \in \mathbb{R}^{t_h \times 64},$$

$$f_e = [f_t, f_i] \in \mathbb{R}^{t_h \times 128}.$$
(2)

Here, [a, b] represents the concatenation for vectors $\{a, b\}$ on the last dimension.

Then, we use a Transformer[49] named T_k to encode agents' behavior representations. The embedded vector f_e is passed to the Transformer encoder, and the spectrums $(a_x, a_y, \phi_x, \phi_y)$ are input to the Transformer decoder. The Transformer here is used as the feature extractor, and it does not contain the final output layer. We employ another MLP encoder (MLP_e) to aggregate features at different frequency nodes [51], thus inferring the behavior feature f:

$$f = \mathrm{MLP}_e\left(\mathrm{T}_k(f_e, (a_x, a_y, \phi_x, \phi_y))\right) \in \mathbb{R}^{t_h \times 128}.$$
(3)

The sub-network finally utilizes a decoder MLP (MLP_d) to predict agents' N_{key} -point ($N_{key} < t_f$) keypoint spectrums [A^{key}, Φ^{key}]. Formally,

$$[A^{key}, \Phi^{key}] = (a_x^{key}, a_y^{key}, \phi_x^{key}, \phi_y^{key}) = \mathrm{MLP}_d(f) \in \mathbb{R}^{N_{key} \times 4}.$$
 (4)

We call N_{key} the number of spectrum keypoints.

When training the sub-network, agents' key spatial coordinates y^{key} (which are gathered from their entire ground-truth future trajectories) will be used as the supervision. Meanwhile, the Inverse DFT (IDFT) will be applied to obtain the predicted spatial keypoints \hat{y}^{key} . The network variables will be tuned by minimizing the average Euclidean distance between y^{key} and the predicted \hat{y}^{key} , therefore learning to predict the corresponding keypoint spectrums. We define the **Average Keypoints Loss** (\mathcal{L}_{AKL}) as:

$$\mathcal{L}_{\text{AKL}} = \|\hat{y}^{key} - y^{key}\|_2 = \frac{1}{N_{key}} \sum_{n=1}^{N_{key}} \|\hat{p}_{t_n^{key}} - p_{t_n^{key}}\|_2,$$
(5)

where

$$\begin{split} \hat{y}^{key} &= \left(\text{IDFT}[a_x^{key} \exp(j\phi_x^{key})], \text{ IDFT}[a_y^{key} \exp(j\phi_y^{key})] \right) \in \mathbb{R}^{N_{key} \times 2}, \\ y^{key} &= \left(p_{t_1^{key}}, p_{t_2^{key}}, ..., p_{t_{N_{key}}^{key}} \right)^T \in \mathbb{R}^{N_{key} \times 2}. \end{split}$$
(6)

Spectrum Interpolation Sub-network. The fine-level spectrum interpolation sub-network reconstructs the complete trajectory spectrums from the keypoint spectrums with *a higher spatio-temporal resolution*. It aims to learn the spectrum biases between the complete spectrums and the keypoint spectrums.

The sub-network takes the N_{key} -point keypoint spectrums $[A^{key}, \Phi^{key}]$ as the input. Similar to Equation 2, we have the representation f_t^{key} :

$$f_t^{key} = \mathrm{MLP}_t((a_x^{key}, a_y^{key}, \phi_x^{key}, \phi_y^{key})) \in \mathbb{R}^{N_{key} \times 64}.$$
 (7)

Note that MLP_t in Equation 2 and Equation 7 do not share weights.

Besides, agents' social and scene interactions will be concerned in this subnetwork. In detail, we use a context embedding MLP (MLP_c) to encode the transferred images C [52] (which encodes both social interactions and scene constraints together in an energy map form by the scene visual image I and trajectories \mathcal{T} via a CNN) into the context feature $f_c = \text{MLP}_c(C)$.

We employ a similar Transformer (called the Interpolation Transformer, T_i) to learn the spectrum biases. We pass the concatenated feature $f_e^{key} = [f_t^{key}, f_c]$ to the transformer encoder, and keypoint spectrums $[A^{key}, \Phi_x^{key}] = (a_x^{key}, a_y^{key}, \phi_x^{key}, \phi_y^{key})$ to the transformer decoder. The transformer here is used to forecast the **complete spectrums** $[\hat{A}, \hat{\Phi}] = (\hat{a}_x, \hat{a}_y, \hat{\phi}_x, \hat{\phi}_y)$. Then, we use the IDFT to obtain the reconstructed trajectory \hat{y}_o . Formally,

$$(\hat{a}_x, \hat{a}_y, \hat{\phi}_x, \hat{\phi}_y) = \mathcal{T}_i(f_e^{key}, (a_x^{key}, a_y^{key}, \phi_x^{key}, \phi_y^{key})) \in \mathbb{R}^{(t_h + t_f) \times 4},$$

$$\hat{y}_o = \left(\text{IDFT}[\hat{a}_x \exp(j\hat{\phi}_x)], \text{ IDFT}[\hat{a}_y \exp(j\hat{\phi}_y)] \right) \in \mathbb{R}^{(t_h + t_f) \times 2}.$$

$$(8)$$

7

Finally, we have one of the V^2 -Net predictions:

$$\hat{y} = \hat{y}_o[t_h:,:] \in \mathbb{R}^{t_f \times 2},\tag{9}$$

where the $[t_h :, :]$ indicates the slicing operation on tensors.

The spectrum interpolation sub-network learns to interpolate the "key" trajectory spectrums into the complete trajectory spectrums, thus reflecting agents' fine interactive details by predicting the remaining spectrum portions. Its variables will be tuned through the **Average Point-wise Loss** (\mathcal{L}_{APL}):

$$\mathcal{L}_{\text{APL}} = \|\hat{y} - y\|_2 = \frac{1}{t_f} \sum_{t=t_h+1}^{t_h+t_f} \|\hat{p}_t - p_t\|_2.$$
(10)

Loss Functions. We use the joint loss function to tain V^2 -Net:

$$\mathcal{L} = \mathcal{L}_{AKL} + \mathcal{L}_{APL}.$$
 (11)

4 Experiments

Datasets. (a) **ETH-UCY Benchmark**: Many previous methods like [1,14,42] take several sub-datasets from ETH[38] and UCY[23] to train and evaluate their models with the "leave-one-out" strategy [1], which is called the ETH-UCY benchmark. It contains 1536 pedestrians with thousands of non-linear trajectories. The annotations are pedestrians' coordinates in meters. (b) **Stanford Drone Dataset**: The Stanford Drone Dataset [41] (SDD) has 60 bird-view videos captured by drones. More than 11,000 different agents are annotated with bounding boxes in pixels. It contains over 185,000 social interactions and 40,000 scene interactions, which is more complex and challenging. ³

Metrics. We employ two metrics to evaluate the prediction accuracy, including the Average Displacement Error (ADE) and the Final Displacement Error (FDE) [1,38]. ADE is the average point-wise Euclidean distance between groundtruth and predictions of all steps, and FDE is the Euclidean distance between the last point's prediction and groundtruth. Following [14], the reported metrics are the minimum value in 20 generations ("best-of-20"). For one prediction, we have:

ADE =
$$\frac{1}{t_f} \sum_{t=t_h+1}^{t_h+t_f} \|p_t - \hat{p}_t\|_2$$
, FDE = $\|p_{t_h+t_f} - \hat{p}_{t_h+t_f}\|_2$. (12)

Baselines. We choose several methods as our baselines, including S-GAN [14], SoPhie [42], Social-BiGAT [21], E-SR-LSTM [60], MANTRA [33], Multiverse [26], SimAug [25], PECNet [29], STAR [57], TPNMS [27], TF [13], Trajectron++ [44], Introvert [45], LB-EBM [37], Agentformer [58], Y-net [28], and SpecTGNN [3].

³ Dataset splits used to train and validate on SDD are the same as [25].

Implementation Details. We predict agents' trajectories in future $t_f = 12$ frames according to their $t_h = 8$ frames' observations. The frame rate is set to 2.5 frames per second when sampling trajectories. We train the entire V²-Net with the Adam optimizer (learning rate lr = 0.0003) on one NVIDIA Tesla P4 graphic processor. V²-Net is trained with the batch size bs = 2500 for 800 epochs on ETH-UCY and 150 epochs on SDD. Detailed layer connections, output units, and activations are listed in Table 1. We employ L = 4 layers of encoder-decoder structure with H = 8 attention heads in each Transformer-based sub-network. The output dimension of fully connected layer used in multi-head attention layers is set to 128. The default number of spectrum keypoints is set to $N_{key} = 3$. When training the network, we set $\{t_1^{key}, t_2^{key}, t_3^{key}\} = \{t_h + 4, t_h + 8, t_h + 12\}$.

4.1 Comparision to State-of-the-Art Methods

ETH-UCY. As shown in Table 2, V²-Net achieves better performance on most ETH-UCY sub-datasets. It improves the metrics on eth by 23.3% and 28.8% compared with LB-EBM. Meanwhile, V²-Net has better performance than Introvert on hotel by about 8.8% FDE. Compared with Introvert, its average ADE and FDE have been improved by 14.2% and 17.6%, respectively. Moreover, V²-Net has achieved comparable average performance to the state-of-the-art Agentformer and Y-net. It shows V²-Net's strong competitiveness on ETH-UCY. Especially, V²-Net dramatically outperforms Agentformer and Y-net about 14.8% ADE on eth, which demonstrates its unique advantages.

SDD. As shown in Table 3, V²-Net has better performance than the previous state-of-the-art methods on SDD. Compared to SpecTGNN, V²-Net has significantly gained by 13.3% and 8.2% in ADE and FDE, respectively. V²-Net also outperforms the state-of-the-art Y-net by 9.3% and 3.9%. In general, V²-Net has a better prediction accuracy than other baselines on SDD. It is worth noting that SDD is more complex than ETH-UCY, which demonstrates V²-Net's robustness and adaptability to larger and more challenging scenarios.

4.2 Quantitative Analysis

We design several model variations and run ablation studies to verify each design in the proposed V²-Net. Quantitative results are shown in Table 4 and Table 5.

DFT on Trajectories. V²-Net implements on the trajectory spectrums to focus on different frequency portions in different prediction stages. Variations a1 and a2 are trained with the same number of spectrum keypoints $N_{key} = 1$, while a1 does not use DFT/IDFT when predicting. It means that variation a2 obtains behavior features and interaction features from the spectrums, while a1 obtains them from the time series. Experiments show that a2 outperforms a1 by about 10% average ADE and FDE on ETH-UCY. In addition, we remove all

Table 2. Comparisons to baselines with the *best-of-20* on ETH-UCY. Metrics are shown in the format of "ADE/FDE" in meters. Lower means better.

Models	eth	hotel	univ	zara1	zara2	Average
Social-BiGAT [21]	0.69/1.29	0.49/1.01	0.55/1.32	0.30/0.62	0.36/0.75	0.48/1.00
E-SR-LSTM [60]	0.44/0.79	0.19/0.31	0.50/1.05	0.32/0.64	0.27/0.54	0.34/0.67
PECNet [29]	0.54/0.87	0.18/0.24	0.35/0.60	0.22/0.39	0.17/0.30	0.29/0.48
STAR [57]	0.56/1.11	0.26/0.50	0.52/1.13	0.40/0.89	0.52/1.13	0.41/0.87
TPNMS $[27]$	0.52/0.89	0.22/0.39	0.55/0.13	0.35/0.70	0.27/0.56	0.38/0.73
TF [13]	0.61/1.12	0.18/0.30	0.35/0.65	0.22/0.38	0.17/0.32	0.31/0.55
Trajectron++[44]	0.43/0.86	0.12/0.19	0.22/0.43	0.17/0.32	0.12/0.25	0.20 /0.39
Introvert [45]	0.42/0.70	0.11/0.17	0.20/0.32	0.16 /0.27	0.16/0.25	0.21/0.34
LB-EBM [37]	0.30/0.52	0.13/0.20	0.27/0.52	0.20/0.37	0.15/0.29	0.21/0.38
Agentformer [58]	0.26/0.39	0.11/0.14	0.26/0.46	0.15/0.23	0.14/0.23	0.18/0.29
Y-net [28]	0.28/ 0.33	0.10/0.14	0.24/0.41	0.17/0.27	0.13/0.22	0.18/0.27
V^2 -Net (Ours)	0.23 /0.37	0.11/0.16	0.21 /0.35	0.19/0.30	0.14/0.24	0.18/0.28

Table 3. Comparisons to baselines with the *best-of-20* on SDD. Lower means better.

Models	S-GAN [14]	SoPhie [42]	Multiverse [26]	SimAug [25]	PECNet
ADE/FDE	27.25/41.44	16.27/29.38	14.78/27.09	12.03/23.98	9.96/15.88
Models	MANTRA [33]	LB-EBM	SpecTGNN [3]	Y-net	V^2 -Net (Ours)
ADE/FDE	8 96/17 76	8 87/15 61	8 21/12 41	7.85/11.85	7.12/11.39

Table 4. Ablation Studies. S1 and S2 represent the keypoints estimation stage and the spectrum interpolation stage, correspondingly. "D" indicates whether different variations model trajectories with spectrums. "N" indicates the number of spectrum keypoints. Results in " \uparrow " are the percentage metric improvements compared to a1.

No.	S1	S2	D	Ν	eth	hotel	univ	zara1	zara2	\uparrow Gain (%)
a1	\checkmark	\checkmark	×	1	0.29/0.47	0.12/0.18	0.23/0.38	0.25/0.37	0.17/0.29	- (base)
a2	\checkmark	\checkmark	\checkmark	1	0.23/0.39	0.12/0.17	0.23/0.37	0.22/0.32	0.16/0.26	7.9%/11.3%
b1	\checkmark	\checkmark	\checkmark	3	0.23/0.37	0.11/0.16	0.21/0.35	0.19/0.30	0.14/0.24	16.2%/16.5%
b2	\checkmark	\checkmark	\checkmark	4	0.24/0.38	0.12/0.17	0.22/0.36	0.21/0.32	0.16/0.26	9.5%/12.4%
b3	\checkmark	\checkmark	\checkmark	6	0.29/0.49	0.13/0.17	0.23/0.38	0.26/0.37	0.21/0.31	-6.0%/-1.2%
с	\checkmark	×	\checkmark	3	0.24/0.37	0.12/0.16	0.23/0.35	0.21/0.30	0.15/0.25	9.6%/15.9%

Table 5. Validation of DFT. We design and run two minimal model variations to show the performance gain brought by DFT. Symbols are the same as Table 4.

No.	S1	S2	D	Ν	eth	hotel	univ	zara1	zara2	\uparrow Gain (%)
$\min 1$	\checkmark	×	Х	-	0.83/1.66	0.25/0.44	0.77/1.39	0.48/0.97	0.38/0.74	- (base)
$\min 2$	\checkmark	\times	\checkmark	-	0.79/1.51	0.22/0.38	0.55/1.10	0.46/0.92	0.34/0.71	12.9%/11.1%

other components from the model and design two minimal model variations min1 and min2 to demonstrate more directly the usefulness of introducing DFT for trajectory prediction. These two variations DO NOT take into account agents' interactions and multimodal properties, but only retain the Transformer backbone to achieve the "simple" trajectory prediction goal. We train both the two variations under the same conditions. Results in Table 5 favorably demonstrate the performance gains (12.9% on ADE and 11.1% on FDE) by using trajectory spectrums, independent of the rest of the designs in the proposed network. It proves the effectiveness of DFT in this task, and further demonstrates the considerable performance gain by implementing on the trajectory spectrums.

Number of Spectrum Keypoints N_{key} . The number of spectrum keypoints to be predicted matters how the keypoints estimation sub-network determines agents' overall motion trends. A smaller N_{key} might cause a looser planning division, making it challenging to reflect the differences between agents' similar future choices. On the contrary, a larger value may lead to a strict trends division, which could be difficult for the subsequent network to reflect agents' multiple uncertain future choices and motion preferences. Ablation experiments on variations a2, b1, b2, and b3 show the quantitative comparisons with different N_{key} configurations. The temporal keypoints are set to the equivalence point among the predicted period $(t_f = 12)$. For example, when $N_{key} = 1$, we set $\{t_1^{key}\} = \{t_h + 12\}$, and when $N_{key} = 4$, we set $\{t_1^{key}, t_2^{key}, t_3^{key}, t_4^{key}\} = \{t_h + 3, t_h + 6, t_h + 9, t_h + 12\}$. It shows that variation b1 $(N_{key} = 3)$ outperforms variation a2 $(N_{key} = 1)$ with additional improvements by about 5% ADE and 6% FDE. However, The performance of variation b3 $(N_{key} = 6)$ reduces by about 5% and 1% compared to variation a1. Please refer to visualized trajectories of different N_{key} variations in section "Qualitative Analysis" to see how they affect the prediction performance qualitatively.

Hierarchical Prediction. Similar to some goal-driven methods, V^2 -Net splits trajectory prediction into a two-stage "keypoints-interpolation" process. However, the most significant difference from these methods is that our model implements the entire process with trajectory spectrums rather than spatial coordinates. Correspondingly, the second stage "goal-conditioned prediction" has turned into the "spectrum interpolation". Although the first stage sub-network has already given a low-resolution prediction composed of a few keypoint spectrums, it is not easy to reflect agents' subtle interactions and activity differences. The spectrum interpolation sub-network aims to interpolate these "key" spectrums into the complete trajectory spectrums, thus reflecting agents' fine-level interactive preferences. Variation c only implements the S1 sub-network but utilizes linear interpolation to finish forecasting. Variations b1 and c could show the improvements brought by the interpolation sub-network. Results point out that b1 improves about 7% in ADE compared with variation c. It effectively reflects the superiority of the "coarse-to-fine" hierarchical prediction with spectrums.



Fig. 3. Keypoints and Interpolation Illustration. We show the predicted spatial keypoints $(N_{key} = 3)$ after IDFT and the corresponding predicted trajectories V²-Net finally outputs. Their corresponding time steps are distinguished with different color masks. Dots connected by white arrows belong to the same prediction steps.



Fig. 4. Number of Spectrum Keypoints and Prediction Styles. We show the visualized predictions with different N_{key} configurations. Each sample has 20 random predictions.

4.3 Qualitative Analysis

Coarse-Fine Prediction. V^2 -Net predicts agents' multiple random keypoint spectrums at the first stage, and then uses a spectrum interpolation method to forecast the complete spectrums by considering interactions and fine activities. Fig.3 shows two visualized results at keypoints estimation and spectrum interpolation stages correspondingly. At the first stage, V^2 -Net aims to predict keypoint spectrums with a low spatio-temporal resolution. We apply IDFT on these keypoint spectrums to show the corresponding spatial keypoints in Fig.3. As shown in Fig.3(a) and (b), these spatial keypoints could only represent the general trend of agents' future activities rather than the specific and precise coordinates. Then, the interpolation sub-network considers both physical constraints and social interactions to "fine-tune" these keypoints. For example, the marked keypoints in Fig.3(b) indicate that the pedestrian could go toward the road. However, the corresponding final prediction shows that he might slow down and stay on the side of the road considering the scene's physical constraints.

Number of Spectrum Keypoints and Prediction Styles. The keypoint sub-network estimate agents' keypoint spectrums to determine their coarse future routes. Meanwhile, it contains a generative module to give agents multiple



Fig. 5. Visualized Predictions. Each sample illustrates 20 random predictions.

predictions by randomly sampling $z \sim N(0, I)$, therefore adapting to agents' uncertain future decisions. The number of spectrum keypoints N_{key} matters directly V²-Net's stochastic performance, which can be seen in Fig.4. It shows that Fig.4(a) exhibits a looser restraint when generating trajectories. Obviously, it is caused by the insufficient trends constraints from the little N_{key} . Although (a) gives a large number of future predictions, most of them may be invalid or do not meet the scene constraints. Visualized results indicate that V²-Net will give more rigorous predictions when N_{key} increases. It will forecast more acceptable predictions when increasing the N_{key} to 3 (Fig.4(b)). However, its predictions will be limited to a small area when N_{key} is set to 6 (Fig.4(c)). At this time, V²-Net could hardly reflect agents' multiple stochastic plannings under the constraints of such excessive keypoints. In short, with a low number of keypoints a wider spectrum of possibilities is covered (higher multimodality) and vice versa by increasing the number the trajectories become more plausible (collisions with static elements are avoided). It also explains the drop of quantitative performance when $N_{key} = 6$ in the above ablation studies. Furthermore, we can set different N_{key} in different prediction scenarios to achieve controllable predictions.

Visualization. As shown in Fig.5, V²-Net could predict multiple trajectories considering both scene constraints and social interactions. For example, V²-Net gives a variety of different future options for bikers crossing the intersection (illustrated in Fig.5(a)(e)(k)). In the prediction case (e) and (k), it considers the physical constraints of the environment and forecast to bypasses the grass. Additionally, it also shows strong adaptability in some unique prediction scenarios. For instance, it gives three kinds of predictions to the biker in case (h) to pass through the traffic circle: turning right, going ahead, and turning left. When turning left, the given predictions are not to turn left directly (just like turning right) but to go left after going around the traffic circle, which is consistent with the traffic rules around the circle. Case (i) presents similar results for the man



Fig. 6. Failure Prediction Cases.

passing the crossroads. It shows V^2 -Net could accurately describe the scene's physical constraints and agents' motion rules in various scenes.

Surprisingly, V^2 -Net could give "smoother" predictions due to the usage of DFT, like turning left in case (h). The spatio-temporal continuity between adjacent points in the predicted time series will be ensured with the superposition of a series of sinusoids. Therefore, the predictions could reflect the real physical characteristics of the target agent while keeping their interactive preferences.

Limitations. As shown in Fig.6, V²-Net has failed predictions for stationary agents and few interactive cases. Although not all predictions are regarded as failed predictions, some of them show not-so-reasonable movement trends for these standing still agents. Additionally, it may provide predictions not that suitable for the interactive context. We will fix this problem in the future.

5 Conclusion

In this work, we focus on giving predictions in two hierarchical steps by modeling their trajectories with spectrums from the global plannings and the interactive preferences. A Transformer-based V²-Net, which consists of the coarselevel keypoints estimation sub-network and the fine-level spectrum interpolation sub-network, is proposed to predict agents' trajectories hierarchically. Experimental results show that V²-Net achieves the competitive performance on both ETH-UCY benchmark and the Stanford Drone Dataset. Although the proposed method shows higher prediction accuracy and provides better visualized results, there are still some failed predictions. We will address this problem and further explore feasible solutions to model and predict agents' possible trajectories.

Acknowledgements. This work was partially supported by the National Natural Science Foundation of China (Grant No. 62172177), and in part by the Fundamental Research Funds for the Central Universities (Grant No. 2021yjsCXCY040).

References

- Alahi, A., Goel, K., Ramanathan, V., Robicquet, A., Fei-Fei, L., Savarese, S.: Social lstm: Human trajectory prediction in crowded spaces. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 961–971 (2016)
- Alahi, A., Ramanathan, V., Goel, K., Robicquet, A., Sadeghian, A.A., Fei-Fei, L., Savarese, S.: Learning to predict human behavior in crowded scenes. In: Group and Crowd Behavior for Computer Vision, pp. 183–207. Elsevier (2017)
- Cao, D., Li, J., Ma, H., Tomizuka, M.: Spectral temporal graph neural network for trajectory prediction. In: 2021 IEEE International Conference on Robotics and Automation (ICRA). pp. 1839–1845. IEEE (2021)
- Cao, D., Wang, Y., Duan, J., Zhang, C., Zhu, X., Huang, C., Tong, Y., Xu, B., Bai, J., Tong, J., et al.: Spectral temporal graph neural network for multivariate time-series forecasting. Advances in Neural Information Processing Systems 33, 17766–17778 (2020)
- Chai, Y., Sapp, B., Bansal, M., Anguelov, D.: Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. arXiv preprint arXiv:1910.05449 (2019)
- Cheng, D., Kou, K.I.: Fft multichannel interpolation and application to image super-resolution. Signal Processing 162, 21–34 (2019)
- Choi, C., Malla, S., Patil, A., Choi, J.H.: Drogon: A trajectory prediction model based on intention-conditioned behavior reasoning. arXiv preprint arXiv:1908.00024 (2019)
- Cui, H., Radosavljevic, V., Chou, F.C., Lin, T.H., Nguyen, T., Huang, T.K., Schneider, J., Djuric, N.: Multimodal trajectory predictions for autonomous driving using deep convolutional networks. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 2090–2096. IEEE (2019)
- Deo, N., Trivedi, M.M.: Convolutional social pooling for vehicle trajectory prediction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 1468–1476 (2018)
- Ergezer, H., Leblebicioğlu, K.: Anomaly detection and activity perception using covariance descriptor for trajectories. In: European Conference on Computer Vision. pp. 728–742. Springer (2016)
- Fernando, T., Denman, S., Sridharan, S., Fookes, C.: Soft+ hardwired attention: An lstm framework for human trajectory prediction and abnormal event detection. Neural networks 108, 466–478 (2018)
- Girase, H., Gang, H., Malla, S., Li, J., Kanehara, A., Mangalam, K., Choi, C.: Loki: Long term and key intentions for trajectory prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9803–9812 (2021)
- Giuliari, F., Hasan, I., Cristani, M., Galasso, F.: Transformer networks for trajectory forecasting pp. 10335–10342 (2021)
- Gupta, A., Johnson, J., Fei-Fei, L., Savarese, S., Alahi, A.: Social gan: Socially acceptable trajectories with generative adversarial networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2255– 2264 (2018)
- Huang, Y., Bi, H., Li, Z., Mao, T., Wang, Z.: Stgat: Modeling spatial-temporal interactions for human trajectory prediction. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 6272–6281 (2019)
- Hug, R., Becker, S., Hübner, W., Arens, M.: Particle-based pedestrian path prediction using lstm-mdl models. In: 2018 21st International Conference on Intelligent Transportation Systems (ITSC). pp. 2684–2691. IEEE (2018)

- 16 C. Wong, B. Xia, et al.
- Ivanovic, B., Pavone, M.: The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2375–2384 (2019)
- Kaur, K., Jindal, N., Singh, K.: Fractional fourier transform based riesz fractional derivative approach for edge detection and its application in image enhancement. Signal Processing 180, 107852 (2021)
- Kim, B., Kang, C.M., Kim, J., Lee, S.H., Chung, C.C., Choi, J.W.: Probabilistic vehicle trajectory prediction over occupancy grid map via recurrent neural network. In: 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC). pp. 399–404. IEEE (2017)
- Komatsu, T., Tyon, K., Saito, T.: 3-d mean-separation-type short-time dft with its application to moving-image denoising. In: 2017 IEEE International Conference on Image Processing (ICIP). pp. 2961–2965 (2017)
- Kosaraju, V., Sadeghian, A., Martín-Martín, R., Reid, I., Rezatofighi, H., Savarese, S.: Social-bigat: Multimodal trajectory forecasting using bicycle-gan and graph attention networks. In: Advances in Neural Information Processing Systems. pp. 137–146 (2019)
- Lee, N., Choi, W., Vernaza, P., Choy, C.B., Torr, P.H., Chandraker, M.: Desire: Distant future prediction in dynamic scenes with interacting agents. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 336–345 (2017)
- Lerner, A., Chrysanthou, Y., Lischinski, D.: Crowds by example. Computer Graphics Forum 26(3), 655–664 (2007)
- Li, S., Zhou, Y., Yi, J., Gall, J.: Spatial-temporal consistency network for lowlatency trajectory forecasting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 1940–1949 (October 2021)
- Liang, J., Jiang, L., Hauptmann, A.: Simaug: Learning robust representations from simulation for trajectory prediction. In: Proceedings of the European conference on computer vision (ECCV) (August 2020)
- Liang, J., Jiang, L., Murphy, K., Yu, T., Hauptmann, A.: The garden of forking paths: Towards multi-future trajectory prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10508– 10518 (2020)
- 27. Liang, R., Li, Y., Li, X., Zhou, J., Zou, W., et al.: Temporal pyramid network for pedestrian trajectory prediction with multi-supervision. arXiv preprint arXiv:2012.01884 (2020)
- Mangalam, K., An, Y., Girase, H., Malik, J.: From goals, waypoints & paths to long term human trajectory forecasting. arXiv preprint arXiv:2012.01526 (2020)
- Mangalam, K., Girase, H., Agarwal, S., Lee, K.H., Adeli, E., Malik, J., Gaidon, A.: It is not the journey but the destination: Endpoint conditioned trajectory prediction. In: European Conference on Computer Vision. pp. 759–776 (2020)
- Manh, H., Alaghband, G.: Scene-lstm: A model for human trajectory prediction. arXiv preprint arXiv:1808.04018 (2018)
- Mao, W., Liu, M., Salzmann, M.: History repeats itself: Human motion prediction via motion attention. In: European Conference on Computer Vision. pp. 474–489. Springer (2020)
- Mao, W., Liu, M., Salzmann, M., Li, H.: Learning trajectory dependencies for human motion prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9489–9497 (2019)

- Marchetti, F., Becattini, F., Seidenari, L., Bimbo, A.D.: Mantra: Memory augmented networks for multiple trajectory prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7143– 7152 (2020)
- Mehran, R., Oyama, A., Shah, M.: Abnormal crowd behavior detection using social force model. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. pp. 935–942. IEEE (2009)
- Mohamed, A., Qian, K., Elhoseiny, M., Claudel, C.: Social-stgcnn: A social spatiotemporal graph convolutional neural network for human trajectory prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14424–14432 (2020)
- Morris, B.T., Trivedi, M.M.: Trajectory learning for activity understanding: Unsupervised, multilevel, and long-term adaptive approach. IEEE transactions on pattern analysis and machine intelligence 33(11), 2287–2301 (2011)
- Pang, B., Zhao, T., Xie, X., Wu, Y.N.: Trajectory prediction with latent belief energy-based model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11814–11824 (2021)
- Pellegrini, S., Ess, A., Schindler, K., Van Gool, L.: You'll never walk alone: Modeling social behavior for multi-target tracking. In: 2009 IEEE 12th International Conference on Computer Vision. pp. 261–268. IEEE (2009)
- Phan-Minh, T., Grigore, E.C., Boulton, F.A., Beijbom, O., Wolff, E.M.: Covernet: Multimodal behavior prediction using trajectory sets. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14074– 14083 (2020)
- Rhinehart, N., McAllister, R., Kitani, K., Levine, S.: Precog: Prediction conditioned on goals in visual multi-agent settings. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2821–2830 (2019)
- Robicquet, A., Sadeghian, A., Alahi, A., Savarese, S.: Learning social etiquette: Human trajectory understanding in crowded scenes. In: European conference on computer vision. pp. 549–565. Springer (2016)
- 42. Sadeghian, A., Kosaraju, V., Sadeghian, A., Hirose, N., Rezatofighi, H., Savarese, S.: 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. pp. 1349–1358 (2019)
- 43. Saleh, F., Aliakbarian, S., Salzmann, M., Gould, S.: Artist: Autoregressive trajectory inpainting and scoring for tracking. arXiv preprint arXiv:2004.07482 (2020)
- 44. Salzmann, T., Ivanovic, B., Chakravarty, P., Pavone, M.: Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16. pp. 683–700. Springer (2020)
- Shafiee, N., Padir, T., Elhamifar, E.: Introvert: Human trajectory prediction via conditional 3d attention. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16815–16825 (2021)
- Sun, J., Jiang, Q., Lu, C.: Recursive social behavior graph for trajectory prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 660–669 (2020)
- Tran, H., Le, V., Tran, T.: Goal-driven long-term trajectory prediction. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 796–805 (2021)

- 18 C. Wong, B. Xia, et al.
- Trautman, P., Krause, A.: Unfreezing the robot: Navigation in dense, interacting crowds. In: 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 797–803. IEEE (2010)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in neural information processing systems. pp. 5998–6008 (2017)
- Vemula, A., Muelling, K., Oh, J.: Social attention: Modeling attention in human crowds. In: 2018 IEEE international Conference on Robotics and Automation (ICRA). pp. 1–7. IEEE (2018)
- 51. Wong, C., Xia, B., Peng, Q., You, X.: Msn: Multi-style network for trajectory prediction. arXiv preprint arXiv:2107.00932 (2021)
- Xia, B., Wong, C., Peng, Q., Yuan, W., You, X.: Cscnet: Contextual semantic consistency network for trajectory prediction in crowded spaces. Pattern Recognition p. 108552 (2022)
- 53. Xie, D., Shu, T., Todorovic, S., Zhu, S.C.: Learning and inferring "dark matter" and predicting human intents and trajectories in videos. IEEE transactions on pattern analysis and machine intelligence 40(7), 1639–1652 (2017)
- 54. Xu, Y., Piao, Z., Gao, S.: Encoding crowd interaction with deep neural network for pedestrian trajectory prediction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5275–5284 (2018)
- Xue, H., Huynh, D.Q., Reynolds, M.: Ss-lstm: A hierarchical lstm model for pedestrian trajectory prediction. In: 2018 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 1186–1194. IEEE (2018)
- Xue, H., Huynh, D.Q., Reynolds, M.: Scene gated social graph: Pedestrian trajectory prediction based on dynamic social graphs and scene constraints. arXiv preprint arXiv:2010.05507 (2020)
- Yu, C., Ma, X., Ren, J., Zhao, H., Yi, S.: Spatio-temporal graph transformer networks for pedestrian trajectory prediction. In: European Conference on Computer Vision. pp. 507–523. Springer (2020)
- Yuan, Y., Weng, X., Ou, Y., Kitani, K.M.: Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9813–9823 (2021)
- Zhang, P., Ouyang, W., Zhang, P., Xue, J., Zheng, N.: Sr-lstm: State refinement for lstm towards pedestrian trajectory prediction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 12085–12094 (2019)
- Zhang, P., Xue, J., Zhang, P., Zheng, N., Ouyang, W.: Social-aware pedestrian trajectory prediction via states refinement lstm. IEEE transactions on pattern analysis and machine intelligence (2020)