

EigenTrajectory: Low-Rank Descriptors for Multi-Modal Trajectory Forecasting

Inhwan Bae^{1*}

Jean Oh²

Hae-Gon Jeon^{1†}

¹GIST AI Graduate School

²Carnegie Mellon University

inhwanbae@gm.gist.ac.kr, jeanoh@cmu.edu, haegonj@gist.ac.kr

Abstract

Capturing high-dimensional social interactions and feasible futures is essential for predicting trajectories. To address this complex nature, several attempts have been devoted to reducing the dimensionality of the output variables via parametric curve fitting such as the Bézier curve and B-spline function. However, these functions, which originate in computer graphics fields, are not suitable to account for socially acceptable human dynamics. In this paper, we present EigenTrajectory (ET), a trajectory prediction approach that uses a novel trajectory descriptor to form a compact space, known here as ET space, in place of Euclidean space, for representing pedestrian movements. We first reduce the complexity of the trajectory descriptor via a low-rank approximation. We transform the pedestrians' history paths into our ET space represented by spatio-temporal principle components, and feed them into off-the-shelf trajectory forecasting models. The inputs and outputs of the models as well as social interactions are all gathered and aggregated in the corresponding ET space. Lastly, we propose a trajectory anchor-based refinement method to cover all possible futures in the proposed ET space. Extensive experiments demonstrate that our EigenTrajectory predictor can significantly improve both the prediction accuracy and reliability of existing trajectory forecasting models on public benchmarks, indicating that the proposed descriptor is suited to represent pedestrian behaviors. Code is publicly available at <https://github.com/inhwanbae/EigenTrajectory>.

1. Introduction

Trajectory prediction involves forecasting the future footsteps of agents based on their past movements. This task is considered one of the core technologies for autonomous navigation, social robot platforms and surveillance systems.

Many existing approaches [1, 2, 5, 10, 15, 19, 28, 35, 36, 37, 40, 48, 55, 69, 73, 74, 77, 101, 108, 109] design their prediction models in the Euclidean space, *i.e.*, to directly infer

*Work done during an internship at Carnegie Mellon University.

†Corresponding author

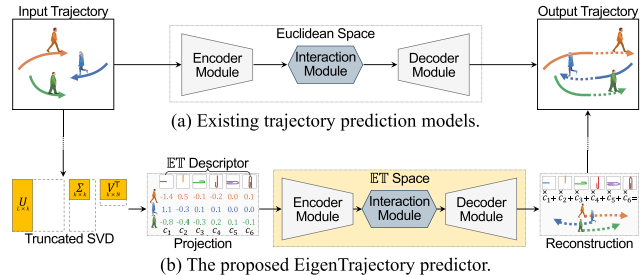


Figure 1. A common pipeline of trajectory prediction models and the proposed EigenTrajectory. For each observation, (a) existing models predict future trajectories using the raw data in Euclidean space; (b) our approach transforms the raw data into our ET space, then captures the social interaction and predicts the coefficient of our trajectory descriptor.

a sequence of 2D coordinates of future frames. These approaches force the models to learn both informative behavioral features and their motion dynamics from raw trajectory data. Such direct predictions can intuitively describe agents' behaviors in the temporal series of spatial coordinates; however, in a higher-dimensional space, it is hard for the models to determine explanatory features.

Recent works have described the pedestrian's movements using trajectory descriptors instead of dealing with all raw coordinate information. Inspired by human beings traveling pathways with a higher level of connotation (*e.g.*, a person who gradually decelerates to turn right, or makes a sharp turn while going straight) [66], parametric curve functions are introduced. In particular, Hug *et al.* [22, 23] introduce the Bézier curve, and Jazayeri and Jahangiri [25] propose a B-spline curve-based representation for effectively modeling continuous-time trajectories. These methods successfully reduce the dimensionality of trajectories by abstractly representing lengthy sequential coordinates in a spatial domain using a smaller set of key points. It is, however, unclear how well these parametric functions can capture human motions and behaviors as they have been designed for computer graphics [17, 62, 71] and part modeling [86, 89, 91].

In this paper, departing from existing parametric curve functions, we present an intuitive trajectory descriptor that is learnable from real-world human trajectory data as illus-

trated in Fig. 1. First, we decompose a stacked trajectory sequence using Singular Value Decomposition (SVD). To represent the data concisely, we reduce the dimensionality by performing the best rank- k approximation. Through this process, all trajectories can be approximated as a linear combination of k eigenvectors, which we call EigenTrajectory ($\mathbb{E}\mathbb{T}$) descriptor. Next, we aggregate the social interaction features and then predict the coefficients of the eigenvectors for the future path in the same space. Here, after clustering the coefficients, a set of trajectory anchors, which can be interpreted as the coefficient candidates, is used to ensure a diversity of prediction paths. Finally, trajectory coordinates can be reconstructed through the matrix multiplication of the eigenvector and the coefficients. Experimental results demonstrate that the proposed EigenTrajectory ($\mathbb{E}\mathbb{T}$) descriptor can successfully represent the pedestrian motion dynamics and significantly improve the prediction accuracy of existing prediction models on various public benchmarks.

2. Related Works

2.1. Pedestrian Trajectory Prediction

Pioneering works [20, 53, 61, 98] model pedestrian movements following the social force theory or energy minimization to describe surrounding interactions and future movements. Since then, there have been significant advances in both social interaction and motion modeling with the introduction of Convolutional Neural Networks (CNNs) and Recurrent neural networks (RNNs) for trajectory prediction. One pioneering work is Social-LSTM [1], which introduces a long short-term memory (LSTM) to recurrently predict future coordinates. Agent’s LSTM hidden states are shared with each other by a social pooling mechanism which leverages a neighbor’s hidden state information inside a spatial grid. Attention mechanisms [16, 24, 69, 84] are also used to directly share social relations based on neighbors’ influence. Better prediction results are achievable when using additional surveillance view images [12, 13, 14, 28, 41, 47, 49, 50, 51, 68, 70, 77, 78, 80, 82, 97, 103, 107]. Recently, graph neural network-based methods, which define a graph representation with pedestrian nodes and interaction edges, are successfully incorporated into the trajectory prediction. Graph Convolutional Networks (GCNs) [2, 27, 55, 78], Graph Attention networks (GATs) [5, 21, 28, 39, 40, 73, 83], and transformers [4, 18, 57, 88, 90, 101, 102] are utilized to update the graph-structured pedestrian features.

From the perspective of trajectory forecasting, probabilistic inferences are studied for multi-modal trajectory prediction. Gaussian modeling [1, 2, 36, 55, 73, 74, 75, 95, 99, 101] estimates the bi-variate parameters of distribution to represent the possibility of moving routes. The Generative Adversarial Network (GAN) [12, 19, 21, 28, 38, 42, 68, 77, 107] introduces additional discriminators to generate realistic paths.

Conditional Variational AutoEncoder (CVAE)-based methods [7, 11, 24, 30, 31, 35, 48, 69, 79, 85, 92, 94] utilize random latent vectors as features for each person’s walking direction and speed. They infer future trajectories recurrently [1, 8, 11, 18, 19, 30, 46, 52, 58, 63, 69, 96, 104, 106] or simultaneously [2, 3, 37, 55, 72, 73] in common. These approaches directly infer spatial coordinates in the Euclidean space without interpretable abstraction, which can suffer from an overfitting problem due to high dimensionality, and noisy paths can be generated as illustrated in Fig. 2.

2.2. Parametric Trajectory Descriptor

To achieve a high-level abstraction of trajectory data, several recent works have introduced parametric trajectory descriptors. Hug *et al.* [23] propose a probabilistic, parametric \mathcal{N} -curve model based on the Bézier curve. To generate multi-modal predictions, multiple curves are computed from the Gaussian mixture model of random variables. Hug *et al.* [22] also introduce a variation of the Mixture Density Networks which is operated in the Bézier curve domain instead of the raw data. A Gaussian process-based refinement framework is adopted in their model. Alternatively, another work in [25] represents each trajectory using a B-spline curve, and predicts the coefficients of this curve with the CVAE and inverse reinforcement learning.

The Bézier curve and B-spline use a Bernstein polynomial and piecewise polynomial as basis functions, respectively. These methods interpolate points on the curve through a weighted sum of a set of control points, but are still costly because it requires a large number of two-dimensional control points. One possible way to tackle this is to use polynomial curves; however, there are limitations on the precise path representations, and the abstract polynomial coefficients are difficult to learn. More fundamentally, although these works can successfully predict a temporally-smooth walking path using well-known curve functions, such parametric descriptors may be hard to fit irregular trajectories of human pedestrians.

3. Methodology

Our approach aims at learning an intuitive abstraction of human movements. To overcome the limited naturalness of parametric curves, we adopt a data-driven approach to leverage the distributions of human trajectories in the real-world data. Additionally, to address the dimensionality issue, we propose a low-rank approximation strategy to succinctly represent human trajectories in a more compact space than the Euclidean space. We achieve our goals by creating a novel trajectory descriptor specialized for representing human movements.

We start with a definition of trajectory prediction in Sec. 3.1. We then introduce the mathematical formulations of $\mathbb{E}\mathbb{T}$ descriptor and transformations between the




Figure 2. Visualization of the trajectory prediction strategies. The most widely used method is to directly predict 12 future raw points. The Bézier curve and B-spline are defined by a set of control points, and roughly model the movements. Our method can accurately represent pedestrian paths with a linear combination of our novel $\mathbb{E}\mathbb{T}$ descriptors.

Euclidean and $\mathbb{E}\mathbb{T}$ spaces in Sec. 3.2. Using our descriptor, we describe how well existing trajectory forecasting baseline models can be operated in our novel space in Sec. 3.3.

3.1. Problem Definition

The problem of trajectory prediction involves forecasting the future paths of agents in an environment from their path histories. Here, a trajectory can be represented by a temporal series of spatial points. Formally, the observation trajectory with length T_{obs} can be represented as $\mathbf{A}_n = \{(x_n^t, y_n^t) \mid t \in [1, \dots, T_{obs}]\}$, where (x_n^t, y_n^t) is the 2D spatial coordinate of a pedestrian n at specific time t . Similarly, a ground truth future trajectory for prediction time length T_{pred} can be defined as $\mathbf{B}_n = \{(x_n^t, y_n^t) \mid t \in [T_{obs} + 1, \dots, T_{pred}]\}$. A trajectory problem is defined as a pair $[\mathbf{A}, \mathbf{B}]$. The goal of a forecasting model is to predict s possible multi-modal future trajectories $\tilde{\mathbf{B}}^s$ given observation \mathbf{A} as input.

3.2. EigenTrajectory ($\mathbb{E}\mathbb{T}$) Descriptor

Inspired by successful uses of Eigenvalue decomposition and Singular Value Decomposition (SVD) in extracting key feature information from raw data in various domains [26, 33, 34, 60, 76, 81, 87], we propose an SVD-based approach for representing the gist of human trajectories.

We assume that pedestrian movements are mainly represented by directions and speed variations that can be captured via principal components in the data. Since a little noise from people’s staggering or tracking can also be added into the representation, we assume that using the top k components in the eigenvectors would be sufficient to correspond to the raw trajectory data. To fully benefit from the successful dimensionality reduction of SVD, we also apply it into our framework to represent spatio-temporal trajectory data as a linear combination of a set of eigenvectors. The proposed $\mathbb{E}\mathbb{T}$ descriptor consists of two elements: k most representative singular vectors of trajectories and their combination coefficients that can explain any given trajectory.




Figure 3. Visualization of the first six $\mathbb{E}\mathbb{T}$ descriptors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_6$ learned from the ETH/UCY dataset. The top row shows the examples of the $\mathbb{E}\mathbb{T}$ descriptors in 2D Euclidean space. The middle and bottom rows illustrate the temporal variations of the x and y positions, respectively.

Mathematical formulation. To make an efficient trajectory descriptor, we need to find the principal components of the trajectory distribution. In mathematical terms, our goal is to obtain the eigenvectors of a covariance matrix from raw trajectory data. We first construct the trajectory matrices \mathbf{A} and \mathbf{B} by stacking all pedestrians’ observations and predicted trajectories from a training set. We then apply SVD to both trajectory matrices \mathbf{A} and \mathbf{B} as follows:

$$\mathbf{A} = \mathbf{U}_{obs} \mathbf{\Sigma}_{obs} \mathbf{V}_{obs}^\top, \quad \mathbf{B} = \mathbf{U}_{pred} \mathbf{\Sigma}_{pred} \mathbf{V}_{pred}^\top, \quad (1)$$

where $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_L]$ and $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_N]$ are orthogonal matrices and $\mathbf{\Sigma}$ is a diagonal matrix, consisting of singular values $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > 0$. Here, L is the number of dimensions (*i.e.*, the number of parameters) of the trajectory, $L = 2 \times T_{obs}$ for observation and $L = 2 \times (T_{pred} - T_{obs})$ for the predicted trajectory. N is the number of pedestrians in the whole dataset, and r is the rank of \mathbf{A} and \mathbf{B} .

Since each trajectory can be more or less dominant to form a unique eigenvector, the resulting eigenvectors can be considered as a set of motion features that can jointly characterize trajectory variations.

Rank- k approximation. Based on the assumption that human trajectories can be described using a few key parameters on directions and speed, we can approximate trajectory $\tilde{\mathbf{A}}_n$ using a linear combination of the first k left singular vectors $\mathbf{U}_k = \mathbf{u}_1, \dots, \mathbf{u}_k$ as:

$$\tilde{\mathbf{A}}_n = \mathbf{U}_{obs,k} \mathbf{c}_{obs,n} = [\mathbf{u}_1, \dots, \mathbf{u}_k]_{obs} \mathbf{c}_{obs,n}, \quad (2)$$

where $\mathbf{c}_{obs,n}$ denotes a set of coefficients specifying the relevance of each principal component and $\tilde{\mathbf{A}}_n$ is known as the best rank- k approximation of \mathbf{A} . Here, k representative vectors are the eigenvectors of trajectories $\mathbf{A}\mathbf{A}^\top$ that is used to summarize the raw trajectories in our trajectory prediction approach. We apply the same approximation to the predicted trajectory \mathbf{B} as follows:

$$\tilde{\mathbf{B}}_n = \mathbf{U}_{pred,k} \mathbf{c}_{pred,n} = [\mathbf{u}_1, \dots, \mathbf{u}_k]_{pred} \mathbf{c}_{pred,n}. \quad (3)$$




Figure 4. Visualization of the rank- k trajectory approximation. Original trajectory sequences including straight, turning left, deceleration, and complex movement are successfully reconstructed by our $\mathbb{E}\mathbb{T}$ descriptor when increasing the k values. In this example, all four GT trajectories are sampled from the ZARA1 scene.

They can be regarded as the k principal components. A pair $(\mathbf{c}_n, \mathbf{U}_k)$ of the coefficients and rank- k singular vectors constitutes the $\mathbb{E}\mathbb{T}$ descriptor.

In Fig. 3, we visualize the $\mathbb{E}\mathbb{T}$ descriptors for predicted trajectory \mathbf{B} and $k = 6$ where the trajectories represented in rank- k vectors explain the spatial displacement of pedestrians over time. Specifically, \mathbf{u}_1 and \mathbf{u}_2 encode the constant velocity motion in the x and y directions, respectively. The reason why both eigenvectors appear first is that most people in the dataset walk at a constant speed while maintaining their directions. Next, \mathbf{u}_3 and \mathbf{u}_4 show velocity changes in the x and y -axis directions, respectively. Using the combination of the vectors, it is feasible to learn complex movements, *e.g.*, a person is going straight in the beginning, but then slows down or turns left/right. In \mathbf{u}_5 and \mathbf{u}_6 , we also observe a velocity change, but $\mathbb{E}\mathbb{T}$ descriptors can represent a more detailed trajectory by combining them with the previous eigenvector. We call the space spanned by $\mathbb{E}\mathbb{T}$ descriptors as the $\mathbb{E}\mathbb{T}$ space.

Transformation between spaces. Using the $\mathbb{E}\mathbb{T}$ descriptors, we can project a trajectory defined in Euclidean space into the $\mathbb{E}\mathbb{T}$ space. Given trajectories \mathbf{A}_n and \mathbf{B}_n , we project it as:

$$\mathbf{c}_{obs,n} = \mathbf{U}_{obs,k}^\top \mathbf{A}_n, \quad \mathbf{c}_{pred,n} = \mathbf{U}_{pred,k}^\top \mathbf{B}_n, \quad (4)$$

where \mathbf{c} is a coefficient vector defined in the $\mathbb{E}\mathbb{T}$ space. Each coefficient determines how much a corresponding $\mathbb{E}\mathbb{T}$ descriptors affects a pedestrian path.

Obviously, the inverse transformation of Eq. (4) is available. At this time, the reconstructed path is low-rank approximated. Since \mathbf{U} is an orthogonal matrix, it can be invertible through a transpose operation as follows:

$$\tilde{\mathbf{A}}_n = \mathbf{U}_{obs,k} \mathbf{c}_{obs,n}, \quad \tilde{\mathbf{B}}_n = \mathbf{U}_{pred,k} \mathbf{c}_{pred,n}. \quad (5)$$

As visualized in Fig. 4, a simple straight path can be expressed with only one $\mathbb{E}\mathbb{T}$ descriptor ($k = 1$). However, the more complex the trajectory, the larger value of k is needed. Nevertheless, most paths can be represented just with k , which is small enough compared to the original trajectory dimension $k \leq L$.




Figure 5. An overview of the trajectory anchor generation and anchor refinement step in our $\mathbb{E}\mathbb{T}$ space.

3.3. Forecasting in the $\mathbb{E}\mathbb{T}$ Space

Using the $\mathbb{E}\mathbb{T}$ descriptor and transformation operators, we optimize each module designed for trajectory forecasting in our $\mathbb{E}\mathbb{T}$ space. In addition, we propose a trajectory anchor to ensure the diversity of the predicted trajectories. This enables the off-the-shelf forecasting model to take the full benefits of the $\mathbb{E}\mathbb{T}$ descriptor in an end-to-end manner.

Trajectory prediction using the $\mathbb{E}\mathbb{T}$ descriptor. Given a trajectory prediction problem $[\mathbf{A}, \mathbf{B}]$, using Eqs. (2) and (3), we use the training data to compute the $\mathbb{E}\mathbb{T}$ descriptors for the observed and predicted trajectories, denoted by $(\mathbf{c}_{obs,n}, \mathbf{U}_{obs,k})$ and $(\mathbf{c}_{pred,n}, \mathbf{U}_{pred,k})$, respectively.

Trajectory anchor generation. Anchor-based methods are widely used in various fields, especially object detection [43, 64]. In the trajectory prediction, social anchors, offering the interpretability of egocentric characteristics, are introduced in [9, 29, 93, 105]. However, these approaches mainly use hand-crafted anchors and work in Euclidean or Polar coordinate systems. In this paper, we introduce a data-driven anchor used in our novel $\mathbb{E}\mathbb{T}$ space, which covers all feasible and diverse multi-modal futures.

We first normalize the translation, rotation, and speed of the trajectories in the training dataset as visualized in Fig. 5. Next, we project all normalized N paths into the $\mathbb{E}\mathbb{T}$ space using Eq. (4). We then cluster them in the $\mathbb{E}\mathbb{T}$ space instead of using them in the original Euclidean space. Since there is a duality between the Euclidean space and the $\mathbb{E}\mathbb{T}$ space with an isometry $\mathbf{U}_{pred,k}^\top$, the distance between the trajectories in the Euclidean space $\|\tilde{\mathbf{B}}_i - \tilde{\mathbf{B}}_j\|$ is equal to those between the corresponding coefficient vectors in the $\mathbb{E}\mathbb{T}$ space $\|\mathbf{c}_{pred,i} - \mathbf{c}_{pred,j}\|$. Hence, the clustering can be performed to yield the same results in both spaces, but it can be done more reliably and more efficiently in the $\mathbb{E}\mathbb{T}$ space. Note that the clustering becomes more difficult because of the well-known curse of dimensionality [6, 54] and this process is executed once at the initialization step before training the model. We use these clustered s centroids $\tilde{\mathbf{c}}_n^s$ as a trajectory coefficient anchor, and compute the correction offset for refinement in the prediction module.

Observations in the $\mathbb{E}\mathbb{T}$ space. As a next step, we project the input observed path into the $\mathbb{E}\mathbb{T}$ space using the $\mathbb{E}\mathbb{T}$ descriptor $\mathbf{U}_{obs,k}$ calculated in Eq. (2) and transformation function in Eq. (4). This higher level of abstraction for the input data representations according to the change of the trajec-




Figure 6. Illustration of the social interaction modeling strategies. (a) Social interactions between pedestrians are captured and gathered for each footstep in the Euclidean space. (b) The intra- and inter-personal interactions are modeled by ET descriptor coefficients in our ET space.

tory descriptor has a fundamental impact on the design of subsequent modules. In addition, the forecasting model can focus on the principal component of physical movements and has an additional effort of trajectory noise suppression by reducing the dimensionality.

History encoder in the ET space. The set of ET descriptor coefficients for the observation is then fed to the trajectory forecasting model as input. Since the Euclidean coordinate sequence \mathcal{A}_n can be transformed into a linear combination of ET descriptors $c_{obs,n}$, the model is able to omit a process for finding physical motion properties from the raw point and to immediately use the movement patterns.

Interaction modeling in the ET space. Since the method for the trajectory description uses the ET space, it is necessary to define social interactions in the same space. The recent models try to design interaction modules through a separation of the social relations between agents into spatial and temporal perspectives [55, 73, 101]. We map these spatial and temporal interactions one-to-one. To do this, we extract both the inter-personal and intra-personal interactions from our ET space, as illustrated in Fig. 6. In the inter-personal point of view, the relation between agents is captured with respect to the coefficients. Here, people with similar coefficients may be peers walking in a same direction; otherwise, they may need to avoid collisions. This information from each coefficient dimension is then aggregated in the intra-personal interaction stage to capture a rich social relation feature. We directly utilize the well-designed GNN and attention mechanism of the existing off-the-shelf models in our system with only a minor change for the dimension.

Extrapolation decoder in the ET space. With the rich features in the ET space, we also predict future trajectories in the same space. To fully reflect the multi-modal and non-linear characteristics of human motions, we introduce the trajectory anchor-based method defined in the preparation stage. Using the trajectory anchor, the model predicts the correction vector f , which has the same shape as the anchor. This correction vector is added to the initial value of the trajectory anchor to obtain the final s trajectory coefficients as follows:

$$\hat{c}_n^s = \bar{c}_n^s + f_n^s. \quad (6)$$

Descriptor	Dim	ETH	HOTEL	UNIV	ZARAI	ZARA2	AVG
Linear interpolation	4	386/703	143/215	163/344	256/426	155/232	221/384
Bézier curve [22, 23]	12	24/47	17/23	18/20	18/19	18/18	19/25
B-spline [25]	12	25/45	19/24	21/21	20/20	21/20	21/26
ET descriptor	4	40/107	20/36	14/45	13/38	08/28	19/51
	6	27/65	14/27	08/23	07/20	04/14	12/30
	8	19/50	11/23	05/15	04/12	03/09	08/22
	10	<u>13/39</u>	<u>07/20</u>	<u>03/11</u>	<u>03/09</u>	<u>02/06</u>	<u>06/17</u>
	12	06/33	03/17	01/08	01/07	01/05	03/14

Table 1. Comparison of the trajectory approximation accuracy of each descriptor. Both the observation and the prediction are curves fitted or approximated to each descriptor, reconstructed in the Euclidean space to fairly measure the error. The averaged L_2 distances between ground truth and reconstructed trajectory points are reported (Observation/prediction, Dim: Dimension of descriptor, Unit: *mm*). **Bold**: Best, underline: Second best.

Reconstruction in the Euclidean space. As a last step, we reconstruct the s refined coefficients $\hat{c}_{pred,n}^s$, and convert them into the full trajectory points. For a fair comparison with the existing methods, we transform the model output to the original Euclidean coordinate system using Eq. (5). Note that in the inference stage, we utilize U_{obs} and U_{pred} obtained from the training datasets for the transformation and reconstruction. The reconstructed trajectories are then used for loss functions and evaluation metrics.

3.4. Implementation Details

We incorporate our ET into six state-of-the-art pedestrian trajectory forecasting baselines [48, 55, 56, 59, 73, 102]. To validate the generality of our ET space, we simply change the size of the input and output channel dimensions of them. We use a loss function which is a linear combination of three terms to train our EigenTrajectory models. First, we measure a Frobenius norm between the refined coefficients $c_{pred,n}$ to regress the ground truth coefficient c_{pred} . Here, the winner-takes-all process [67] is chosen to back-propagate only to the closest trajectory anchors for the training:

$$\mathcal{L}_{coeff} = \frac{1}{N} \sum_{n=1}^N \min_{i \in [1, \dots, s]} \left\| \hat{c}_{pred,i}^s - c_{pred,i} \right\|. \quad (7)$$

Next, we calculate a Euclidean norm between the reconstructed path and the ground-truth raw coordinate, and then average them along a time axis as shown below:

$$\mathcal{L}_{dist} = \frac{1}{N(T_{pred} - T_{obs})} \sum_{n=1}^N \sum_{t=T_{obs}+1}^{T_{pred}} \min_{i \in [1, \dots, s]} \left\| \hat{B}_{i,t}^s - B_{i,t} \right\|. \quad (8)$$

Lastly, we impose an additional penalty for the last prediction frame. The ground-truth coefficient c_{pred} obtained by the low-rank approximation tends to be more careless for the endpoint because it is designed to minimize the error of the whole path. We thus penalize the endpoint coordinate so

	STGCNN [55]				EigenTrajectory - STGCNN					SGCN [73]				EigenTrajectory - SGCN				
	ADE↓	FDE↓	TCC↑	COL↓	ADE↓	FDE↓	TCC↑	COL↓	Gain↑	ADE↓	FDE↓	TCC↑	COL↓	ADE↓	FDE↓	TCC↑	COL↓	Gain↑
ETH	0.650	1.097	0.510	1.804	0.365	0.582	0.471	1.050	46.9%	0.567	0.997	0.545	1.686	0.360	0.565	0.438	1.160	43.3%
HOTEL	0.496	0.858	0.270	3.936	0.147	0.220	0.272	1.700	74.3%	0.308	0.533	0.295	2.523	0.131	0.210	0.269	1.752	60.7%
UNIV	0.441	0.798	0.637	9.691	0.246	0.427	0.751	8.548	46.5%	0.374	0.668	0.689	6.846	0.244	0.428	0.791	8.362	35.9%
ZARA1	0.341	0.532	0.710	2.528	0.217	0.393	0.808	1.396	26.1%	0.285	0.508	0.746	0.791	0.200	0.347	0.841	1.147	31.6%
ZARA2	0.305	0.482	0.394	7.150	0.168	0.290	0.645	6.212	39.7%	0.225	0.422	0.491	2.234	0.153	0.261	0.611	5.892	38.1%
AVG	0.447	0.753	0.504	5.022	0.229	0.383	0.589	3.781	49.2%	0.352	0.626	0.553	2.816	0.218	0.362	0.590	3.663	42.1%
SDD	20.76	33.18	0.471	0.679	8.11	13.35	0.578	0.541	59.8%	11.42	18.96	0.570	4.450	8.05	13.25	0.582	0.541	30.1%
GCS	14.72	23.87	0.698	3.921	8.45	14.49	0.834	3.354	39.3%	11.18	20.65	0.777	1.450	7.86	13.38	0.849	3.145	35.2%

	PECNet [48]				EigenTrajectory - PECNet					AgentFormer [102]				EigenTrajectory - AgentFormer				
	ADE↓	FDE↓	TCC↑	COL↓	ADE↓	FDE↓	TCC↑	COL↓	Gain↑	ADE↓	FDE↓	TCC↑	COL↓	ADE↓	FDE↓	TCC↑	COL↓	Gain↑
ETH	0.610	1.073	0.596	3.076	0.365	0.572	0.580	1.215	46.7%	0.456	0.797	0.594	1.105	0.362	0.568	0.487	1.105	28.7%
HOTEL	0.222	0.390	0.335	5.689	0.132	0.211	0.298	1.192	45.9%	0.142	0.222	0.363	0.579	0.147	0.222	0.267	1.866	-0.1%
UNIV	0.335	0.558	0.752	3.804	0.244	0.432	0.765	8.310	22.6%	0.254	0.454	0.775	4.636	0.244	0.430	0.747	8.416	5.4%
ZARA1	0.250	0.448	0.808	2.993	0.195	0.348	0.828	0.996	22.4%	0.176	0.303	0.839	0.235	0.216	0.397	0.808	1.416	-31.1%
ZARA2	0.186	0.332	0.596	4.910	0.143	0.250	0.628	2.817	24.8%	0.141	0.237	0.565	1.186	0.166	0.290	0.731	6.010	-22.1%
AVG	0.321	0.560	0.617	4.094	0.216	0.362	0.620	2.906	35.3%	0.234	0.403	0.627	1.548	0.227	0.381	0.608	3.763	5.3%
SDD	9.97	15.89	0.647	1.444	8.12	13.10	0.575	2.970	17.6%	8.68	14.92	0.608	0.379	8.10	13.43	0.590	0.562	10.0%
GCS	17.08	29.30	0.708	2.866	7.42	12.49	0.888	2.970	57.4%	10.18	16.91	0.840	2.319	8.41	14.56	0.889	3.263	13.9%

Table 2. Comparison between EigenTrajectory and the Euclidian space for four state-of-the-art multi-modal trajectory prediction models, Social-STGCNN [55], SGCN [73], PECNet [48] and AgentFormer [102]. The models are evaluated on the ETH [61], UCY [32], SDD [65] and GCS [100] datasets. Gain: performance improvement w.r.t FDE over the baseline models, Unit for ADE and FDE: meter, **Bold**: Best.

that the model can learn to infer the correct destination:

$$\mathcal{L}_{end} = \frac{1}{N} \sum_{n=1}^N \min_{i \in [1, \dots, s]} \left\| \hat{B}_{i, T_{pred}}^s - B_{i, T_{pred}} \right\|. \quad (9)$$

The final loss function is a linear combination of the three losses $\mathcal{L} = \mathcal{L}_{coeff} + \alpha \mathcal{L}_{dist} + \beta \mathcal{L}_{end}$. We empirically set α and β to 1. We train our EigenTrajectory models with the AdamW optimizer [45] with a batch size of 128 and a learning rate of 0.001 for 256 epochs. The training time takes about a day on a machine with an NVIDIA 3090 GPU.

4. Experiments

In this section, we conduct comprehensive experiments on public benchmark datasets to verify the efficiency of our $\mathbb{E}\mathbb{T}$ descriptor and the effectiveness of the $\mathbb{E}\mathbb{T}$ space for the trajectory prediction. We first describe our experimental setup briefly in Sec. 4.1. We then provide comparison results with other trajectory descriptors, various baselines, and state-of-the-art models in Sec. 4.2. Finally, we present the results of an extensive ablation study demonstrating the effect of each component of our method in Sec. 4.3.

4.1. Experimental Setup

Datasets. We conduct experiments on four public datasets: ETH [61], UCY [32], Stanford Drone Dataset (SDD) [65], and the Grand Central Station (GCS) [100] datasets to compare our EigenTrajectory with state-of-the-art baselines and to check the performance improvement for trajectory forecasting. The ETH and UCY datasets consist of pedestrian trajectories across five unique scenes (ETH, Hotel, Univ, ZARA1 and ZARA2) with 1,536 pedestrians recorded in the world coordinates. Following previous works [1, 19], we adopt the standard leave-one-out strategy for the training and evaluation.

SDD has 5,232 pedestrians in eight different university campus scenes at a top-down drone view. GCS shows a highly congested terminal scene with 12,684 pedestrians. We use a standard training and evaluation protocol [5, 19, 21, 55, 73] in which the first 3.2 seconds ($T_{obs} = 8$ frames) are observed, and succeeding 4.8 seconds ($T_{pred} - T_{obs} = 12$ frames) are used for the trajectory prediction.

Baseline models. We evaluate an efficiency and a generality of our EigenTrajectory by incorporating it into the following state-of-the-art baseline models: Social-STGCNN [55], SGCN [73], PECNet [48], AgentFormer [102], LB-EBM [59], and Social-Implicit [56]. For strictly fair comparison, we directly utilize authors' provided official source code for the vanilla baseline model designed in the Euclidean space. To validate the effectiveness of our $\mathbb{E}\mathbb{T}$ space, we modify the input and output shape of the baseline models to take and predict the coefficients of the $\mathbb{E}\mathbb{T}$ descriptor, instead of the direct use of the raw Euclidean coordinates like the conventional manner.

Evaluation metrics. To measure the prediction performance of the baseline models with our $\mathbb{E}\mathbb{T}$ space, we use four metrics: Average Displacement Error (ADE), Final Displacement Error (FDE), Temporal Correlation Coefficient (TCC) [80], and Collision rate (COL) [44]. The ADE and FDE compute the Euclidean distance between a predicted and a ground-truth trajectory and their end-points, respectively. The TCC measures the Pearson correlation coefficient of motion patterns between a predicted and ground-truth trajectory, and the COL calculates the percentage of collision cases between agents on the predicted path. We use both the ADE and FDE as accuracy measures, and both the TCC and COL as reliability measures. Following [19], we generate $s = 20$ samples, and then choose the best path to evaluate the multi-modal trajectory prediction performance.

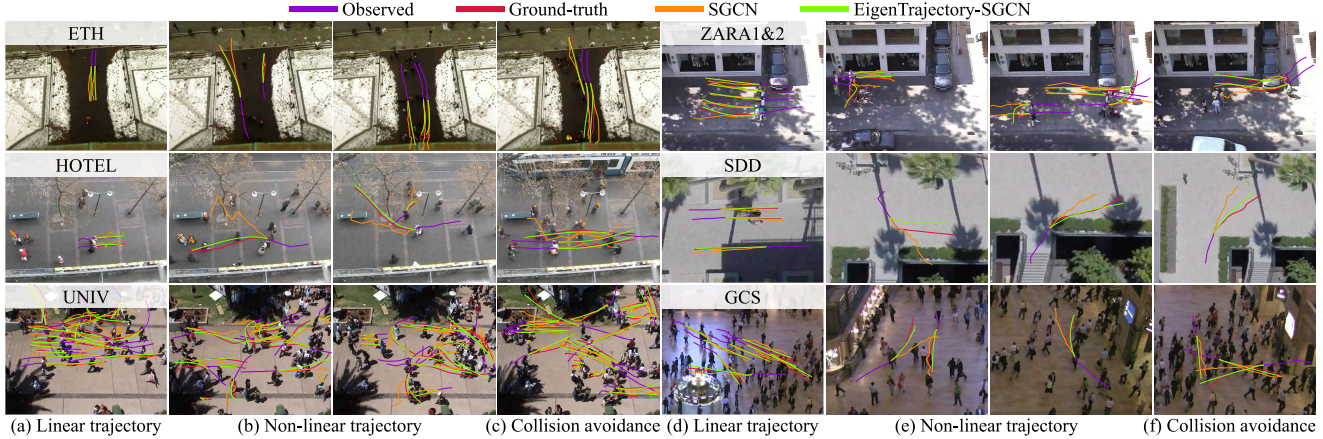


Figure 7. Examples of the prediction results from our EigenTrajectory predictor, compared to the baseline model in the conventional Euclidean space. To aid visualization, the SGCN [73] is used, and we report one trajectory with the best FDE of $s = 20$ samples each.

4.2. Evaluation Results

Evaluation of the trajectory descriptors. First, we check the efficiency of our low-rank approximation method using the $\mathbb{E}\mathbb{T}$ descriptor over other linear interpolation-based and parametric curve-based methods. The linear interpolation-based method predicts the coordinates of the first frame and the last frame of the predicted trajectory, and then equally divides and interpolates them to reconstruct the full frame points. Bézier curve and B-spline predict a set of control points and then compute the full coordinates through matrix multiplication with the Bernstein polynomial and piecewise polynomial basis, respectively. Following [22, 23, 25], the order of both curves is set to 5.

Tab. 1 shows that our descriptor is 87% more efficient than the results from the linear interpolation. This is because our $\mathbb{E}\mathbb{T}$ descriptor can cover linear and non-linear movements of pedestrians with only 2-dimensional and 4-dimensional representations, respectively. Similarly, the $\mathbb{E}\mathbb{T}$ descriptor shows superior low-rank approximation results even with the same $k = 12$ dimension as both the Bézier curve and B-spline. We choose $k = 6$ in our space design which shows similar or better approximation accuracy than that of the existing parametric curve.

Evaluation of the trajectory prediction. Next, we evaluate the prediction space with the public trajectory prediction benchmarks. As reported in Tab. 2, our EigenTrajectory framework achieves consistent performance improvements with all the baseline models. In particular, the prediction accuracy improves significantly, up to 74.3%. These results demonstrate that the trajectory forecasting models can benefit from our $\mathbb{E}\mathbb{T}$ space rather than dealing with raw data, and make the prediction task tractable when using trajectory anchors. In addition, most of the trajectory reliability metrics also achieve improvements by predicting accurate and stable movement patterns, as visualized in Fig. 7.

However, there are some limitations with our EigenTrajectory. The TCC metric is slightly worse in the ETH and hotel scenes because we remove the noisy motions through the truncation, but there are a lot of wobbling people in both scenes. In addition, due to the low-rank approximation, the macroscopic movements are leveraged, but the microscopic movements are removed. This is why there are more collisions in very complex UNIV and GCS scenes. In this case, it is important to find a good trade-off between the accuracy and reliability by adjusting k .

Comparison with the state-of-the-art models. We compare our EigenTrajectory models with the state-of-the-art models. As shown in Tab. 3, our EigenTrajectory achieves a better accuracy by taking full advantage of our data-driven trajectory descriptor than those of the previous models. Furthermore, our EigenTrajectory shows a higher improvement than other generalized approaches such as NCE [44], Causal [10], GP-Graph [4], and NPSN [5]. While all of those approaches adhere to the conventional Euclidean space for the input and output, the introduction of our efficient descriptor can significantly improve the network ability.

4.3. Ablation Studies

Effectiveness of each component. We validate the effectiveness of each module optimized in our $\mathbb{E}\mathbb{T}$ space. In Tab. 4, we only replace the input of the baseline models with the $\mathbb{E}\mathbb{T}$ space and keep the output space as it is. In this case, we observe the marginal improvements. Next, we replace both the input and output space with the $\mathbb{E}\mathbb{T}$ space, and achieve significant performance improvements. Here, we observe that the existing baseline network is struggling to predict all the points of the future path. Lastly, the best performance is obtained when our trajectory anchor is incorporated into the baseline model.

Trajectory anchors. Next, we demonstrate the effective-

Model	ETH	HOTEL	UNIV	ZARA1	ZARA2	AVG
Social-LSTM [1]	1.09/2.35	0.79/1.76	0.67/1.40	0.47/1.00	0.56/1.17	0.72/1.54
Social-GAN [19]	0.87/1.62	0.67/1.37	0.76/1.52	0.35/0.68	0.42/0.84	0.61/1.21
STGAT [21]	0.65/1.12	0.35/0.66	0.52/1.10	0.34/0.69	0.29/0.60	0.43/0.83
Social-STGCNN [55]	0.64/1.11	0.49/0.85	0.44/0.79	0.34/0.53	0.30/0.48	0.44/0.75
PECNet [†] [48]	0.61/1.07	0.22/0.39	0.34/0.56	0.25/0.45	0.19/0.33	0.32/0.56
Trajectron++ [†] [69]	0.61/1.03	0.20/0.28	0.30/0.55	0.24/0.41	0.18/0.32	0.31/0.52
SGCN [73]	0.57/1.00	0.31/0.53	0.37/0.67	0.29/0.51	0.22/0.42	0.35/0.63
LB-EBM [†] [59]	0.60/1.06	0.21/0.38	0.28/0.54	0.21/0.39	0.15/0.30	0.29/0.53
AgentFormer [†] [102]	0.46/0.80	0.14/0.22	0.25/0.45	0.18/ 0.30	0.14/ 0.24	0.23/0.40
ExpertTraj [106]	<u>0.37/0.65</u>	0.11/0.15	0.20/0.44	0.15/0.31	0.12/0.26	0.19/0.36
MemoNet [93]	0.40/0.61	0.11/0.17	0.24/0.43	0.18/0.32	0.14/ 0.24	<u>0.21/0.35</u>
Social-Implicit [56]	0.66/1.44	0.20/0.36	0.31/0.60	0.25/0.50	0.22/0.43	0.33/0.67
MID [18]	0.39/0.66	0.13/0.22	<u>0.22/0.45</u>	<u>0.17/0.30</u>	<u>0.13/0.27</u>	<u>0.21/0.38</u>
NCE-STGCNN [44]	0.67/1.22	0.44/0.68	0.47/0.88	0.33/0.52	0.29/0.48	0.44/0.76
Causal-STGCNN [10]	0.64/1.00	0.38/0.45	0.49/0.81	0.34/0.53	0.32/0.49	0.43/0.66
GP-Graph-STGCNN [4]	0.48/0.77	0.24/0.40	0.29/0.47	0.24/0.40	0.23/0.40	0.29/0.49
NPSN-STGCNN [5]	0.44/0.65	0.21/0.34	0.28/0.44	0.25/0.43	0.22/0.38	0.28/0.45
EigenTrajectory-STGCNN	0.36/0.58	0.15/0.22	0.25/0.43	0.22/0.39	0.17/0.29	0.23/0.38
EigenTrajectory-AgentFormer	0.36/0.57	0.15/0.22	0.24/0.43	0.22/0.40	0.17/0.29	0.23/0.38
EigenTrajectory-Implicit	0.36/0.57	0.13/0.21	0.24/0.43	0.21/0.37	0.15/0.26	0.22/0.37
EigenTrajectory-SGCN	0.36/0.57	0.13/0.21	0.24/0.43	0.20/0.35	0.15/0.26	0.22/0.36
EigenTrajectory-PECNet	0.36/0.57	0.13/0.21	0.24/0.43	0.19/0.35	0.14/ <u>0.25</u>	0.22/0.36
EigenTrajectory-LB-EBM	0.36/0.53	<u>0.12/0.19</u>	0.24/0.43	0.19/0.33	0.14/ 0.24	<u>0.21/0.34</u>

Table 3. Comparison EigenTrajectory methods with other state-of-the-art stochastic models (ADE/FDE, Unit: meter). †: Issues raised in the authors’ GitHubs are fixed, **Bold**: Best, Underline: 2nd Best.

Model	ETH	HOTEL	UNIV	ZARA1	ZARA2	AVG
Baseline	0.57 / 1.00	0.31 / 0.53	0.37 / 0.67	0.29 / 0.51	0.23 / 0.42	0.35 / 0.63
+ Observed trajectory	0.50 / 0.79	0.21 / 0.34	<u>0.33/0.62</u>	0.27 / 0.48	0.26 / 0.44	0.31 / 0.53
+ Predicted trajectory	<u>0.48 / 0.70</u>	<u>0.17 / 0.27</u>	<u>0.33 / 0.61</u>	<u>0.26 / 0.47</u>	<u>0.22 / 0.40</u>	<u>0.29 / 0.49</u>
+ Trajectory anchor	0.36 / 0.57	0.13 / 0.21	0.24 / 0.43	0.20 / 0.35	0.15 / 0.26	0.22 / 0.36
Euclidean anchor	0.47 / 0.76	0.18 / 0.30	0.36 / 0.69	0.35 / 0.67	0.25 / 0.48	0.32 / 0.58
$\mathbb{E}\mathbb{T}$ anchor	0.36 / 0.57	0.14 / 0.23	0.24 / 0.43	0.22 / 0.40	0.17 / 0.29	0.23 / 0.38

Table 4. Ablation study on trajectory anchor generated in different spaces on SGCN [73] (ADE/FDE, Unit: meter). **Bold**: Best.

ness of our $\mathbb{E}\mathbb{T}$ space for trajectory anchor generation. Tab. 4 presents the results of an ablation study on the clustering space to obtain the data-driven trajectory anchors. The accuracy of the anchors from the $\mathbb{E}\mathbb{T}$ space is much better compared to the Euclidean space. Additionally, the results are better than those of most state-of-the-art works, only with anchors without any refinement method. We think that the trajectory anchors offer an efficient initial trajectory candidate from the initial data, and better anchors can be estimated when clustering with lower dimensions.

Non-linear trajectories. We evaluate how well the $\mathbb{E}\mathbb{T}$ descriptor represent and predict non-linear trajectories. Following [19], we evaluate a case in which the linear approximation error of the future path is more than 0.02m. In Tab. 5, our $\mathbb{E}\mathbb{T}$ space has a smaller performance drop, compared to the conventional Euclidean space. The prediction ability of a linear path is already saturated, so it is important to design a model to handle non-linear paths well. Since most of the paths in the dataset are straight, the output trajectories tend to be smooth during training. Nevertheless, our $\mathbb{E}\mathbb{T}$ space shows a robust performance for these non-linear trajectories because our $\mathbb{E}\mathbb{T}$ descriptor can explicitly represent non-linearity with the combinations of $\mathbb{E}\mathbb{T}$ coefficients.

Trajectory perturbation. Lastly, we examine the robustness of the prediction to input noise. To this end, we measure

Space	Type	ETH	HOTEL	UNIV	ZARA1	ZARA2	AVG
Euclidean	All	0.57 / 1.00	0.31 / 0.53	0.37 / 0.67	0.29 / 0.51	0.23 / 0.42	0.35 / 0.63
	NL	0.65 / 1.16	0.46 / 0.82	0.49 / 0.92	0.39 / 0.78	0.50 / 1.03	0.50 / 0.94
	Diff.	-0.08/-0.17	-0.15/-0.29	-0.12/-0.25	-0.11/-0.28	-0.27/-0.61	-0.15/-0.32
$\mathbb{E}\mathbb{T}$	All	0.36 / 0.57	0.13 / 0.21	0.24 / 0.43	0.20 / 0.35	0.15 / 0.26	0.22 / 0.36
	NL	0.41 / 0.65	0.19 / 0.31	0.32 / 0.55	0.25 / 0.43	0.32 / 0.53	0.30 / 0.49
	Diff.	-0.05/-0.08	-0.06/-0.10	-0.08/-0.12	-0.05/-0.08	-0.17/-0.27	-0.08/-0.13

Table 5. Ablation study on non-linear trajectories on SGCN [73]. NL: evaluations only with non-linear trajectories, Diff: Performance difference (ADE/FDE, Unit: meter). **Bold**: Best.

Space	Noise	ETH	HOTEL	UNIV	ZARA1	ZARA2	AVG
Euclidean	-	0.57 / 1.00	0.31 / 0.53	0.37 / 0.67	0.29 / 0.51	0.23 / 0.42	0.35 / 0.63
	0.02	0.58 / 1.01	0.35 / 0.59	0.39 / 0.69	0.30 / 0.54	0.27 / 0.48	0.38 / 0.66
	0.05	0.63 / 1.09	0.42 / 0.69	0.45 / 0.79	0.37 / 0.65	0.35 / 0.61	0.44 / 0.77
	0.10	0.75 / 1.28	0.60 / 1.00	0.61 / 1.05	0.56 / 0.98	0.53 / 0.89	0.61 / 1.04
	Diff.	-0.18/-0.29	-0.29/-0.46	-0.24/-0.38	-0.27/-0.47	-0.31/-0.47	-0.26/-0.42
$\mathbb{E}\mathbb{T}$	-	0.36 / 0.57	0.13 / 0.21	0.24 / 0.43	0.20 / 0.35	0.15 / 0.26	0.22 / 0.36
	0.02	0.37 / 0.58	0.14 / 0.22	0.25 / 0.44	0.21 / 0.37	0.17 / 0.28	0.23 / 0.38
	0.05	0.41 / 0.62	0.17 / 0.25	0.29 / 0.48	0.26 / 0.42	0.20 / 0.32	0.26 / 0.41
	0.10	0.47 / 0.68	0.25 / 0.33	0.36 / 0.55	0.36 / 0.53	0.27 / 0.39	0.34 / 0.50
	Diff.	-0.11/-0.11	-0.11/-0.12	-0.11/-0.12	-0.16/-0.18	-0.12/-0.13	-0.12/-0.13

Table 6. Ablation study on trajectory perturbation result of SGCN [73]. Noise: Standard deviation of Gaussian noise, Diff: Performance difference (ADE/FDE, Unit: meter). **Bold**: Best.

the change in prediction accuracy after adding a little Gaussian noise. In Tab. 6, our EigenTrajectory has a marginal performance drop when noise exists, compared to the baselines. In our $\mathbb{E}\mathbb{T}$ space, because the principal motion pattern components are only left through the rank- k approximation, the negative effect of noise can be mitigated.

5. Conclusion

In this work, we introduce a low-rank approximation-based trajectory descriptor trained in a data-driven manner to make a low-dimensional representation of pedestrian paths. While the existing architectures working in the Euclidean space suffer from the curse of dimensionality, we define a new operating space, the $\mathbb{E}\mathbb{T}$ space, that unfolds highly-conjugated feature relations. We then cluster the coefficients of the $\mathbb{E}\mathbb{T}$ descriptor coefficients on the training set, and utilize them as trajectory anchors. The architectures learn to refine this data-driven anchor to infer structurally-diverse trajectories that can cover all travelable paths. A variety of experiments demonstrate that it provides great applicability and stability, which can be applied to off-the-shelf trajectory forecasting models with consistent performance improvements on most public datasets.

Acknowledgement This work is in part supported by the Institute of Information & communications Technology Planning & Evaluation (IITP) (No.2021-0-02068, Artificial Intelligence Innovation Hub), the National Research Foundation of Korea (NRF) (No.2020R1C1C1012635) grant funded by the Korea government (MSIT), GIST-MIT Research Collaboration grant funded by the GIST in 2023, and the Ministry of Trade, Industry and Energy (MOTIE) and Korea Institute of Advancement of Technology (KIAT) through the International Cooperative R&D program: P0019782, Embedded AI Based fully autonomous driving software and Maas technology development.

References

- [1] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1, 2, 6, 8
- [2] Inhwon Bae and Hae-Gon Jeon. Disentangled multi-relational graph convolutional network for pedestrian trajectory prediction. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2021. 1, 2
- [3] Inhwon Bae and Hae-Gon Jeon. A set of control points conditioned pedestrian trajectory prediction. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2023. 2
- [4] Inhwon Bae, Jin-Hwi Park, and Hae-Gon Jeon. Learning pedestrian group representations for multi-modal trajectory prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 2, 7, 8
- [5] Inhwon Bae, Jin-Hwi Park, and Hae-Gon Jeon. Non-probability sampling network for stochastic human trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 6, 7, 8
- [6] Richard Ernest Bellman. *Dynamic programming*. Princeton University Press, USA, 1957. 4
- [7] Apratim Bhattacharyya, Michael Hanselmann, Mario Fritz, Bernt Schiele, and Christoph-Nikolas Straehle. Conditional flow variational autoencoders for structured sequence prediction. *arXiv preprint arXiv:1908.09008*, 2020. 2
- [8] Niccoló Bisagno, Bo Zhang, and Nicola Conci. Group lstm: Group trajectory prediction in crowded scenarios. In *Proceedings of the European Conference on Computer Vision Workshop (ECCVW)*, 2018. 2
- [9] Yuning Chai, Benjamin Sapp, Mayank Bansal, and Dragomir Anguelov. Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2019. 4
- [10] Guangyi Chen, Junlong Li, Jiwen Lu, and Jie Zhou. Human trajectory prediction via counterfactual analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 1, 7, 8
- [11] Guangyi Chen, Junlong Li, Nuoxing Zhou, Liangliang Ren, and Jiwen Lu. Personalized trajectory prediction via distribution discrimination. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2
- [12] Patrick Dendorfer, Sven Elflein, and Laura Leal-Taixé. Mgan: A multi-generator model preventing out-of-distribution samples in pedestrian trajectory prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2
- [13] Patrick Dendorfer, Aljosa Osep, and Laura Leal-Taixé. Goalgan: Multimodal trajectory prediction based on goal position estimation. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, 2020. 2
- [14] Nachiket Deo and Mohan M. Trivedi. Trajectory forecasts in unknown environments conditioned on grid-based plans. *arXiv preprint arXiv:2001.00735*, 2020. 2
- [15] Tharindu Fernando, Simon Denman, Sridha Sridharan, and Clinton Fookes. Gd-gan: Generative adversarial networks for trajectory prediction and group detection in crowds. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, 2018. 1
- [16] Tharindu Fernando, Simon Denman, Sridha Sridharan, and Clinton Fookes. Soft+ hardwired attention: An lstm framework for human trajectory prediction and abnormal event detection. *Neural Networks*, 108:466–478, 2018. 2
- [17] Matthias Fey, Jan Eric Lenssen, Frank Weichert, and Heinrich Müller. Splinecnn: Fast geometric deep learning with continuous b-spline kernels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1
- [18] Tianpei Gu, Guangyi Chen, Junlong Li, Chunze Lin, Yongming Rao, Jie Zhou, and Jiwen Lu. Stochastic trajectory prediction via motion indeterminacy diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 8
- [19] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social gan: Socially acceptable trajectories with generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2, 6, 8
- [20] Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. *Physical review E*, 51(5):4282, 1995. 2
- [21] 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 (ICCV)*, 2019. 2, 6, 8
- [22] Ronny Hug, Stefan Becker, Wolfgang Hübner, Michael Arens, and Jürgen Beyerer. Bézier curve gaussian processes. *arXiv preprint arXiv:2205.01754*, 2022. 1, 2, 5, 7
- [23] Ronny Hug, Wolfgang Hübner, and Michael Arens. Introducing probabilistic bézier curves for n-step sequence prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2020. 1, 2, 5, 7
- [24] 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 (ICCV)*, 2019. 2
- [25] Mohammad Sadegh Jazayeri and Arash Jahangiri. Utilizing b-spline curves and neural networks for vehicle trajectory prediction in an inverse reinforcement learning framework. *Journal of Sensor and Actuator Networks*, 2022. 1, 2, 5, 7
- [26] Dongkwon Jin, Wonhui Park, Seong-Gyun Jeong, Heeyeon Kwon, and Chang-Su Kim. Eigenlanes: Data-driven lane descriptors for structurally diverse lanes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [27] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *Proceedings of*

- the International Conference on Learning Representations (ICLR)*, 2017. 2
- [28] Vineet Kosaraju, Amir Sadeghian, Roberto Martín-Martín, Ian Reid, Hamid Rezatofighi, and Silvio Savarese. Socialbigat: Multimodal trajectory forecasting using bicycle-gan and graph attention networks. In *Proceedings of the Neural Information Processing Systems (NeurIPS)*, 2019. 1, 2
- [29] Parth Kothari, Brian Siffringer, and Alexandre Alahi. Interpretable social anchors for human trajectory forecasting in crowds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 4
- [30] Mihee Lee, Samuel S. Sohn, Seonghyeon Moon, Sejong Yoon, Mubbasir Kapadia, and Vladimir Pavlovic. Muse-vae: Multi-scale vae for environment-aware long term trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [31] Namhoon Lee, Wongun Choi, Paul Vernaza, Christopher B. Choy, Philip H. S. Torr, and Manmohan Chandraker. Desire: Distant future prediction in dynamic scenes with interacting agents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [32] Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. Crowds by example. *Computer Graphics Forum*, 26(3):655–664, 2007. 6
- [33] Anat Levin, Dani Lischinski, and Yair Weiss. A closed-form solution to natural image matting. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2007. 3
- [34] Anat Levin, Alex Rav-Acha, and Dani Lischinski. Spectral matting. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2008. 3
- [35] Jiachen Li, Hengbo Ma, and Masayoshi Tomizuka. Conditional generative neural system for probabilistic trajectory prediction. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019. 1, 2
- [36] Jiachen Li, Fan Yang, Masayoshi Tomizuka, and Chihoh Choi. Evolvegraph: Multi-agent trajectory prediction with dynamic relational reasoning. In *Proceedings of the Neural Information Processing Systems (NeurIPS)*, 2020. 1, 2
- [37] Shijie Li, Yanying Zhou, Jinhui Yi, and Juergen Gall. Spatial-temporal consistency network for low-latency trajectory forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 1, 2
- [38] Yuke Li. Which way are you going? imitative decision learning for path forecasting in dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [39] Junwei Liang, Lu Jiang, and Alexander Hauptmann. Simaug: Learning robust representations from simulation for trajectory prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2
- [40] Junwei Liang, Lu Jiang, Kevin Murphy, Ting Yu, and Alexander Hauptmann. The garden of forking paths: Towards multi-future trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2
- [41] Junwei Liang, Lu Jiang, Juan Carlos Niebles, Alexander G Hauptmann, and Li Fei-Fei. Peeking into the future: Predicting future person activities and locations in videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [42] Rongqin Liang, Yuanman Li, Xia Li, Yi Tang, Jiantao Zhou, and Wenbin Zou. Temporal pyramid network for pedestrian trajectory prediction with multi-supervision. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2021. 2
- [43] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016. 4
- [44] Yuejiang Liu, Qi Yan, and Alexandre Alahi. Social nce: Contrastive learning of socially-aware motion representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 6, 7, 8
- [45] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018. 6
- [46] Yuexin Ma, Xinge Zhu, Xinjing Cheng, Ruigang Yang, Jiming Liu, and Dinesh Manocha. Autotrajectory: Label-free trajectory extraction and prediction from videos using dynamic points. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2
- [47] Karttikeya Mangalam, Yang An, Harshayu Girase, and Jitendra Malik. From goals, waypoints & paths to long term human trajectory forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2
- [48] Karttikeya Mangalam, Harshayu Girase, Shreyas Agarwal, Kuan-Hui Lee, Ehsan Adeli, Jitendra Malik, and Adrien Gaidon. It is not the journey but the destination: Endpoint conditioned trajectory prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 5, 6, 8
- [49] Huynh Manh and Gita Alaghand. Scene- lstm: A model for human trajectory prediction. *arXiv preprint arXiv:1808.04018*, 2018. 2
- [50] Francesco Marchetti, Federico Becattini, Lorenzo Seidenari, and Alberto Del Bimbo. Mantra: Memory augmented networks for multiple trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [51] Francesco Marchetti, Federico Becattini, Lorenzo Seidenari, and Alberto Del Bimbo. Multiple trajectory prediction of moving agents with memory augmented networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2020. 2
- [52] Francesco Marchetti, Federico Becattini, Lorenzo Seidenari, and Alberto Del Bimbo. Smemo: Social memory for trajectory forecasting. *arXiv preprint arXiv:2203.12446*, 2022. 2
- [53] Ramin Mehran, Alexis Oyama, and Mubarak Shah. Abnormal crowd behavior detection using social force model.

- In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 2
- [54] M. Minsky and S. Papert. *Perceptrons*. MIT Press, USA, 1969. 4
- [55] Abdullallah 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 (CVPR)*, 2020. 1, 2, 5, 6, 8
- [56] Abdullallah Mohamed, Deyao Zhu, Warren Vu, Mohamed Elhoseiny, and Christian Claudel. Social-implicit: Rethinking trajectory prediction evaluation and the effectiveness of implicit maximum likelihood estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 5, 6, 8
- [57] Alessio Monti, Angelo Porrello, Simone Calderara, Pasquale Coscia, Lamberto Ballan, and Rita Cucchiara. How many observations are enough? knowledge distillation for trajectory forecasting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [58] Ingrid Navarro and Jean Oh. Social-patternn: Socially-aware trajectory prediction guided by motion patterns. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022. 2
- [59] Bo Pang, Tianyang Zhao, Xu Xie, and Ying Nian Wu. Trajectory prediction with latent belief energy-based model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 5, 6, 8
- [60] Wonhui Park, Dongkwon Jin, and Chang-Su Kim. Eigencontours: Novel contour descriptors based on low-rank approximation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [61] Stefano Pellegrini, Andreas Ess, Konrad Schindler, and Luc Van Gool. You'll never walk alone: Modeling social behavior for multi-target tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2009. 2, 6
- [62] Weilong Peng, Zhiyong Feng, Chao Xu, and Yong Su. Parametric t-spline face morphable model for detailed fitting in shape subspace. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1
- [63] Mark Pfeiffer, Giuseppe Paolo, Hannes Sommer, Juan I. Nieto, Roland Y. Siegwart, and César Cadena. A data-driven model for interaction-aware pedestrian motion prediction in object cluttered environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2018. 2
- [64] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 4
- [65] Alexandre Robicquet, Amir Sadeghian, Alexandre Alahi, and Silvio Savarese. Learning social etiquette: Human trajectory understanding in crowded scenes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016. 6
- [66] Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M Kitani, Dariu M Gavrila, and Kai O Arras. Human motion trajectory prediction: A survey. *International Journal of Robotics Research (IJRR)*, 2020. 1
- [67] Christian Rupprecht, Iro Laina, Robert DiPietro, Maximilian Baust, Federico Tombari, Nassir Navab, and Gregory D Hager. Learning in an uncertain world: Representing ambiguity through multiple hypotheses. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017. 5
- [68] Amir Sadeghian, Vineet Kosaraju, Ali Sadeghian, Noriaki Hirose, Hamid Rezaatofghi, 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 (CVPR)*, 2019. 2
- [69] Tim Salzman, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 8
- [70] Nasim Shafiee, Taskin Padir, and Ehsan Elhamifar. Introvert: Human trajectory prediction via conditional 3d attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [71] Gopal Sharma, Difan Liu, Subhransu Maji, Evangelos Kalogerakis, Siddhartha Chaudhuri, and Radomír Měch. Parsenet: A parametric surface fitting network for 3d point clouds. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1
- [72] Liushuai Shi, Le Wang, Chengjiang Long, Sanping Zhou, Fang Zheng, Nanning Zheng, and Gang Hua. Social interpretable tree for pedestrian trajectory prediction. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2022. 2
- [73] Liushuai Shi, Le Wang, Chengjiang Long, Sanping Zhou, Mo Zhou, Zhenxing Niu, and Gang Hua. Sgcn: Sparse graph convolution network for pedestrian trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 5, 6, 7, 8
- [74] Xiaodan Shi, Xiaowei Shao, Zipei Fan, Renhe Jiang, Haoran Zhang, Zhiling Guo, Guangming Wu, Wei Yuan, and Ryosuke Shibasaki. Multimodal interaction-aware trajectory prediction in crowded space. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2020. 1, 2
- [75] Xiaodan Shi, Xiaowei Shao, Guangming Wu, Haoran Zhang, Zhiling Guo, Renhe Jiang, and Ryosuke Shibasaki. Social-dpf: Socially acceptable distribution prediction of futures. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2021. 2
- [76] Jisu Shin, Seunghyun Shin, and Hae-Gon Jeon. Task-specific scene structure representations. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2023. 3
- [77] Hao Sun, Zhiqun Zhao, and Zhihai He. Reciprocal learning networks for human trajectory prediction. In *Proceedings of*

- the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2
- [78] Jianhua Sun, Qinhong Jiang, and Cewu Lu. Recursive social behavior graph for trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [79] Jianhua Sun, Yuxuan Li, Hao-Shu Fang, and Cewu Lu. Three steps to multimodal trajectory prediction: Modality clustering, classification and synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2
- [80] Chaofan Tao, Qinhong Jiang, and Lixin Duan. Dynamic and static context-aware lstm for multi-agent motion prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2, 6
- [81] Matthew A Turk and Alex P Pentland. Face recognition using eigenfaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1991. 3
- [82] Daksh Varshneya and G. Srinivasaraghavan. Human trajectory prediction using spatially aware deep attention models. *arXiv preprint arXiv:1705.09436*, 2017. 2
- [83] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018. 2
- [84] Anirudh Vemula, Katharina Muelling, and Jean Oh. Social attention: Modeling attention in human crowds. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2018. 2
- [85] Chuhua Wang, Yuchen Wang, Mingze Xu, and David J Crandall. Stepwise goal-driven networks for trajectory prediction. *IEEE Robotics and Automation Letters (RA-L)*, 2022. 2
- [86] Xiaogang Wang, Yuelang Xu, Kai Xu, Andrea Tagliasacchi, Bin Zhou, Ali Mahdavi-Amiri, and Hao Zhang. Pie-net: Parametric inference of point cloud edges. *Proceedings of the Neural Information Processing Systems (NeurIPS)*, 2020. 1
- [87] Yu Wang, Yi Niu, Peiyong Duan, Jianwei Lin, and Yuanjie Zheng. Deep propagation based image matting. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI)*, 2018. 3
- [88] Song Wen, Hao Wang, and Dimitris Metaxas. Social ode: Multi-agent trajectory forecasting with neural ordinary differential equations. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 2
- [89] Karl DD Willis, Pradeep Kumar Jayaraman, Joseph G Lambourne, Hang Chu, and Yewen Pu. Engineering sketch generation for computer-aided design. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1
- [90] Conghao Wong, Beihao Xia, Ziming Hong, Qinmu Peng, Wei Yuan, Qiong Cao, Yibo Yang, and Xinge You. View vertically: A hierarchical network for trajectory prediction via fourier spectrums. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 2
- [91] Rundi Wu, Chang Xiao, and Changxi Zheng. Deepcad: A deep generative network for computer-aided design models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 1
- [92] Chenxin Xu, Maosen Li, Zhenyang Ni, Ya Zhang, and Siheng Chen. Groupnet: Multiscale hypergraph neural networks for trajectory prediction with relational reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [93] Chenxin Xu, Weibo Mao, Wenjun Zhang, and Siheng Chen. Remember intentions: Retrospective-memory-based trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 4, 8
- [94] Pei Xu, Jean-Bernard Hayet, and Ioannis Karamouzas. Socialvae: Human trajectory prediction using timewise latents. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 2
- [95] Yi Xu, Lichen Wang, Yizhou Wang, and Yun Fu. Adaptive trajectory prediction via transferable gnn. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [96] Yi Xu, Jing Yang, and Shaoyi Du. Cf-lstm: Cascaded feature-based long short-term networks for predicting pedestrian trajectory. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2020. 2
- [97] Hao Xue, Du Q Huynh, and Mark Reynolds. Ss-lstm: A hierarchical lstm model for pedestrian trajectory prediction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2018. 2
- [98] Kota Yamaguchi, Alexander C Berg, Luis E Ortiz, and Tamara L Berg. Who are you with and where are you going? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011. 2
- [99] Yu Yao, Ella Atkins, Matthew Johnson-Roberson, Ram Vasudevan, and Xiaoxiao Du. Bitrap: Bi-directional pedestrian trajectory prediction with multi-modal goal estimation. *IEEE Robotics and Automation Letters (RA-L)*, 2021. 2
- [100] Shuai Yi, Hongsheng Li, and Xiaogang Wang. Understanding pedestrian behaviors from stationary crowd groups. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 6
- [101] Cunjun Yu, Xiao Ma, Jiawei Ren, Haiyu Zhao, and Shuai Yi. Spatio-temporal graph transformer networks for pedestrian trajectory prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 5
- [102] Ye Yuan, Xinshuo Weng, Yanglan Ou, and Kris Kitani. Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2, 5, 6, 8
- [103] Jiangbei Yue, Dinesh Manocha, and He Wang. Human trajectory prediction via neural social physics. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 2
- [104] 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 (CVPR)*, 2019. 2

- [105] Hang Zhao, Jiyang Gao, Tian Lan, Chen Sun, Benjamin Sapp, Balakrishnan Varadarajan, Yue Shen, Yi Shen, Yuning Chai, Cordelia Schmid, Congcong Li, and Dragomir Anguelov. Tnt: Target-driven trajectory prediction. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2020. 4
- [106] He Zhao and Richard P. Wildes. Where are you heading? dynamic trajectory prediction with expert goal examples. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2, 8
- [107] Tianyang Zhao, Yifei Xu, Mathew Monfort, Wongun Choi, Chris Baker, Yibiao Zhao, Yizhou Wang, and Ying Nian Wu. Multi-agent tensor fusion for contextual trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [108] Bolei Zhou, Xiaoou Tang, and Xiaogang Wang. Learning collective crowd behaviors with dynamic pedestrian-agents. *International Journal on Computer Vision (IJCV)*, 2015. 1
- [109] Bolei Zhou, Xiaogang Wang, and Xiaoou Tang. Understanding collective crowd behaviors: Learning a mixture model of dynamic pedestrian-agents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012. 1