

# Personalized Trajectory Prediction via Distribution Discrimination

Guangyi Chen<sup>1,2,3</sup>, Junlong Li<sup>1,2,3</sup>, Nuoxing Zhou<sup>1,2,3</sup>, Liangliang Ren<sup>1,2,3</sup>, Jiwen Lu<sup>1,2,3,\*</sup>

<sup>1</sup>Department of Automation, Tsinghua University, China

<sup>2</sup>State Key Lab of Intelligent Technologies and Systems, China

<sup>3</sup>Beijing National Research Center for Information Science and Technology, China

chen-gy16@mails.tsinghua.edu.cn; {ljlong.leo, nuoxingzhou}@gmail.com;

renliangliang@cvte.com; lujiwen@tsinghua.edu.cn

## Abstract

Trajectory prediction is confronted with the dilemma to capture the multi-modal nature of future dynamics with both diversity and accuracy. In this paper, we present a distribution discrimination (DisDis) method to predict personalized motion patterns by distinguishing the potential distributions. Motivated by that the motion pattern of each person is personalized due to his/her habit, our DisDis learns the latent distribution to represent different motion patterns and optimize it by the contrastive discrimination. This distribution discrimination encourages latent distributions to be more discriminative. Our method can be integrated with existing multi-modal stochastic predictive models as a plug-and-play module to learn the more discriminative latent distribution. To evaluate the latent distribution, we further propose a new metric, probability cumulative minimum distance (PCMD) curve, which cumulatively calculates the minimum distance on the sorted probabilities. Experimental results on the ETH and UCY datasets show the effectiveness of our method.<sup>1</sup>

## 1. Introduction

Human trajectory prediction aims at forecasting future pedestrian trajectories in complex dynamic environments with the observed history behaviors [1, 12, 18, 31]. It plays a critical role in developing safe human-interactive autonomous systems such as self-driving vehicles [28], social robotics [4], and intelligent surveillance applications [5, 3, 6].

Forecasting the trajectories of humans faces the serious challenges to model the indeterminacy of human behavior. With the same environment and social interactions, the



Figure 1. Illustration of the motivation with an example of a family of three. The family members (father, mother, and child) with the similar history trajectories might have different motion patterns due to their habits. (Best viewed in color)

human might take different plausible actions. For example, to avoid collisions with other pedestrians, one can choose to stop for a moment or speed up. To generate plausible multi-modal future states instead of a deterministic trajectory, many methods [12, 18, 32, 50] employ the Generative Adversarial Networks [11] (GAN) to spread the distribution of the prediction and cover the space of possible paths, while some methods [19, 16, 34] apply the Conditional Variational Auto-encoder (CVAE) [35] to explicitly encode the multi-modal distribution with a latent variable.

With the increase of motion diversity and complexity, stochastic prediction with the prior Gaussian distribution is insufficient to cover the wide spectrum of future possibilities. As shown in Figure 1, considering a family of three with different habits, the motion patterns are always different even given the similar trajectories. For example, the child tends to go straight to school, while the father might turn right for the bus. Although some methods [9, 16, 34] attempt to use different latent distributions for different persons, discriminative ability of these latent distributions are limited. The model can't always generate appropriate

\*Corresponding author

<sup>1</sup>Code and a video demo are available at <https://github.com/CHENGY12/DisDis>

outputs via a most-likely manner. A piece of obvious evidence is the large performance gap in existing generative methods[9, 34] between most likely single output and the posterior selected best output. The poor performance of the most likely single output indicates that the learned distribution cannot represent the personalized tendency. The self-driving system samples possible trajectories from learned distribution to make decisions. A good latent distribution is of critical importance for reducing the cost of sampling with no performance drop.

To overcome the above issues, in this paper, we propose a distribution discrimination method (DisDis) to learn the personalized multi-modal distribution, where the behavior pattern of each person is modeled as the latent distribution. Different from other unified prior distribution based generative prediction methods, we learn a distribution discriminator to distinguish the potential behavior patterns. Without any extra supervisory signals, we optimize the distribution discriminator in a self-supervised manner, which encourages the latent variable distributions of the same motion pattern to be similar while pushing the ones of the different motion patterns away. However, blindly increasing the discriminative abilities of latent distributions might break the accuracy of prediction. Therefore, we also optimize the latent distribution with the prediction accuracy via the policy gradient algorithm. To evaluate whether the learned latent variable can represent motion pattern, we further propose a new evaluation metric for stochastic trajectory prediction methods. We calculate the probability cumulative minimum distance (PCMD) curve to comprehensively and stably evaluate the learned model and latent distribution, which cumulatively selects the minimum distance between sampled trajectories and ground-truth trajectories from high probability to low probability. We highlight that DisDis is a plug-and-play module which could be integrated with existing multi-modal stochastic predictive models to enhance the discriminative ability of latent distribution. Note that DisDis method only changes the loss function, which doesn't raise any extra parameters. We show that our DisDis method can achieve improved performance on the ETH [21] and UCY [27] datasets. We summarize the main contributions as follows:

- We propose a DisDis method to learn the discriminative personalized latent distribution via the self-supervised contrastive learning and the discrete optimization of consistency constraint.
- We further propose a new metric for stochastic methods which comprehensively and stably evaluate predicted trajectories under the multi-modal distribution.
- The proposed DisDis method could be integrated with existing stochastic predictors and achieve improved performance in the experiments.

## 2. Related Work

**Social Interactions:** Many previous methods model the complex social human interaction with energy potentials based on handcrafted rules, e.g. Social Force Model [14], continuum dynamics [41], Discrete Choice framework [2], Gaussian processes [39, 43], crowd analysis [30, 46], and social sensitivity [29]. With the development of deep learning, some methods [1, 12] use the social pooling layer to aggregate the clues from neighboring trajectories, while attention models [42, 32, 18, 10, 23, 22] are also introduced to distinguish the importance of different neighbors or cues. To further analyze social interactions, many recent methods [48, 47, 16, 34, 18, 24, 15, 37] employ the spatial temporal graph as the encoder of social interactions. Besides the human interactions, many studies [19, 33, 45, 18] incorporate environment interactions (e.g. physical scene) as additional clues.

**Deterministic Prediction:** Most previous methods [14, 1, 20, 48, 25, 45, 42] forecast human trajectory in a deterministic manner. They regard the trajectory forecasting as the sequence prediction problem and apply models like Recurrent Neural Networks (RNN) [1, 45, 42, 48], Temporal Convolution Neural Network (TCNN) [25], and Inverse Reinforcement Learning (IRL) [20]. For example, Social LSTM [1] encodes each history trajectory with an LSTM network and connects neighboring trajectories with a social pooling layer, then applies an LSTM decoder for sequential prediction, while Nikhil *et al.* [25] employ the Temporal CNN to model the temporal connections of human motion and predict future motions. IRL method [20] regards the motion as a Markov Decision Process for optimization. However, these deterministic models are infeasible to handle multiple possibilities of human behaviors.

**Stochastic Prediction:** To explore the indeterminacy of future states, many stochastic prediction methods are proposed to predict multiple plausible paths. These methods always incorporate a latent variable into original predictor, such as GAN [11] model or CVAE [35]. GAN based methods [12, 18, 32, 50, 8, 36] implicitly model the multi-modality and optimize generated trajectories with a discriminator, while CVAE-based methods [19, 16, 38, 34, 23] explicitly represent the multi-modal distribution and learn it with the constraint between prior and posterior distributions. Stochastic methods generate a distribution of potential motion paths, instead of predicting a most-likely single output. However, there is a large performance gap in stochastic models[12, 34] between the most-likely single output and the best posterior selection from distribution. It indicates that the latent distribution is not well learned to represent personalized future behaviors. To address this problem, we propose a DisDis method, which distinguishes the latent distributions in a self-supervised manner.

### 3. Approach

#### 3.1. Problem Definition

We formulate the pedestrian trajectory prediction task as a sequential prediction problem, which extracts the clues from prior movements and social interactions to predict the possible future navigation movements of pedestrians. The inputs of the prediction model are the observed  $N$  history human trajectories within the scene,  $\mathbf{x}_i = \{(x_i^t, y_i^t) \in \mathbb{R}^2 | t = 1, 2, \dots, T_{obs}\}$  for  $\forall i \in \{1, 2, \dots, N\}$ , where the  $(x_i^t, y_i^t)$  is the 2D location at time-instant  $t$ . Given all above inputs, the goal of our model is to learn the potential distribution to generate plausible future trajectories  $\mathbf{y}_i = \{(x_i^t, y_i^t) \in \mathbb{R}^2 | t = T_{obs} + 1, T_{obs} + 2, \dots, T_{pred}\}$ .

For the sake of simplicity, we omit the index  $i$  in the following descriptions and use  $\mathbf{x}, \mathbf{y}$  to respectively denote observed, and future trajectories.

The variational trajectory prediction methods [19, 16, 38, 34, 23] always learn a latent variable to encode the multi-modal distribution and generate the stochastic future states. The predictive process is as follows: the latent variable  $\mathbf{z}$  is generated from a distribution  $p_\theta(\mathbf{z}|\mathbf{x})$  and the future trajectory  $\mathbf{y}$  is predicted by the generative distribution  $g_\psi(\mathbf{y}|\mathbf{x}, \mathbf{z})$ . This generative process can be formulated as a target distribution  $p(\mathbf{y}|\mathbf{x})$ :

$$p(\mathbf{y}|\mathbf{x}) = \int g_\psi(\mathbf{y}|\mathbf{x}, \mathbf{z})p_\theta(\mathbf{z}|\mathbf{x})d\mathbf{z}. \quad (1)$$

In the training process, the distribution  $p_\theta(\mathbf{z}|\mathbf{x})$  is constrained with Kullback-Leibler of a prior distribution. Conventionally, the prior distribution is fixed, i.e., a Gaussian distribution  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ . It indicates the predicted trajectories  $\mathbf{y}$  is controlled by a fixed  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ , instead of the real-world personalized motion distribution.

#### 3.2. Distribution Discrimination

In this work, we present that the latent variable denotes the personalized motion patterns. The core of our work is to learn a discriminative latent variable distribution. Unlike original variational prediction methods which learn the latent variable with only the consistency constraint between proposal distribution and the prior distribution, we learn a distribution discriminator to distinguish the potential behavior patterns. We enforce three important criteria in the optimization process of latent distribution:

- The latent variable distribution should help to predict the accurate pedestrian trajectories. Besides the diversity of multi-modal motions, the accuracy of the prediction model is also of critical importance in the trajectory forecasting system.
- The latent variable distribution should be consistent for the history trajectories and corresponding future

predictions. We assume that the motion pattern of the same person is consistent due to his/her habit and character.

- The latent variable distributions should be discriminative. To model the personalized motion patterns, we encourage the latent variable distributions of the same motion pattern to be similar while pushing the ones of the different motion patterns away.

To achieve these criteria, we formulate the following loss function to learn the latent variable distribution:

$$\begin{aligned} \mathcal{L}_{DisDis}(\psi, \phi, \theta) &= \mathcal{L}_1(\phi, \theta) + \lambda \mathcal{L}_2(\phi, \psi) + \mu \mathcal{L}_3(\phi) \\ &= \mathcal{L}_{KL}(q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y}) || p_\theta(\mathbf{z}|\mathbf{x})) \\ &\quad - \lambda \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y})}[\log g_\psi(\mathbf{y}|\mathbf{x}, \mathbf{z})] \quad (2) \\ &\quad - \mu \mathbb{E}_{q_\phi(\mathbf{z})} \left[ \log \frac{h(\mathbf{z}, \mathbf{x})}{\sum_{q_\phi(\mathbf{z})} h(\mathbf{z}, \mathbf{x})} \right], \end{aligned}$$

where  $h(\mathbf{z}, \mathbf{x})$  denotes the density ratio which preserves the mutual information between  $\mathbf{z}$  and  $\mathbf{x}$ ,  $\lambda$  and  $\mu$  are two hyper-parameters to balance the effects of different terms in the objective function for a good trade-off. The  $q_\phi(\mathbf{z})$  is the marginal distribution evaluated with unbiased samples from the distribution  $q_\phi(\mathbf{z}) = q_\phi(\mathbf{z}|\mathbf{x})p_{\mathcal{D}}(\mathbf{x})$  [49], where  $p_{\mathcal{D}}(\mathbf{x})$  denotes the true underlying distribution approximated by a training set (or the samples selected in the batch). We explain each term of our DisDis loss function as follows:

- Inspired by CVAE [35], we reduce the KL divergence between the proposal distribution  $q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y})$  and prior distribution  $p_\theta(\mathbf{z}|\mathbf{x})$  in the first term of (2). The proposal distribution encodes the future motions, while the prior distribution only observes the history trajectories. This negative KL divergence encourages the consistency of latent variables generated from history and future motions.
- We also encourage the learned latent variable to predict the real future trajectory. Different from the reconstruction part of conventional VAE methods which sample latent variables from distribution to optimize the predictor  $g_\psi$  independently, we jointly learn the latent variable distribution and the predictor. We apply the discrete optimization algorithms, e.g., policy gradient, to learn the latent variable distribution, since the sampling process can not be directly optimized with the gradient back-propagation,

$$\nabla \mathcal{L}_2(\phi) \approx -\frac{1}{N} \sum_{i=1}^N [\nabla_\phi \log q_\phi(\mathbf{z}_i|\mathbf{x}_i, \mathbf{y}_i)R], \quad (3)$$

where the policy is defined as the sampling process of the latent variable, and the reward  $R$  denotes the negative  $\mathcal{L}_1$ . The policy gradients are estimated with the Monte Carlo method with selected  $N$  samples.

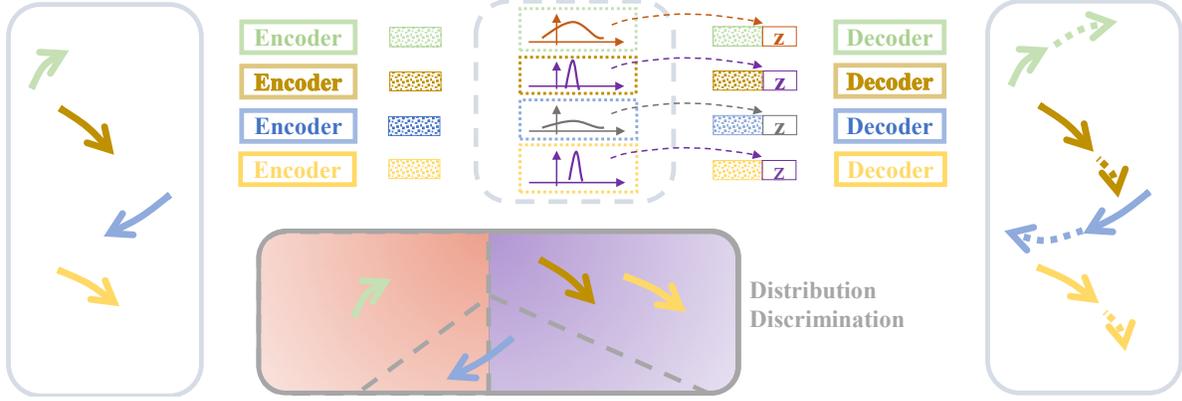


Figure 2. Training process for the DisDis method. DisDis regards the latent distribution as the motion pattern and optimizes the trajectories with the same motion pattern to be close while the ones with different patterns are pushed away, where the same latent distributions are in the same color. For a given history trajectory, DisDis predicts a latent distribution as the motion pattern, and takes the latent distribution as the discrimination to jointly optimize the embeddings of trajectories and latent distributions. (Best viewed in color.)

- The last term of (2) optimizes the discriminative ability of latent variables. The main intuition behind the distribution discrimination is to distinguish personalized motion patterns (the latent variable). As shown in Figure 2, we optimize a discriminative latent space, where the trajectories with the same motion pattern have similar representation while negative trajectories are pushed away. We predict a latent distribution as the motion pattern to optimize the embeddings of trajectories and latent distributions as pseudo labels. Note that profiting from the consistency constraint in  $\mathcal{L}_2$ , we optimize prior distribution  $p_\theta(\mathbf{z}|\mathbf{x})$  to replace  $q_\phi(\mathbf{z}|\mathbf{x})$ , (from  $\mathcal{L}_3(\phi)$  to  $\mathcal{L}_3(\theta)$ ). In the following, we give two perspectives to understand our distribution discrimination formulation in  $\mathcal{L}_3(\theta)$ , including the contrastive metric learning and mutual information optimization.

*The perspective of contrastive metric learning:* We encourage the model to learn a discriminative latent variable space, where the latent distributions (motion patterns) denote a recognition goal of trajectories. Considering the embedding  $\mathbf{f}$  of the trajectory  $\mathbf{x}$ , we define the  $h(\mathbf{z}, \mathbf{x})$  as an energy-based model:

$$h(\mathbf{z}, \mathbf{x}) = \exp(\mathbf{z}^T W^T \mathbf{f}), \quad (4)$$

where  $W^T \mathbf{f}$  is a linear transformation to predict the motion pattern  $\mathbf{z}$  with input  $\mathbf{x}$ . We assume that the latent variable denotes the motion pattern, and call the trajectories with same motion pattern as positive trajectories  $\mathbf{f}^+$  while the others as negative ones  $\mathbf{f}^-$ . When using this linear transformation to parameterize the  $p_\theta(\mathbf{z}|\mathbf{x})$ , we can obtain that  $\mathbf{z} = W^T \mathbf{f} = W^T \mathbf{f}^+$ , and write the original  $\mathcal{L}_3(\theta)$  as:

$$\begin{aligned} \mathcal{L}_3(\theta) &= -\log \frac{\exp(d(\mathbf{f}, \mathbf{f}^+))}{\sum_{\mathbf{f}_i \sim p_D(\mathbf{f})} \exp(d(\mathbf{f}, \mathbf{f}_i))} \\ &= \log \left( 1 + \sum_{\mathbf{f}_i \sim p_D(\mathbf{f})} \exp(d(\mathbf{f}, \mathbf{f}_i) - d(\mathbf{f}, \mathbf{f}^+)) \right), \end{aligned} \quad (5)$$

where,  $d(\mathbf{f}, \mathbf{f}^+) = \mathbf{f}^T W W^T \mathbf{f}^+$  is the Mahalanobis distance of embedding  $\mathbf{f}$ , which can be replaced with other distances, i.e. cosine distance. We represent the latent variable  $\mathbf{z}$  with a projection of the trajectory embedding or the positive one. This formulation is regarded as a typical metric learning objective which focuses entirely on shortening the distance between the embeddings of the trajectories with the same motion pattern while enlarging the distances between those of different motion patterns.

*The perspective of mutual information optimization:* The distribution discrimination is also equal to optimizing the mutual information  $I(\mathbf{z}, \mathbf{x})$ , which denotes the dependence between observed trajectories  $\mathbf{x}$  and corresponding latent motion pattern  $\mathbf{z}$ . When the energy-based function  $h(\mathbf{z}, \mathbf{x})$  is defined as the density ratio between  $p_\theta(\mathbf{z}|\mathbf{x})$  and  $p_\theta(\mathbf{z})$ :

$$h(\mathbf{z}, \mathbf{x}) \propto \frac{p_\theta(\mathbf{z}|\mathbf{x})}{p_\theta(\mathbf{z})}, \quad (6)$$

we reformulate the  $\mathcal{L}_3$  as a lower bound of the mutual information as:

$$\begin{aligned} \mathcal{L}_3(\theta) &= -\mathbb{E}_{\mathbf{z}} \log \left[ \frac{\frac{p_\theta(\mathbf{z}|\mathbf{x})}{p_\theta(\mathbf{z})}}{\frac{p_\theta(\mathbf{z}|\mathbf{x})}{p_\theta(\mathbf{z})} + \sum_{\mathbf{z}_i \neq \mathbf{z}} \frac{p_\theta(\mathbf{z}_i|\mathbf{x})}{p_\theta(\mathbf{z}_i)}} \right] \\ &\geq -I(\mathbf{z}, \mathbf{x}) + \log(N'), \end{aligned} \quad (7)$$

where  $N'$  is the number of the selected negative samples. Thus minimizing the distribution discrimination loss is equal to increasing the mutual information between  $\mathbf{x}$  and  $\mathbf{z}$ , which reduces the information preference problem [7]. The formal proof is given by [26].

### 3.3. Discussion

**Relation to InfoVAE:** The InfoVAE [49] method optimizes the mutual information by approximating  $I(\mathbf{z}, \mathbf{x})$  with the KL divergence between  $q_\phi(\mathbf{z})$  and  $p(\mathbf{z})$ . However, evaluating  $q_\phi(\mathbf{z})$  with the sampling of  $\mathbf{x} \sim p(\mathbf{x})$  might

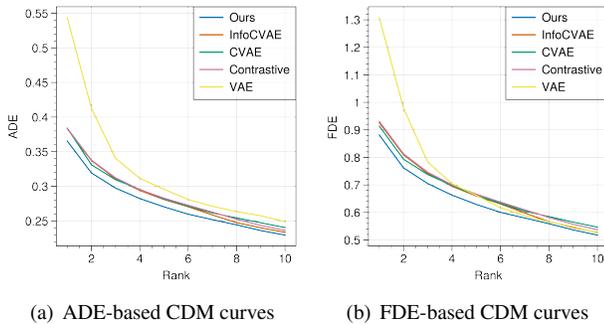


Figure 3. Ablation comparisons with baseline methods on ADE-based and FDE-based CDM curves. We set sampling  $M = 80$  and plot top 10 ranks. The lower is the better.

introduce extra noises, which influence the calculation of the KL divergence. In fact, we are primarily interested in maximizing mutual information instead of its precise value. Thus, in our DisDis method, we optimize the lower bound of mutual information instead of calculating KL divergences. In the experiments, we maintain the mutual information term in InfoVAE, and add our contrastive discrimination. The results show the advantages of adding distribution discrimination.

**Relation to Contrastive Learning:** Contrastive learning [13, 40, 44] is a kind of self-supervised methods to encourage the close embeddings under different views of the same scene. There are two main differences between conventional contrastive learning methods and our DisDis method: 1) the goal of our method is to learn the latent variable distribution instead of unsupervised representation learning; 2) we encourage the trajectory embeddings to be close to the ones of the same motion pattern, instead of different views of the same trajectory. An intuitive explanation is that we argue the latent variable as pseudo labels while conventional contrastive learning methods regard the augmented samples of same trajectory as a class.

## 4. Experiment

In this section, we first propose a new evaluation metric for stochastic prediction algorithms, which considers the comprehensive evaluation under the learned distribution. Then, we evaluate our DisDis method on two publicly available human trajectory prediction datasets: ETH [27] and UCY [21]. Quantitatively, we conduct the ablation studies to compare our method with baseline models including VAE, CVAE, InfoVAE and the combination of VAE and contrastive learning. Besides, we also compare our method with other state-of-the-art human trajectory prediction approaches. Qualitatively, we provide the visualization explanation under the real world environment and show that the proposed DisDis method can learn the discriminative latent pattern models.

### 4.1. Datasets and Experimental Settings

**Datasets:** Our experiments are conducted on two publicly available datasets: ETH [27] and UCY [21], which serve as the major benchmark for human trajectory prediction task. These datasets contain 1536 detected pedestrians in five unique scenes: Zara1, Zara2, Univ, ETH, and Hotel. We follow the commonly used leave one set out cross-validation evaluation strategy, i.e., training on four scenes and testing on the remaining one [12, 18, 15, 34]. All trajectories are sampled in 0.4 seconds (20 frames), where the first 3.2 seconds correspond to observed data and the next 4.8 seconds correspond to predicted future data.

**Evaluation Metric:** The same evaluation metrics as prior methods [12, 18, 15, 34] are adopted, including Average Displacement Error (ADE) and Final Displacement Error (FDE). ADE computes the mean square error (MSE) overall estimated positions in the predicted trajectory and ground-truth trajectory, while FDE computes the distance between the predicted final destination and the ground-truth final destination.

Instead of the deterministic prediction, stochastic predicted methods generate the future motions under a multi-modal distribution. There are two widely used strategies to compute the ADE and FDE for the predicted multi-modal trajectories:

- **“Most-likely” strategy** predicts the most-likely single output under the multi-modal distribution and uses it to calculate the ADE and FDE. It changes the stochastic prediction to deterministic one to obtain the most-likely performance.
- **“Best-of-N” strategy** generates  $N$  samples (i.e.  $N = 20$  [12, 34, 18]) based on the predicted distribution and selects the closest sample to the ground truth for computing ADE and FDE metrics. It evaluates the highest performance of the model by considering as many samples as possible.

However, both two strategies ignore the evaluation of learned latent distribution. It is of critical importance for stochastic prediction algorithms since the latent variable denotes the personalized motion pattern of human. In this paper, we propose a new evaluation strategy, called probability cumulative minimum distance (PCMD) curve, to evaluate the prediction model under the multi-modal distribution. There are two motivations of our PCMD evaluation metric: 1) Though the Most-likely strategy considers the latent distribution, it only evaluates the most-likely point. The single point cant represent the latent distribution. 2) Best-of-N strategy samples multiple points from the latent distribution. However, it only uses the best one for evaluation, which ignores the probability of each sample. Different from these two strategies, PCMD considers the predictions with corre-

Table 1. Comparison with several state-of-the-art models with PCMD curves. All the models are reproduced with the code previously published online. We show the values on PCMD curves at rank 1, 5 and 20. The lower is the better.

Methods	$PCMD_{ADE}@{\frac{1}{M}/\frac{5}{M}/\frac{20}{M}} M=80$					
	ETH	HOTEL	ZARA1	ZARA2	UNIV	AVG
Social-GAN [12]	0.98/0.82/0.73	0.63/0.54/0.48	0.47/0.38/0.33	0.39/0.33/0.3	0.64/0.58/0.55	0.62/0.53/0.48
STGAT [15]	1.03/0.89/0.78	0.59/0.45/0.38	0.53/0.39/0.33	0.43/0.34/0.29	0.68/0.59/0.56	0.65/0.53/0.47
Social-STGCNN [24]	1.01/0.87/0.76	0.74/0.53/0.42	0.57/0.42/0.34	0.51/0.38/0.31	0.71/0.59/0.51	0.71/0.56/0.47
Trajectron++ [34]	0.73/0.58/0.43	0.27/0.19/0.11	0.30/0.20/0.13	<b>0.22/0.16/0.10</b>	<b>0.39/0.27/0.17</b>	0.38/0.28/0.19
DisDis(Ours)	<b>0.71/0.55/0.38</b>	<b>0.25/0.17/0.11</b>	<b>0.28/0.20/0.13</b>	<b>0.22/0.16/0.10</b>	<b>0.36/0.27/0.17</b>	<b>0.36/0.27/0.18</b>

Methods	$PCMD_{FDE}@{\frac{1}{M}/\frac{5}{M}/\frac{20}{M}} M=80$					
	ETH	HOTEL	ZARA1	ZARA2	UNIV	AVG
Social-GAN [12]	1.98/1.58/1.40	1.36/1.14/1.02	1.02/0.79/0.67	0.87/0.70/0.63	1.38/1.24/1.17	1.32/1.09/1.03
STGAT [15]	2.20/1.86/1.52	1.21/0.89/0.73	1.17/0.81/0.67	0.94/0.70/0.60	1.49/1.28/1.20	1.40/1.11/0.94
Social-STGCNN [24]	1.83/1.54/1.29	1.42/1.00/0.74	1.13/0.79/0.57	0.97/0.69/0.52	1.38/1.15/0.98	1.35/1.03/0.82
Trajectron++ [34]	1.73/1.36/0.92	0.57/0.37/0.18	0.76/0.49/0.27	0.57/0.41/0.23	0.99/0.67/0.36	0.93/0.66/0.39
DisDis(Ours)	<b>1.67/1.24/0.75</b>	<b>0.52/0.33/0.18</b>	<b>0.73/0.50/0.27</b>	<b>0.56/0.41/0.22</b>	<b>0.92/0.66/0.37</b>	<b>0.88/0.63/0.36</b>

Table 2. Comparison with SOTA models with Best-of-20 strategy

Methods	ADE	FDE
Social-GAN [12]	0.58	1.18
STGAT [15]	0.43	0.83
Social-STGCNN [24]	0.44	0.75
Trajectron++ [34]	0.19	0.41
DisDis(Ours)	<b>0.17</b>	<b>0.37</b>

sponding probabilities, which evaluate the model under the whole latent distribution. Inspired by the Cumulative Match Characteristic (CMC) curve in the recognition tasks, we cumulatively calculate the minimum ADE and FDE from high probability to low probability.

Formally, given the distribution  $p_\theta(\mathbf{z}|\mathbf{x})$ ,  $\mathbf{z} \in \Omega$ , we can define a projection as:

$$F(\tau) = \mathbb{E}_{\mathbf{z} \in \Omega} \mathbb{I}(p_\theta(\mathbf{z}|\mathbf{x}) \geq \tau), \quad (8)$$

where the input denotes a selected probability value  $\tau \in (\min p_\theta(\mathbf{z}|\mathbf{x}), \max p_\theta(\mathbf{z}|\mathbf{x}))$ ,  $\mathbb{I}$  is an indicator function, while the output  $F(\tau) \in (0, \|\Omega\|)$  denotes the interval length of  $\mathbf{z}$  satisfying conditions. We define  $k = \frac{F(\tau)}{\|\Omega\|}$  to normalize it into  $(0, 1)$ , if the latent space  $\Omega$  is finite. (The condition  $\|\Omega\| \rightarrow \infty$  can be ignored when we discretely approximate it by the Monte Carlo methods. ) We obtain the values of PCMD curve as:

$$PCMD(k) = \mathbb{E}_{\mathbf{x}} \min\{\mathcal{D}(\mathbf{z}) | p_\theta(\mathbf{z}|\mathbf{x}) \geq \tau, \mathbf{z} \in \Omega\}, \quad (9)$$

where  $\mathcal{D}(\mathbf{z})$  denotes the ADE or FDE distance between the ground-truth and the sampled predicted trajectories based

on  $\mathbf{z}$ . The  $PCMD(k)$  denotes the minimum ADE or FDE distance from all sampled latent variable satisfying condition  $p_\theta(\mathbf{z}|\mathbf{x}) \geq \tau$ . The lower PCMD value indicates we can obtain better performance with the same sampling number. It is critical for auto-driving system since more sampling means more time-delay.

To numerically calculate the PCMD, we sample  $M$  variables  $Z = \{\mathbf{z}_i \in \Omega | i = 1, 2, \dots, M\}$  and sort these variables with the probability  $p_\theta(\mathbf{z}_i|\mathbf{x})$  from large to small to obtain  $Z_{sort} = \{\mathbf{z}_i^*\}$ . Then, the values of PCMD are calculated as:

$$PCMD(k) = \mathbb{E}_{\mathbf{x}} \min\{\mathcal{D}(\mathbf{z}) | \mathbf{z} \in \{\mathbf{z}_1^*, \mathbf{z}_2^*, \dots, \mathbf{z}_m^*\}\}, \quad (10)$$

where  $k = \frac{m}{M}$  denotes the ranking rate in the numerical calculation, e.g. we calculate the minimum ADE/FDE of top 20 probability trajectories, when  $k = \frac{20}{M}$ .

The proposed PCMD curve has two significant advantages over existing evaluation metrics: more comprehensive and more stable. First, PCMD comprehensively analyzes the latent distribution by the accumulation of probability, while other evaluation metrics can be regarded as a part of the PCMD curve.  $PCMD(1/M)$  denotes the most-likely single output ADE/FDE performance, while  $PCMD(N/M)$  denotes the best ADE/FDE performance of trajectories with top  $N$  probability. Note that in the experiments we used the same sampling number for all methods for fair comparison.

Besides, the proposed PCMD curve is more stable than the ‘‘Best-of- $N$ ’’ evaluation since random sampling brings large performance jitter due to the randomness. PCMD reduces this randomness by the larger sampling  $M$  ( $M \gg N$ )

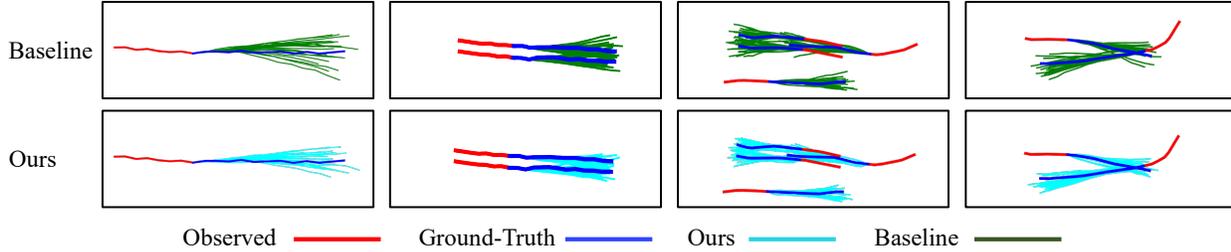


Figure 4. Illustration of the learned distributions. We compare the learned distributions of our method and Trajectron++ [34] baseline on 4 different scenarios including a pedestrian walking alone; pedestrians walking in parallel; pedestrians following the people ahead; and pedestrians meeting from different directions. We plot 20 top rank trajectories to show the distributions.

and more stable expression (a curve rather than a number). For the discrete latent distribution, we can obtain a totally fixed evaluation due to the complete sampling. Sampling  $M$  trajectories is only conducted when we evaluate different models and select the better one. In a real-life scenario, we can only generate a part of trajectories with higher probabilities using the selected model.

## 4.2. Implementation Details

We select the SOTA Trajectron++ [34]<sup>2</sup> method, which is based on CVAE and InfoVAE, as the baseline and incorporate DisDis. We follow the encoder and decoder networks of Trajectron++, and make some modifications for convenient and fair comparison: 1) we did not use the scene clues (e.g. local map) and removed complex dynamics integration to obtain simpler embedding; 2) we modeled the latent variable with a higher dimensional discrete variable ( $|Z| = 80$ ) which better evaluates the learned latent variable distribution; 3) we used a fixed most likely output to replace the Gaussian output structure, which guarantees that the diversity is only from the learned latent distribution. We adopt the same data preprocessing strategies with Trajectron++ [34], which rotates trajectories in a scene around the scene’s origin in a interval of  $15^\circ$ .

## 4.3. Quantitative Evaluation

**Comparison with Baseline Methods:** We conduct ablation study to compare our method with the following baseline methods with the same network implementation and different latent distribution learning: 1) VAE [17] optimizes the latent distribution towards a fixed prior distribution. 2) CVAE [35] learns the latent variable with the KL divergence between the proposal distribution  $q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y})$  and prior distribution  $p_\theta(\mathbf{z}|\mathbf{x})$ . 3) CVAE + InfoVAE [49] further considers the mutual information between  $\mathbf{x}$  and  $\mathbf{z}$  by approximating the distribution  $q_\phi(\mathbf{z})$ . 4) CVAE + Contrastive Learning optimizes to the similar latent variable for different views of the same trajectory.

<sup>2</sup>We reproduced the Trajectron++ [34] method with the official released code from <https://github.com/StanfordASL/Trajectron-plus-plus>.

We evaluate our method and the baselines in our proposed PCMD curve metric. As shown in Figure 3, the original VAE baseline obtains the worst performance on top probability ranks, which indicates bias between the fixed prior distribution and the real-world personalized motion distribution, while our DisDis obtains best top rank performance due to the more discriminative distribution. Compared with CVAE and Info+CVAE, our DisDis also achieves an improved performance, which demonstrates optimizing the mutual information by distribution discrimination is a better manner than approximately calculating the  $q_\phi(\mathbf{z})$ . Besides, the DisDis method outperforms directly applying contrastive self supervised learning on VAE. It demonstrates that using the latent distributions to represent motion patterns is better than regarding each trajectory as one motion pattern. In our assumption, the trajectories with the same motion pattern are clustered in the embedding space. When increasing the number of cumulative probabilities, performance gaps among all methods become small, which is because we use the same encoder and decoder networks for all baseline methods and our method. *Note that* the performances of all methods will converge to the same point when we sample enough trajectories since we use the same networks. The main difference among these methods is the optimization strategy of the latent variables. In this setting, the latent distribution of our DisDis method outperforms the ones of other methods since we can achieve same performance with fewer sampling.

**Comparison with SOTA Stochastic Methods:** We always compare our method with a wide range of SOTA stochastic methods, including: Social-GAN [12], STGAT [15], Social-STGCNN [24], and Trajectron++ [34]. All the models are reproduced with the code previously published online, and **fairly** evaluated with PCMD curve metric. As shown in Table 1, we provide some detailed points including the performance of  $PCMD@{\frac{1}{M}, \frac{5}{M}, \frac{20}{M}}$ . For convenience, we define the  $PCMD@{\frac{1}{M}, \frac{5}{M}, \frac{20}{M}}$  as Rank 1, Rank 5 and Rank 20. Rank 1 of ADE/FDE denotes the most-likely single output ADE/FDE performance, while Rank 20 ADE/FDE denotes the best ADE/FDE performance of trajectories with top 20

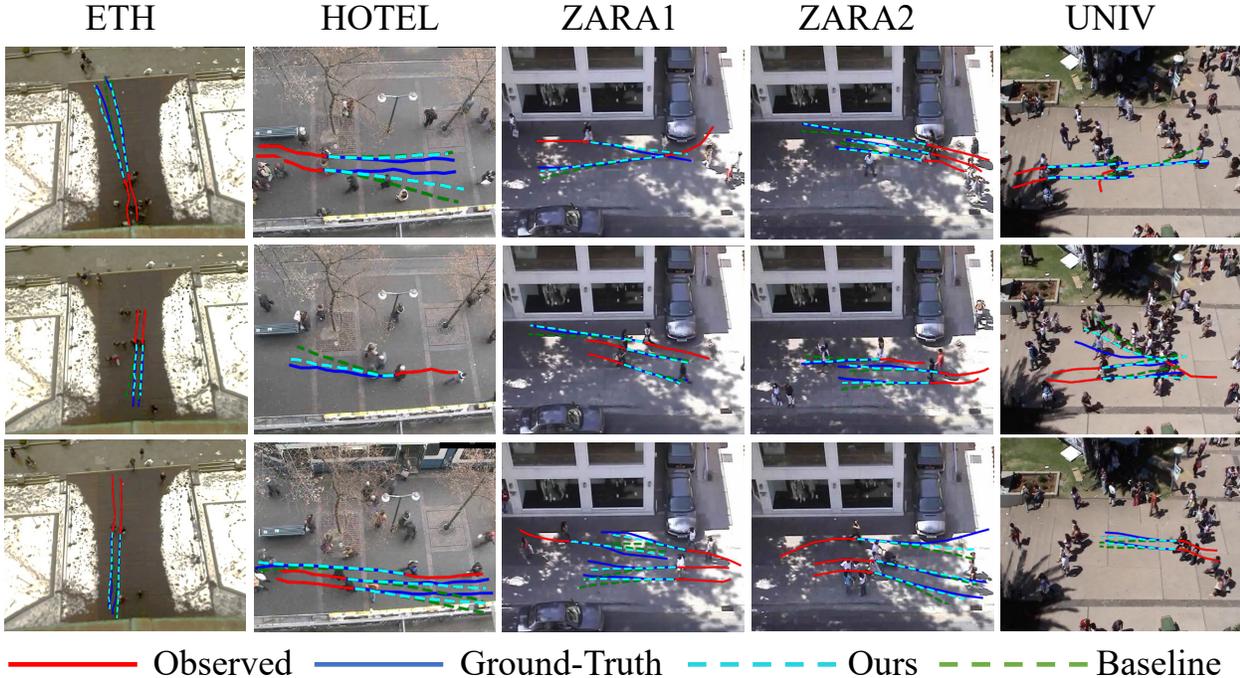


Figure 5. Visualization examples of our DisDis method and baseline Trajectron++ [34] method in the different scenes in the both ETH and UCY datasets. We plot the most likely predictions as the results, which quantitatively demonstrate the effectiveness of distribution discrimination to learn the discriminative latent distribution. (Best viewed in color.)

probabilities. We observe that the baseline Trajectron++ is the current SOTA method with the strong representation ability of social interactions. And with a better latent variable, our DisDis method achieves further improvement on the Rank 1 with no cost of the performance of Rank 20. It is because DisDis can select better predicted trajectories in higher probability. We also report the original “best of 20” performance in Table 2.

#### 4.4. Qualitative Evaluation

We qualitatively analyze how our DisDis method learns the personalized latent variable and improves motion predictions. Figure 4 visualizes the learned latent distributions of our method and baseline method in 4 different scenarios including a pedestrian walking alone; pedestrians walking in parallel; pedestrians following the people ahead; and pedestrians meeting from different directions. We plot 20 top rank trajectories to represent the learned distributions. For all scenarios, we observe that our predicted trajectories are closer to ground-truth trajectories than the ones of baseline, which indicates more discriminative latent distributions. Taking the first image as the example, our predicted trajectories are obviously more concentrative than the ones of baseline method.

Besides, we also provide the visualization examples with different scenes in the both ETH and UCY datasets in Figure 5. We plot the most-likely predictions of our DisDis method and the baseline Trajectron++ [34]. We observe that

predicted trajectories are more similar to ground-truth trajectories, which demonstrates our method can obtain better prediction with higher probability.

## 5. Conclusion

In this paper, we have proposed a distribution discrimination (DisDis) method to distinguish latent distribution via self-supervised contrastive learning. The DisDis method encourages the model to learn a discriminative latent variable space, where the trajectories with the same motion pattern have similar latent variable distribution while negative trajectories are pushed away. To evaluate the learned latent distribution, we further propose a probability cumulative minimum distance (PCMD) curve as the metric of stochastic trajectory prediction methods, which cumulatively calculates the minimum distance on the sorted probabilities. Finally, we show that our DisDis method could be integrated with existing stochastic predictors and obtain improvement in the learning of latent distribution.

**Acknowledgments** This work was supported in part by the National Natural Science Foundation of China under Grant 61822603, Grant U1813218, and Grant U1713214, in part by a grant from the Beijing Academy of Artificial Intelligence (BAAI), and in part by a grant from the Institute for Guo Qiang, Tsinghua University.

## 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 *CVPR*, pages 961–971, 2016.
- [2] Gianluca Antonini, Michel Bierlaire, and Mats Weber. Discrete choice models of pedestrian walking behavior. *Transportation Research Part B: Methodological*, 40(8):667–687, 2006.
- [3] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning discriminative model prediction for tracking. In *ICCV*, pages 6182–6191, 2019.
- [4] Changan Chen, Yuejiang Liu, Sven Kreiss, and Alexandre Alahi. Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning. In *ICRA*, pages 6015–6022, 2019.
- [5] Guangyi Chen, Chunze Lin, Liangliang Ren, Jiwen Lu, and Zhou Jie. Self-critical attention learning for person re-identification. In *ICCV*, 2019.
- [6] Guangyi Chen, Jiwen Lu, Ming Yang, and Jie Zhou. Spatial-temporal attention-aware learning for video-based person re-identification. *TIP*, 28(9):4192–4205, 2019.
- [7] Xi Chen, Diederik P Kingma, Tim Salimans, Yan Duan, Prafulla Dhariwal, John Schulman, Ilya Sutskever, and Pieter Abbeel. Variational lossy autoencoder. *arXiv preprint arXiv:1611.02731*, 2016.
- [8] Liangji Fang, Qinhong Jiang, Jianping Shi, and Bolei Zhou. Tpnnet: Trajectory proposal network for motion prediction. In *CVPR*, June 2020.
- [9] Panna Felsen, Patrick Lucey, and Sujoy Ganguly. Where will they go? predicting fine-grained adversarial multi-agent motion using conditional variational autoencoders. In *ECCV*, pages 732–747, 2018.
- [10] 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.
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, pages 2672–2680, 2014.
- [12] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social gan: Socially acceptable trajectories with generative adversarial networks. In *CVPR*, pages 2255–2264, 2018.
- [13] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, June 2020.
- [14] Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. *Physical Review E*, 51(5):4282, 1995.
- [15] Yingfan Huang, HuiKun Bi, Zhaoxin Li, Tianlu Mao, and Zhaoqi Wang. Stgat: Modeling spatial-temporal interactions for human trajectory prediction. In *ICCV*, pages 6272–6281, 2019.
- [16] Boris Ivanovic and Marco Pavone. The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs. In *ICCV*, pages 2375–2384, 2019.
- [17] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014.
- [18] Vineet Kosaraju, Amir Sadeghian, Roberto Martín-Martín, Ian Reid, Hamid Rezaatofghi, and Silvio Savarese. Socialbigat: Multimodal trajectory forecasting using bicycle-gan and graph attention networks. In *NeurIPS*, pages 137–146, 2019.
- [19] Namhoon Lee, Wongun Choi, Paul Vernaza, Christopher B Choy, Philip HS Torr, and Manmohan Chandraker. Desire: Distant future prediction in dynamic scenes with interacting agents. In *CVPR*, pages 336–345, 2017.
- [20] Namhoon Lee and Kris M Kitani. Predicting wide receiver trajectories in american football. In *WACV*, pages 1–9, 2016.
- [21] Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. Crowds by example. In *Computer Graphics Forum*, volume 26, pages 655–664, 2007.
- [22] 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 *CVPR*, pages 5725–5734, 2019.
- [23] 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 *ECCV*, 2020.
- [24] Abdullh Mohamed, Kun Qian, Mohamed Elhoseiny, and Christian Claudel. Social-stgcn: A social spatio-temporal graph convolutional neural network for human trajectory prediction. In *CVPR*, 2020.
- [25] Nishant Nikhil and Brendan Tran Morris. Convolutional neural network for trajectory prediction. In *ECCVW*, pages 0–0, 2018.
- [26] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [27] Stefano Pellegrini, Andreas Ess, and Luc Van Gool. Improving data association by joint modeling of pedestrian trajectories and groupings. In *ECCV*, pages 452–465, 2010.
- [28] Scott Drew Pendleton, Hans Andersen, Xinxin Du, Xiaotong Shen, Malika Meghjani, You Hong Eng, Daniela Rus, and Marcelo H Ang. Perception, planning, control, and coordination for autonomous vehicles. *Machines*, 5(1):6, 2017.
- [29] Alexandre Robicquet, Amir Sadeghian, Alexandre Alahi, and Silvio Savarese. Learning social etiquette: Human trajectory understanding in crowded scenes. In *ECCV*, pages 549–565, 2016.
- [30] Mikel Rodriguez, Josef Sivic, Ivan Laptev, and Jean-Yves Audibert. Data-driven crowd analysis in videos. In *ICCV*, pages 1235–1242, 2011.
- [31] Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M Kitani, Dariu M Gavrila, and Kai O Arras. Human motion trajectory prediction: A survey. *IJRR*, 2020.
- [32] 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 *CVPR*, June 2019.
- [33] 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 *CVPR*, pages 1349–1358, 2019.
- [34] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *ECCV*, 2020.
- [35] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *NeurIPS*, pages 3483–3491, 2015.
- [36] Hao Sun, Zhiqun Zhao, and Zhihai He. Reciprocal learning networks for human trajectory prediction. In *CVPR*, June 2020.
- [37] Jianhua Sun, Qinrong Jiang, and Cewu Lu. Recursive social behavior graph for trajectory prediction. In *CVPR*, June 2020.
- [38] Charlie Tang and Russ R Salakhutdinov. Multiple futures prediction. In *NeurIPS*, pages 15398–15408, 2019.
- [39] Meng Keat Christopher Tay and Christian Laugier. Modelling smooth paths using gaussian processes. In *Field and Service Robotics*, pages 381–390, 2008.
- [40] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019.
- [41] Adrien Treuille, Seth Cooper, and Zoran Popović. Continuum crowds. In *TOG*, volume 25, pages 1160–1168, 2006.
- [42] Anirudh Vemula, Katharina Muelling, and Jean Oh. Social attention: Modeling attention in human crowds. In *ICRA*, pages 1–7, 2018.
- [43] Jack M Wang, David J Fleet, and Aaron Hertzmann. Gaussian process dynamical models for human motion. *TPAMI*, 30(2):283–298, 2007.
- [44] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, pages 3733–3742, 2018.
- [45] Hao Xue, Du Q Huynh, and Mark Reynolds. Ss-lstm: A hierarchical lstm model for pedestrian trajectory prediction. In *WACV*, pages 1186–1194, 2018.
- [46] Kota Yamaguchi, Alexander C Berg, Luis E Ortiz, and Tamara L Berg. Who are you with and where are you going? In *CVPR*, pages 1345–1352, 2011.
- [47] Cunjun Yu, Xiao Ma, Jiawei Ren, Haiyu Zhao, and Shuai Yi. Spatio-temporal graph transformer networks for pedestrian trajectory prediction. In *ECCV*, 2020.
- [48] Pu Zhang, Wanli Ouyang, Pengfei Zhang, Jianru Xue, and Nanning Zheng. Sr-lstm: State refinement for lstm towards pedestrian trajectory prediction. In *CVPR*, pages 12085–12094, 2019.
- [49] Shengjia Zhao, Jiaming Song, and Stefano Ermon. Infovae: Balancing learning and inference in variational autoencoders. In *AAAI*, pages 5885–5892, 2019.
- [50] 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 *CVPR*, pages 12126–12134, 2019.