

MBZUAI

Digital.Commons@MBZUAI

Machine Learning Faculty Publications

Scholarly Works

9-27-2022

Stochastic Trajectory Prediction via Motion Indeterminacy Diffusion

Tianpei Gu

University of California, Los Angeles & SenseTime Research

Guangyi Chen

Mohamed Bin Zayed University of Artificial Intelligence & Carnegie Mellon University, United States

Junlong Li

Tsinghua University

Chunze Lin

SenseTime Research

Yongming Rao

Tsinghua University

See next page for additional authors

Follow this and additional works at: <https://dclibrary.mbzuai.ac.ae/mlfp>



Part of the [Artificial Intelligence and Robotics Commons](#)

Open Access version, provided by the [Computer Vision Foundation](#).

Uploaded: January 18, 2023

Recommended Citation

T. Gu et al., "Stochastic Trajectory Prediction via Motion Indeterminacy Diffusion," *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, New Orleans, LA, USA, 2022, pp. 17092-17101, doi: 10.1109/CVPR52688.2022.01660.

This Conference Proceeding is brought to you for free and open access by the Scholarly Works at Digital.Commons@MBZUAI. It has been accepted for inclusion in Machine Learning Faculty Publications by an authorized administrator of Digital.Commons@MBZUAI. For more information, please contact libraryservices@mbzuai.ac.ae.

Authors

Tianpei Gu, Guangyi Chen, Junlong Li, Chunze Lin, Yongming Rao, Jie Zhou, and Jiwen Lu

Stochastic Trajectory Prediction via Motion Indeterminacy Diffusion

Tianpei Gu^{*1,5}, Guangyi Chen^{*,2,3}, Junlong Li⁴, Chunze Lin⁵, Yongming Rao⁴, Jie Zhou⁴, Jiwen Lu^{†,4}
¹University of California, Los Angeles, ²MBZUAI, ³Carnegie Mellon University,
⁴Tsinghua University, ⁵SenseTime Research

Abstract

Human behavior has the nature of indeterminacy, which requires the pedestrian trajectory prediction system to model the multi-modality of future motion states. Unlike existing stochastic trajectory prediction methods which usually use a latent variable to represent multi-modality, we explicitly simulate the process of human motion variation from indeterminate to determinate. In this paper, we present a new framework to formulate the trajectory prediction task as a reverse process of motion indeterminacy diffusion (MID), in which we progressively discard indeterminacy from all the walkable areas until reaching the desired trajectory. This process is learned with a parameterized Markov chain conditioned by the observed trajectories. We can adjust the length of the chain to control the degree of indeterminacy and balance the diversity and determinacy of the predictions. Specifically, we encode the history behavior information and the social interactions as a state embedding and devise a Transformer-based diffusion model to capture the temporal dependencies of trajectories. Extensive experiments on the human trajectory prediction benchmarks including the Stanford Drone and ETH/UCY datasets demonstrate the superiority of our method. Code is available at <https://github.com/gutianpei/MID>.

1. Introduction

Human trajectory prediction plays a crucial role in human-robot interaction systems such as self-driving vehicles and social robots, since human is omnipresent in their environments. Although significant progresses have been achieved over past few years [6, 28, 29, 32, 38, 45, 49, 53], predicting the future trajectories of pedestrians remains challenging due to the multi-modality of human motion.

The future trajectories of pedestrians are full of indeterminacy, because human can change future motion according to their will or adjust their movement direction based on the surroundings. Given a history of observed trajectory,

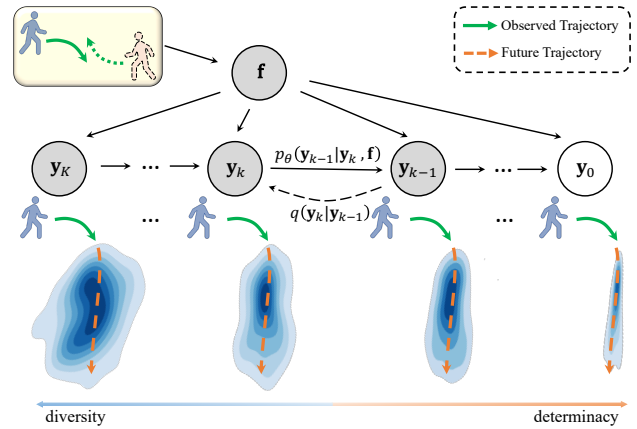


Figure 1. Illustration of the reverse diffusion process for human motion indeterminacy variation. Under high indeterminacy, the trajectory distribution can be regarded as a noise Gaussian distribution which denotes ambiguous walkable areas. With the decreasing of indeterminacy, this distribution gradually approximates the real data distribution to generate desired trajectory. This process from high indeterminacy to low indeterminacy is defined as a reverse diffusion process, in which we learn a Markov chain to progressively discard the indeterminacy. By adjusting the length of the chain, we can make a trade-off between diversity and determinacy, where the longer chain leads to lower diversity and higher determinacy. Best viewed in color.

ries, there exist many plausible paths that pedestrians could move in the future. Facing this challenge, most of prior researches apply the generative model to represent multi-modality by a latent variable. For instance, some methods [6, 9, 12, 19, 37, 43, 54] utilize generative adversarial networks (GANs) to spread the distribution over all possible future trajectories, while other methods [3, 16, 20, 25, 38, 46] exploit conditional variational auto-encoder (CVAE) to encode the multi-modal distribution of future trajectories. Despite the remarkable progress, these methods still face inherent limitations, *e.g.*, training process could be unstable for GANs due to adversarial learning, and CVAE tends to produce unnatural trajectories.

In this paper, we propose a new trajectory prediction framework, called motion indeterminacy diffusion (MID), to model the indeterminacy of human behavior. Inspired by

^{*}Equal contribution. [†]Corresponding author.

non-equilibrium thermodynamics, we consider the future positions as particles in thermodynamics in our framework. The particles (positions) gather and deform to a clear trajectory under low indeterminacy, while stochastically spread over all walkable areas under high indeterminacy. The process of particles evolving from low indeterminacy to high indeterminacy is defined as the diffusion process. This process can be simulated by gradually adding noise to the trajectory until the path is corrupted as Gaussian noise. The goal of our MID is to reverse this diffusion process by progressively discarding indeterminacy, and converting the ambiguous prediction regions into a deterministic trajectory. We illustrate the reverse diffusion process of motion indeterminacy in Figure 1. Contrary to other stochastic prediction methods that add a noise latent variable on the trajectory feature to obtain indeterminacy, we explicitly simulate the motion indeterminacy variation process. Our MID learns a Markov chain with parameterized Gaussian transition to model this reverse diffusion process and train it using variational inference conditioned on the observed trajectories. By choosing different lengths of the chain, we can obtain the predictions with a flexible indeterminacy that is capable of adapting to dynamic environment. Moreover, our method is more efficient to train than GANs, and is capable of producing more high-quality samples than CVAEs.

To be more specific, we encode the history human trajectories and the social interactions as state embedding via a spatial-temporal graph network. Then, we exploit this state embedding as condition in the Markov chain to guide the learning of reverse diffusion process. To model the temporal dependencies in trajectories, we carefully design a Transformer-based architecture as the core network of MID framework. In the training process, we optimize the model with the variational lower bound, and during the inference, we sample the reasonable trajectories by progressive denoising from a noise distribution. Extensive experiments demonstrate that our method accurately forecasts reasonable future trajectories with multi-modality, achieving state-of-the-art results on Stanford Drone and ETH/UCY datasets. We summarize the main contributions of this paper as follows:

- We present a new stochastic trajectory prediction framework with motion indeterminacy diffusion, which gradually discards the indeterminacy to obtain desired trajectory from ambiguous walkable areas.
- We devise a Transformer-based architecture for the proposed framework to capture the temporal dependencies in trajectories.
- The proposed method achieves state-of-the-art performance on widely used human trajectory prediction benchmarks and provides a potential direction for balancing the diversity and accuracy of predictions.

2. Related Work

Pedestrian Trajectory Prediction: Given the observed paths, human trajectory forecasting system aims to estimate the future positions. Most existing methods formulate trajectory forecasting as a sequential prediction problem and focus on modeling the complex social interaction. For instance, Social Forces [13] introduces attractive and repulsive forces to model human interaction. With the success of deep learning, many methods design ingenious networks to model the social interactions. For example, Social-LSTM [1] devises a social pooling layer to aggregate the interaction information of neighborhoods. Some methods apply the attention models [10, 19, 37, 47, 51] to explore the key interactions of the crowd. In addition, the spatial-temporal graph model is applied to jointly model the temporal clues and social interactions [15, 16, 30, 38, 44, 50]. Beyond social interactions, many methods incorporate the physical environment interactions by introducing the map images [6, 19, 20, 28, 37]. Recently, some methods analyze the effect of social interaction and show it is biased [2, 27].

Stochastic Prediction Model: Due to the inherent indeterminacy of human behavior, Many stochastic prediction methods are proposed to model the multi-modality of future motions. Some methods [6, 9, 12, 19, 37, 43, 54] employ GANs [11] to model the multi-modality with a noise variable, and another line of methods [3, 16, 20, 25, 38, 46] apply the CVAE [41] instead. Besides, some methods [7, 23, 24] propose to learn the grid-based location encoder for multi-modal probability prediction. Recently, the goals of pedestrians [28, 29, 52, 53] are introduced in the trajectory prediction system as condition to analyze the probability of multiple plausible endpoints. While remarkable progress have been made, these stochastic prediction methods have some inherent limitations, *e.g.*, the unstable training or unnatural trajectories. In this paper, we propose a new stochastic framework with motion indeterminacy diffusion, which formulates the trajectory prediction problem as a process from an ambiguous walkable region to the desired trajectory.

Denoising Diffusion Probabilistic Models: Denoising diffusion probabilistic models (DDPM) [14, 40], as known as diffusion models for brevity, are a class of deep generative models inspired by non-equilibrium thermodynamics. It is first proposed by Sohl-Dickstein *et al.* [40] and attracts much attention recently due to state-of-the-art performance in various generation tasks including image generation [5, 8, 14, 31], 3D point cloud generation [26, 55], and audio generation [4, 18, 34]. The diffusion models generally learn a parameterized Markov chain to gradually denoise from an original common distribution to a specific data distribution. In this paper, we introduce the diffusion model to simulate the variation of indeterminacy for trajectory prediction, and design a Transformer-based architecture for the temporal dependency of trajectories.

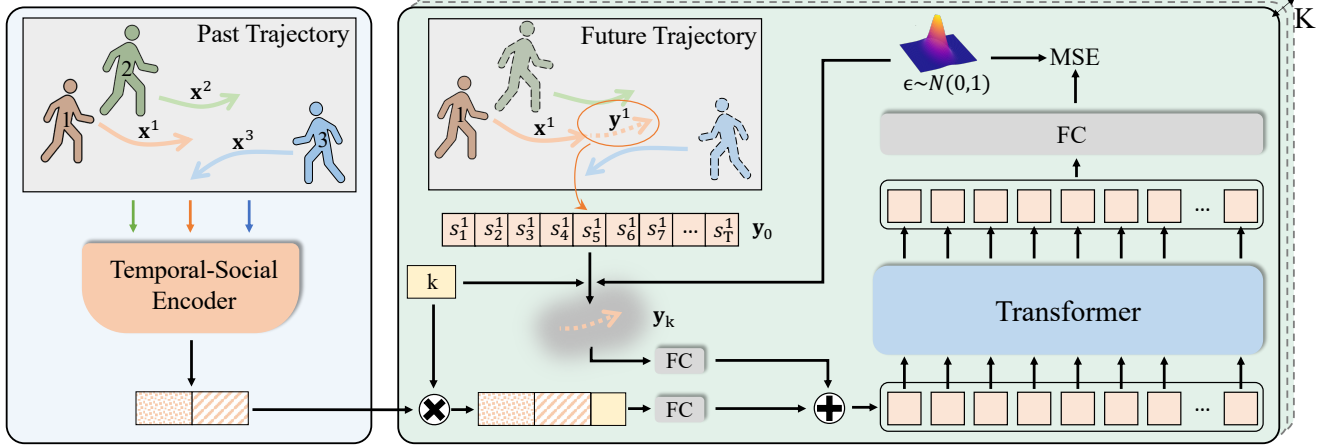


Figure 2. The architecture of our MID framework. MID consists of a temporal-social encoder network and a Transformer-based decoder network. The encoder maps the history path and social interaction clues into a state embedding. The decoder then takes \mathbf{y}_k along with state embedding and the time embedding, where \mathbf{y}_k is corrupted k times by a noise variable from the ground truth trajectory \mathbf{y}_0 . We learn the model with the MSE loss between the model output and a noise variable in standard Gaussian distribution.

3. Proposed Approach

In this section, we introduce our MID method, which models stochastic trajectory prediction task by motion indeterminacy diffusion. We first explicitly formulate the indeterminacy variation as a reverse diffusion process. Then we describe how to train this diffusion model using the variational inference. Finally, we present the detailed network architecture of our method shown in Figure 2.

3.1. Problem Formulation

The goal of pedestrian trajectory prediction is to generate plausible future trajectories for pedestrians based on their prior movements. The input of the prediction system is the N history trajectories in a scene such that $\mathbf{x}^i = \{s_t^i \in \mathbb{R}^2 | t = -T_{init}, -T_{init} + 1, \dots, 0\}$, $\forall i \in \{1, 2, \dots, N\}$, where the s_t^i is the 2D location at timestamp t , T_{init} denotes the length of the observed trajectory, and the current timestamp is $t = 0$. Similarly, the predicted future trajectories can be written as $\mathbf{y}^i = \{s_t^i \in \mathbb{R}^2 | t = 1, 2, \dots, T_{pred}\}$. For clarity, we use \mathbf{x} and \mathbf{y} without the superscript i for the history and future trajectory in the following subsections.

3.2. Motion Indeterminacy Diffusion

Due to the indeterminacy of human behavior, one person has multiple plausible paths in future state. Thus, we present a new framework to formulate the stochastic trajectory prediction by motion indeterminacy diffusion. Unlike other stochastic prediction methods that add a latent variable on the trajectory feature to obtain indeterminacy, our MID generates the trajectory by gradually reducing the indeterminacy from all walkable areas to the determinate prediction with a parameterized Markov chain.

As shown in Figure 1, given the initial ambiguous region

\mathbf{y}_K under the noise distribution and the desired trajectory \mathbf{y}_0 under the data distribution, we define the diffusion process as $(\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_K)$, where K is the maximum number of diffusion steps. This process aims to gradually add the indeterminacy until the ground truth trajectory is corrupted into a noisy walkable region. On the contrary, we learn the reverse process as $(\mathbf{y}_K, \mathbf{y}_{K-1}, \dots, \mathbf{y}_0)$ to gradually reduce the indeterminacy from \mathbf{y}_K to generate the trajectories. Both diffusion and reverse diffusion processes are formulated by a Markov chain with Gaussian transitions.

First, we formulate the posterior distribution of the diffusion process from \mathbf{y}_0 to \mathbf{y}_K as:

$$q(\mathbf{y}_{1:K} | \mathbf{y}_0) := \prod_{k=1}^K q(\mathbf{y}_k | \mathbf{y}_{k-1}) \quad (1)$$

$$q(\mathbf{y}_k | \mathbf{y}_{k-1}) := \mathcal{N}(\mathbf{y}_k; \sqrt{1 - \beta_k} \mathbf{y}_{k-1}, \beta_k \mathbf{I}),$$

where $\beta_1, \beta_2, \dots, \beta_K$ are fixed variance schedulers that control the scale of the injected noise. Due to the notable property of the Gaussian transitions, we calculate the diffusion process at any step k in a closed form as:

$$q(\mathbf{y}_k | \mathbf{y}_0) := \mathcal{N}(\mathbf{y}_k; \sqrt{\bar{\alpha}_k} \mathbf{y}_0, (1 - \bar{\alpha}_k) \mathbf{I}), \quad (2)$$

where $\alpha_k = 1 - \beta_k$ and $\bar{\alpha}_k = \prod_{s=1}^k \alpha_s$. Therefore, when K is large enough, we approximately obtain that $\mathbf{y}_K \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. It indicates that the signal is corrupted into a Gaussian noise distribution when gradually adding noise, which conforms to the non-equilibrium thermodynamics phenomenon of diffusion process.

Next, we formulate the trajectories generation process as a reverse diffusion process from noise distribution. We model this reverse process by parameterized Gaussian transitions with the observed trajectories as condition. Given a state feature \mathbf{f} learned by a temporal-social encoder \mathcal{F}_ψ parameterized by ψ with the history trajectories \mathbf{x} as input, we

formulate the reverse diffusion process as:

$$p_\theta(\mathbf{y}_{0:K}|\mathbf{f}) := p(\mathbf{y}_K) \prod_{k=1}^K p_\theta(\mathbf{y}_{k-1}|\mathbf{y}_k, \mathbf{f})$$

$$p_\theta(\mathbf{y}_{k-1}|\mathbf{y}_k, \mathbf{f}) := \mathcal{N}(\mathbf{y}_{k-1}; \boldsymbol{\mu}_\theta(\mathbf{y}_k, k, \mathbf{f}); \boldsymbol{\Sigma}_\theta(\mathbf{y}_k, k)), \quad (3)$$

where $p(\mathbf{y}_K)$ is an initial noise Gaussian distribution, and θ denotes the parameter of the diffusion model. Both parameters of diffusion model θ and encoder network ψ are trained using the trajectory data. Note that we share the network parameters for all transitions. As shown the previous work [14], the variance term of the Gaussian transition can be set as $\boldsymbol{\Sigma}_\theta(\mathbf{y}_K, k) = \sigma_k^2 \mathbf{I} = \beta_k \mathbf{I}$. This setting denotes the upper bound on reverse process entropy for data and shows good performance in practice [40].

3.3. Training Objective

Having formulated diffusion and reverse diffusion processes, we describe how to train the diffusion model. To predict the real trajectory \mathbf{y}_0 , the desired training should optimize the log-likelihood $\mathbb{E}[\log p_\theta(\mathbf{y}_0)]$ in the reverse process. However, the exact log-likelihood is intractable, we thus maximize the variational lower bound for optimization:

$$\mathbb{E}[\log p_\theta(\mathbf{y}_0)] \geq \mathbb{E}_q[\log \frac{p_\theta(\mathbf{y}_{0:K}, \mathbf{f})}{q(\mathbf{y}_{1:K}|\mathbf{y}_0)}]$$

$$= \mathbb{E}_q[\log p(\mathbf{y}_K) + \sum_{k=1}^K \log \frac{p_\theta(\mathbf{y}_{k-1}|\mathbf{y}_k, \mathbf{f})}{q(\mathbf{y}_k|\mathbf{y}_{k-1})}]. \quad (4)$$

We utilize the negative bound as the loss function and perform the training by optimizing it as:

$$L(\theta, \psi) = \mathbb{E}_q[\sum_{k=2}^K D_{KL}(q(\mathbf{y}_{k-1}|\mathbf{y}_k, \mathbf{y}_0) \| p_\theta(\mathbf{y}_{k-1}|\mathbf{y}_k, \mathbf{f})) - \log p_\theta(\mathbf{y}_0|\mathbf{y}_1, \mathbf{f})]. \quad (5)$$

In this loss function, we ignore the term with $\mathbb{E}_q \log p(\mathbf{y}_K)$ in (4), since $p(\mathbf{y}_K)$ is a standard Gaussian and $q(\mathbf{y}_K|\mathbf{y}_0)$ has no learnable parameters as shown in (2).

Here we describe how to calculate the first term D_{KL} . The posterior $q(\mathbf{y}_{k-1}|\mathbf{y}_k, \mathbf{y}_0)$ in D_{KL} is tractable and can be represented by Gaussian distribution as:

$$q(\mathbf{y}_{k-1}|\mathbf{y}_k, \mathbf{y}_0) = \mathcal{N}(\mathbf{y}_{k-1}; \tilde{\boldsymbol{\mu}}_k(\mathbf{y}_k, \mathbf{y}_0), \tilde{\boldsymbol{\beta}}_k \mathbf{I}), \quad (6)$$

where the closed form of $\tilde{\boldsymbol{\mu}}_k(\mathbf{y}_k, \mathbf{y}_0)$ and $\tilde{\boldsymbol{\beta}}_k$ is calculated as:

$$\tilde{\boldsymbol{\mu}}_k(\mathbf{y}_k, \mathbf{y}_0) = \frac{\sqrt{\bar{\alpha}_{k-1}}\beta_k}{1 - \bar{\alpha}_k} \mathbf{y}_0 + \frac{\sqrt{\bar{\alpha}_k}(1 - \bar{\alpha}_{k-1})}{1 - \bar{\alpha}_k} \mathbf{y}_k$$

$$\tilde{\boldsymbol{\beta}}_k = \frac{1 - \bar{\alpha}_{k-1}}{1 - \bar{\alpha}_k} \beta_k \mathbf{I}. \quad (7)$$

Since both diffusion process (6) and reverse process (3) are Gaussian, we can calculate the D_{KL} by the difference be-

tween the means of $\tilde{\boldsymbol{\mu}}_k$ and $\boldsymbol{\mu}_\theta$ as:

$$D_{KL} = \mathbb{E}_q [\lambda \|\tilde{\boldsymbol{\mu}}_k(\mathbf{y}_k, \mathbf{y}_0) - \boldsymbol{\mu}_\theta(\mathbf{y}_k, k, \mathbf{f})\|^2] + C, \quad (8)$$

where λ and C are coefficients with no effect on the gradient direction. Note that the second term $-\log p_\theta(\mathbf{y}_0|\mathbf{y}_1, \mathbf{f})$ can also be formulated as the form in (8) when $k = 1$. Finally, we apply the parameterization method as shown in [14] to reparameterize:

$$\boldsymbol{\mu}_\theta(\mathbf{y}_k, k, \mathbf{f}) = \frac{1}{\sqrt{\bar{\alpha}_k}} (\mathbf{y}_k - \frac{\beta_k}{\sqrt{1 - \bar{\alpha}_k}} \epsilon_\theta(\mathbf{y}_k, k, \mathbf{f})), \quad (9)$$

and obtain a simplified loss function as:

$$L(\theta, \psi) = \mathbb{E}_{\epsilon, \mathbf{y}_0, k} \|\epsilon - \epsilon_{(\theta, \psi)}(\mathbf{y}_k, k, \mathbf{x})\|, \quad (10)$$

where $\epsilon \sim \mathcal{N}(0, \mathbf{I})$, $\mathbf{y}_k = \sqrt{\bar{\alpha}_k} \mathbf{y}_0 + \sqrt{1 - \bar{\alpha}_k} \epsilon$ and the training is performed at each step $k \in 1, 2, \dots, K$. (Please see the detailed derivation and detail algorithms in the supplementary material.)

3.4. Inference

Once the reverse process is trained, we can generate the plausible trajectories by a noise Gaussian $\mathbf{y}_K \sim \mathcal{N}(0, \mathbf{I})$ through the reverse process p_θ . With the reparameterization in (9), we generate the trajectories from \mathbf{y}_K to \mathbf{y}_0 as:

$$\mathbf{y}_{k-1} = \frac{1}{\sqrt{\bar{\alpha}_k}} (\mathbf{y}_k - \frac{\beta_k}{\sqrt{1 - \bar{\alpha}_k}} \epsilon_\theta(\mathbf{y}_k, k, \mathbf{f})) + \sqrt{\beta_k} \mathbf{z}, \quad (11)$$

where \mathbf{z} is a random variable in standard Gaussian distribution and ϵ_θ is the trained network whose inputs include the previous step's prediction \mathbf{y}_k , state embedding \mathbf{f} , and step k .

3.5. Network Architecture

Different from the widely used UNet [36] in image-based diffusion models [8, 14, 31], we design a new Transformer-based network architecture for our MID. With the Transformer, the model can better explore the temporal dependency of paths for the trajectory prediction task. To be specific, MID consists of two key networks: an encoder network with parameters ψ which learns the state embedding by observed history trajectories and their social interactions, and a Transformer-based decoder parameterized by θ for the reverse diffusion process. An overview of the whole architecture is depicted in Figure 2. We will introduce each part in detail in the following.

The encoder network models the history behaviors and social interactions as the state embedding \mathbf{f} . This embedding is fed into the decoder network as the condition of the diffusion model. Note that, designing the network to model social interactions is not the main focus of this work, and MID is an encoder-agnostic framework which can directly equip with different encoders introduced in previous methods. In the experiments, we apply the encoder of Trajec-tron++ [38] for its superior representation ability.

For the decoder, we design a Transformer-based architecture to model the Gaussian transitions in Markov chain. As shown in Figure 2, the inputs of decoder include the ground truth trajectory \mathbf{y}_0 , the noise variable $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, the condition feature \mathbf{f} from the encoder, and a time embedding. In step k , we first add noise into trajectory to get $\mathbf{y}_k = \sqrt{\alpha_k}\mathbf{y}_0 + \sqrt{1 - \alpha_k}\epsilon$. Simultaneously, we calculate the time embedding and concatenate it with the feature of observed trajectory. Then, we apply fully-connected layers to upsample both trajectory \mathbf{y}_k and condition \mathbf{f} , then sum up the outputs as the fused feature. We also introduce the positional embedding in the form of sinusoidal functions on the summation to emphasize the positional relation at different trajectory timestamp t . Finally, the fused feature with positional embedding is fed into the Transformer network to learn the complex spatial-temporal clues. The Transformer-based decoder network consists of three self-attention layers to sufficiently model the temporal dependencies in trajectories, which takes the high dimension sequence as input and outputs the sequence with the same dimension. With a fully-connected layer, we downsample the output sequence to the trajectory dimension. We finally perform mean square error (MSE) loss between the output and a random Gaussian as (10) for current iteration to optimize the network. Please see the network details in the supplementary material.

4. Experiments

In this section, we first compared the proposed method with state-of-the-art approaches on two widely-used pedestrian trajectories prediction benchmarks, then conducted ablation studies to analyze the effectiveness of key components of our MID framework and provided an analysis regarding the reverse diffusion process.

4.1. Experimental Setup

Datasets: We evaluated our method on two public pedestrian trajectories forecasting benchmarks including Stanford Drone Dataset (SDD) [35] and UCY/ETH [21, 33].

Stanford Drone Dataset: Stanford Drone Dataset [35] is a well established benchmark for human trajectory prediction in bird’s eye view. The dataset consists of 20 scenes captured using a drone in top down view around the university campus containing several moving agents like humans and vehicles.

ETH/UCY: The ETH [33] and UCY [21] dataset group consists of five different scenes – ETH & HOTEL (from ETH) and UNIV, ZARA1, & ZARA2 (from UCY). All the scenes report the position of pedestrians in world-coordinates and hence the results we report are in metres. The scenes are captured in unconstrained environments with few objects blocking pedestrian paths.

Table 1. Quantitative results on the Stanford Drone dataset with Best-of-20 strategy in ADE/FDE metric. “T” denotes the method only using the trajectory position information, and “T + I” denotes the method using both position and visual image information. † means the results are reproduced by us with the official released code. Lower is better.

Methods	Input	Sampling	ADE	FDE
CGNS [22]	T + I	20	15.60	28.20
SimAug [23]	T + I	20	10.27	19.71
†Y-Net [28]	T + I	20	8.97	14.61
Y-Net [28]+ TTST	T + I	10000	7.85	11.85
Social-GAN [12]	T	20	27.23	41.44
PECNet [29]	T	20	9.96	15.88
†Trajectron++ [38]	T	20	8.98	19.02
LB-EBM [32]	T	20	8.87	15.61
PCCSNET [45]	T	20	8.62	16.16
†Expert [53]	T	20	10.67	14.38
† Expert [53]+GMM	T	20×20	7.65	14.38
MID	T	20	7.61	14.30

Evaluation Metric: We adopted the widely-used evaluation metrics Average Displacement Error (ADE) and Final Displacement Error (FDE). ADE computes the average error between all the ground truth positions and the estimated positions in the trajectory, and FDE computes the displacement between the end points of ground truth and predicted trajectories. The trajectories are sampled 0.4 seconds interval, where the first 3.2 seconds of a trajectory is used as observed data to predict the next 4.8 seconds future trajectory. For the ETH/UCY dataset, we followed the leave one out cross-validation evaluation strategy such that we trained our model on four scenes and tested on the remaining one [12, 15, 19, 38]. Considering the stochastic property of our method, we used Best-of-N strategy to compute the final ADE and FDE with $N = 20$.

Implementation Details: We devised a three-layers Transformer as the core network for our MID, where the Transformer dimension is set to 512, and 4 attention heads are applied. We employed one fully-connected layer to upsample the input of the model from dimension 2 to the Transformer dimension, and another fully-connected layer to upsample the observed trajectory feature \mathbf{f} to the same dimension. We utilized three fully-connected layers to progressively downsample the Transformer output sequence to the predicted trajectory, such that 512d-256d-2d. The training was performed with Adam optimizer, with a learning rate of 0.001 and batch size of 256. All the experiments were conducted on a single Tesla V100 GPU.

4.2. Comparison with state-of-the-art methods

We quantitatively compare our method with a wide range of current methods. As shown in Table 1, we provide

Table 2. Quantitative results on the ETH/UCY dataset with Best-of-20 strategy in ADE/FDE metric. Lower is better.

	Input	Sampling	ETH		HOTEL		UNIV		ZARA1		ZARA2		AVG	
			ADE	FDE	ADE	FDE	ADE	FDE	ADE	FDE	ADE	FDE	ADE	FDE
SoPhie [37]	T + I	20	0.70	1.43	0.76	1.67	0.54	1.24	0.30	0.63	0.38	0.78	0.54	1.15
CGNS [22]	T + I	20	0.62	1.40	0.70	0.93	0.48	1.22	0.32	0.59	0.35	0.71	0.49	0.97
Social-BiGAT [19]	T + I	20	0.69	1.29	0.49	1.01	0.55	1.32	0.30	0.62	0.36	0.75	0.48	1.00
MG-GAN [6]	T + I	20	0.47	0.91	0.14	0.24	0.54	1.07	0.36	0.73	0.29	0.60	0.36	0.71
Y-Net [28] + TTST	T + I	10000	0.28	0.33	0.10	0.14	0.24	0.41	0.17	0.27	0.13	0.22	0.18	0.27
Social-GAN [12]	T	20	0.81	1.52	0.72	1.61	0.60	1.26	0.34	0.69	0.42	0.84	0.58	1.18
Causal-STGCNN [2]	T	20	0.64	1.00	0.38	0.45	0.49	0.81	0.34	0.53	0.32	0.49	0.43	0.66
PECNet [29]	T	20	0.54	0.87	0.18	0.24	0.35	0.60	0.22	0.39	0.17	0.30	0.29	0.48
STAR [49]	T	20	0.36	0.65	0.17	0.36	0.31	0.62	0.26	0.55	0.22	0.46	0.26	0.53
Trajectron++ [38]	T	20	0.39	0.83	0.12	0.21	0.20	0.44	0.15	0.33	0.11	0.25	0.19	0.41
LB-EBM [32]	T	20	0.30	0.52	0.13	0.20	0.27	0.52	0.20	0.37	0.15	0.29	0.21	0.38
PCCSNET [45]	T	20	0.28	0.54	0.11	0.19	0.29	0.60	0.21	0.44	0.15	0.34	0.21	0.42
†Expert [53]	T	20	0.37	0.65	0.11	0.15	0.20	0.44	0.15	0.31	0.12	0.26	0.19	0.36
†Expert [53]+GMM	T	20×20	0.29	0.65	0.08	0.15	0.15	0.44	0.11	0.31	0.09	0.26	0.14	0.36
MID	T	20	0.39	0.66	0.13	0.22	0.22	0.45	0.17	0.30	0.13	0.27	0.21	0.38

Table 3. Ablation studies on MID and network architecture designs. Trans is the abbreviation of Transformer

Group	Method	Architecture	ADE	FDE
1	MID	Trans-512d	7.61	14.30
2	MID	Trans-256d	7.91	14.50
	MID	Trans-1024d	7.64	14.37
3	MID	Linear	8.85	17.25
	MID	LSTM	8.41	16.57
4	Trajectron++	LSTM	8.98	19.02
	Trajectron++	Trans-256d	9.86	19.56

the comparison between our method and existing methods on the Stanford Drone dataset. We categorize methods as Trajectory-Only methods (T) and Trajectory-and-Image (T+I) methods, as the additional image information may be crucial at certain circumstances yet increases computation cost. Besides, we also report the sampling number since adding the sampling number can effectively promote the performance. We provide the results under standard 20 samplings of MID and other methods for a fair comparison. We observe that our method achieves an average ADE/FDE of 7.61/14.30 in pixel coordinate, which achieves the best performance among all the current methods, regardless of the involvement of image data. Specifically, our MID outperforms the current state-of-the-art T+I method Y-Net+TTST on the ADE metric. Note that our method did not use the image data and apply any post-processing such as the Test-Time Sampling Trick (TTST) [28]. We provide the results with the sampling trick in the supplementary.

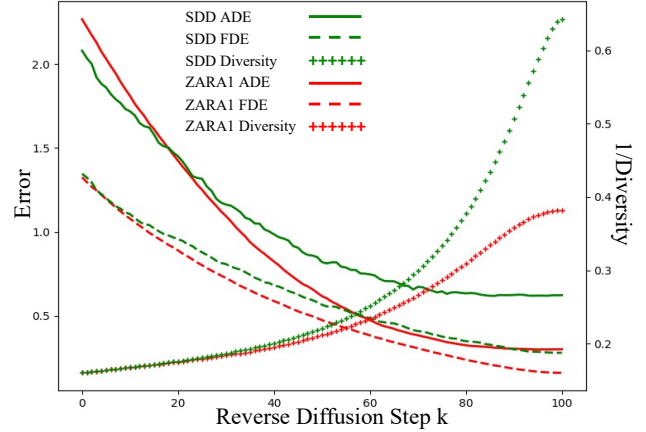


Figure 3. Trade-off between determinacy (ADE/FDE) and diversity within reverse diffusion steps from 0 to 100.

We also conducted experiments on the ETH-UCY dataset and tabulated the results in Table 2. Our method achieves a comparable performance with only trajectory input in 20 sampling, with an average performance of 0.21 ADE and 0.38 FDE. We found that MID are benefited more on the larger dataset (e.g. the SDD dataset).

4.3. Ablation Studies

In this subsection, we conducted ablation studies to investigate the effectiveness of each key component including diffusion model and Transformer architecture. Then, we provided a detailed analysis for reverse diffusion process.

Diffusion Model: In order to examine the importance of our diffusion model, we degraded our MID into a CVAE

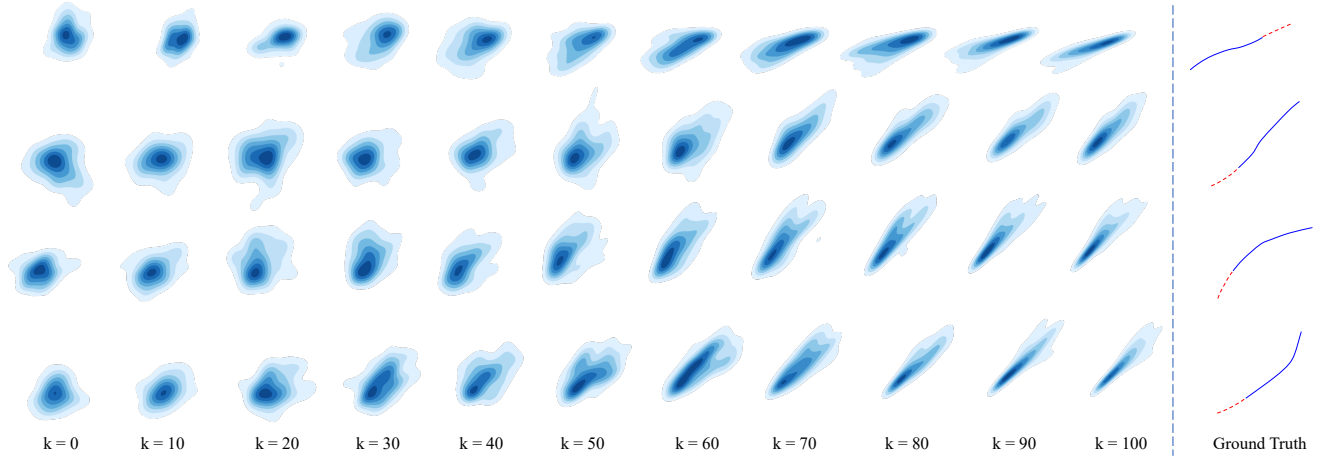


Figure 4. Visualization of generated trajectories at each diffusion time step t . We can see that the reverse diffusion process progressively reduces the indeterminacy and reaches the desired trajectory through the time step. Starting from a normal distribution at $t = 0$ that corresponds to all walkable areas, and the observed paths (red dash), our MID method successfully eliminates improbable trajectories and gradually fits the ground truth future paths (blue line). Best viewed in color.

based framework, Trajectron++. We replaced the decoder from commonly-used LSTM to our Transformer in this CVAE based framework to verify whether the performance boost comes from the Transformer. The group 2 and 4 in Table 3 show the performance comparison. We observe that using the same encoder and decoder but without our diffusion model, the results degrade significantly, demonstrating the effectiveness of our diffusion model. Besides, only replacing the decoder with our Transformer architecture in the CVAE based framework does not improve performance as shown in group 4 in Table 3.

Transformer Architecture: We also conducted experiments on the decoder architecture of MID. According to group 1 and 3 in Table 3, Transformer outperforms the Linear and LSTM architecture by a large margin. It indicates the Transformer architecture is effective for MID to model the temporal dependencies of trajectory. Besides, we evaluated the Transformer architectures with different dimensions. As tabulated in Table 3 group 1 and 2, we observe that the Transformer with 512 dimensions leads to the best performance, and further increasing the Transformer dimension or model parameters does not yield better results.

Analysis of Reverse Diffusion Process: To further explore the reverse diffusion process, we generated 20 trajectories at each reverse diffusion step and analyzed the gradual change of the distribution. We provide an analysis between the reverse diffusion step and the corresponding diversity and ADE/FDE, as illustrated in Figure 3. The trajectory diversity is calculated as the average of Euclidean distance between any of the two in the generated 20 trajectories. When the reverse diffusion step is small, the trajectory distribution is more indeterminate and produces highly diverse trajectories. As the reverse diffusion step increases,

we observe the decline in diversity and the rise of determinacy. With our MID framework, we can control the degree of indeterminacy by adjusting the step numbers, and achieve a flexible trade-off between the diversity and determinacy of the generated trajectories.

In addition, we visualize the distribution of trajectories as contours in Figure 4 and each contour map is sampled by ten steps interval. We see that the contours are diverse at the early stage of the diffusion process, and deform gradually to be more concentrated and fit to the ground truth trajectory.

4.4. Qualitative Evaluation

We further investigated the ability of our framework by the qualitative results. Figure 5 illustrates the most-likely predictions of our MID and Trajectron++ [38] on all five scenes on the ETH/UCY dataset. The qualitative results show that both MID and Trajectron++ fits the ground truth paths well. We observe that Trajectron++ performs similarly to MID for short-term forecasting yet a little deviates from the ground truth path for longer prediction. Besides, we visualize multiple predicted trajectories on SDD in Figure 6. We observe that all predictions show their feasibility conditioned on the observed trajectories. Though reducing the ambiguity with the reverse diffusion model, We found that generated trajectories are still full of the diversity in a walkable region.

5. Conclusion & Discussion

In this paper, we introduced a new MID framework to formulate trajectory prediction with motion indeterminacy diffusion. In this framework, we learned a parameterized Markov chain conditioned on the observed trajectories to gradually discard the indeterminacy from ambiguous areas

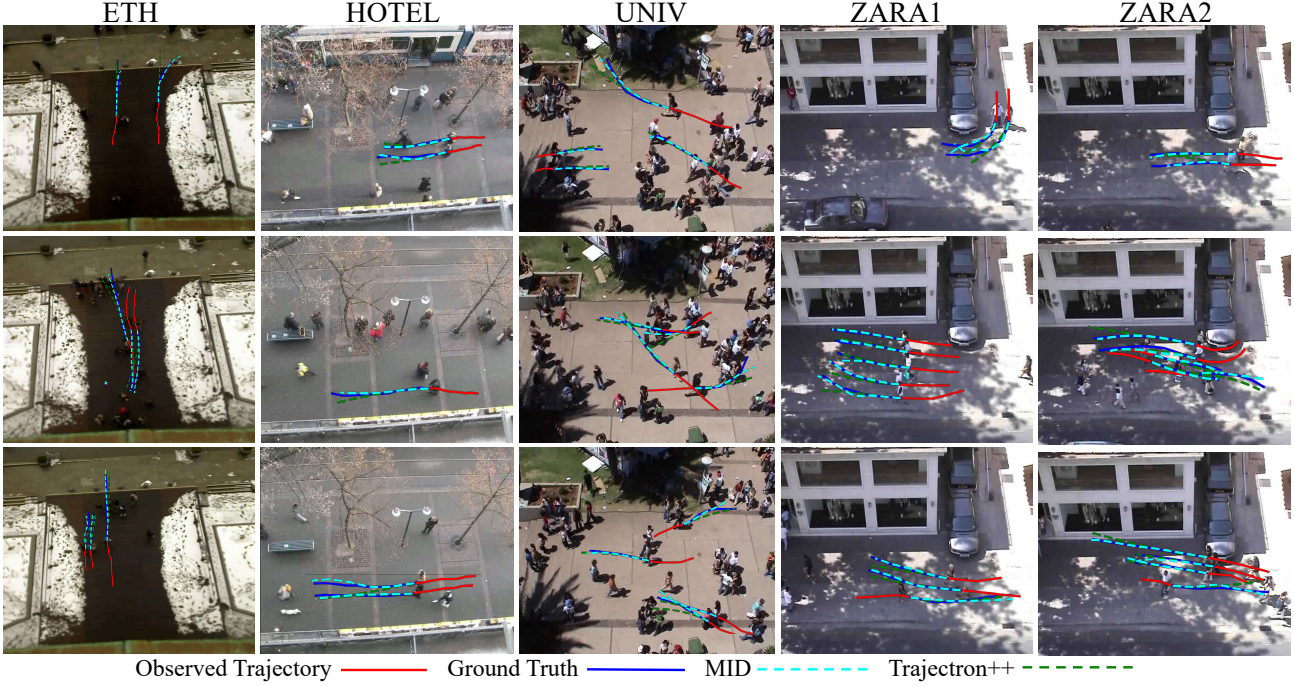


Figure 5. Visualization of predicted trajectories on the ETH/UCY Dataset. Given the observed trajectories (red), we illustrate the ground truth paths (blue) and predicted trajectories by MID (dashed cyan) and Trajectron++ (green) for five different scenes. We see that our results are much closer to the ground truth compared with Trajectron++. Best viewed in color and zoom-in for more clarity.




Figure 6. Visualization of generated trajectories in the Stanford Drone Dataset. Given the observed trajectories (red), we illustrate the ground truth paths (blue) and predicted best-of-20 trajectories by MID (dashed cyan) in different scenes. The blue line is covered by the cyan dashes in all scenes. Best viewed in color and zoom-in for more clarity.

to acceptable trajectories. By adjusting the length of the chain, we can achieve the trade-off between diversity and determinacy. Besides, we designed a Transformer-based architecture as the core network of our method to model complex temporal dependency in trajectories. Experimental results demonstrate the superiority of our method which achieves state-of-the-art performance on the Stanford Drone and ETH/UCY benchmarks.

Broader Impact: MID could be applied to a wide range of applications with human-robots interaction. With indeterminacy modeling, we can generate accurate and reasonable future trajectories, which helps much with decision making in auto-driving. Besides, MID can adjust the degree of indeterminacy, which has the potential to be applied in dynamic and interactive environments.

Limitations: Despite the promising performance and an applicable trade-off nature, the time cost at reverse diffusion process could be expensive due to multiple steps (100 steps in our experiments). When evaluated with 512 trajectories on the ZARA1 dataset, Trajectron++ needs 0.443s but MID will need 17.368s with 100 diffusion steps setting. Fortunately, many recent efforts have been made to significantly reduce the sampling cost while keeping the high generation performance [17, 31, 39, 42, 48]. However, plugging these methods in our MID is not trivial. We leave it as future work to build a more efficient system.

Acknowledgments This work was supported in part by the National Natural Science Foundation of China under Grant 62125603, and Grant U1813218, in part by a grant from the Beijing Academy of Artificial Intelligence (BAAI).

References

- [1] Alexandre Alahi, Kratharth 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](#)
- [2] Guangyi Chen, Junlong Li, Jiwen Lu, and Jie Zhou. Human trajectory prediction via counterfactual analysis. In *ICCV*, pages 9824–9833, 2021. [2](#), [6](#)
- [3] Guangyi Chen, Junlong Li, Nuoxing Zhou, Liangliang Ren, and Jiwen Lu. Personalized trajectory prediction via distribution discrimination. In *ICCV*, pages 15580–15589, 2021. [1](#), [2](#)
- [4] Nanxin Chen, Yu Zhang, Heiga Zen, Ron J Weiss, Mohammad Norouzi, and William Chan. Wavegrad: Estimating gradients for waveform generation. In *ICLR*, 2021. [2](#)
- [5] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. In *ICCV*, 2021. [2](#)
- [6] Patrick Dendorfer, Sven Elflein, and Laura Leal-Taixé. Mrgan: A multi-generator model preventing out-of-distribution samples in pedestrian trajectory prediction. In *ICCV*, pages 13158–13167, 2021. [1](#), [2](#), [6](#)
- [7] Nachiket Deo and Mohan M Trivedi. Trajectory forecasts in unknown environments conditioned on grid-based plans. *arXiv preprint arXiv:2001.00735*, 2020. [2](#)
- [8] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, 2021. [2](#), [4](#)
- [9] Liangji Fang, Qinhong Jiang, Jianping Shi, and Bolei Zhou. Tpnnet: Trajectory proposal network for motion prediction. In *CVPR*, 2020. [1](#), [2](#)
- [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. [2](#)
- [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. [2](#)
- [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. [1](#), [2](#), [5](#), [6](#)
- [13] Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. *Physical Review E*, 51(5):4282, 1995. [2](#)
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. [2](#), [4](#)
- [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. [2](#), [5](#)
- [16] Boris Ivanovic and Marco Pavone. The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs. In *ICCV*, pages 2375–2384, 2019. [1](#), [2](#)
- [17] Alexia Jolicœur-Martineau, Ke Li, Rémi Piché-Taillefer, Tal Kachman, and Ioannis Mitliagkas. Gotta go fast when generating data with score-based models. *arXiv preprint arXiv:2105.14080*, 2021. [8](#)
- [18] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. In *ICLR*, 2021. [2](#)
- [19] 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. [1](#), [2](#), [5](#), [6](#)
- [20] 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. [1](#), [2](#)
- [21] Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. Crowds by example. In *Computer Graphics Forum*, volume 26, pages 655–664, 2007. [5](#)
- [22] Jiachen Li, Hengbo Ma, and Masayoshi Tomizuka. Conditional generative neural system for probabilistic trajectory prediction. In *IROS*, pages 6150–6156. IEEE, 2019. [5](#), [6](#)
- [23] Junwei Liang, Lu Jiang, and Alexander Hauptmann. Simaug: Learning robust representations from simulation for trajectory prediction. In *ECCV*, pages 275–292, 2020. [2](#), [5](#)
- [24] Junwei Liang, Lu Jiang, Kevin Murphy, Ting Yu, and Alexander Hauptmann. The garden of forking paths: Towards multi-future trajectory prediction. In *CVPR*, pages 10508–10518, 2020. [2](#)
- [25] Yuejiang Liu, Qi Yan, and Alexandre Alahi. Social nce: Contrastive learning of socially-aware motion representations. In *ICCV*, pages 15118–15129, 2021. [1](#), [2](#)
- [26] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *CVPR*, pages 2837–2845, 2021. [2](#)
- [27] Osama Makansi, Julius von Kügelgen, Francesco Locatello, Peter Gehler, Dominik Janzing, Thomas Brox, and Bernhard Schölkopf. You mostly walk alone: Analyzing feature attribution in trajectory prediction. *arXiv preprint arXiv:2110.05304*, 2021. [2](#)
- [28] Karttikeya Mangalam, Yang An, Harshayu Girase, and Jitendra Malik. From goals, waypoints & paths to long term human trajectory forecasting. In *ICCV*, pages 15233–15242, 2021. [1](#), [2](#), [5](#), [6](#)
- [29] 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. [1](#), [2](#), [5](#), [6](#)
- [30] Abdullah Mohamed, Kun Qian, Mohamed Elhoseiny, and Christian Claudel. Social-stgcnn: A social spatio-temporal graph convolutional neural network for human trajectory prediction. In *CVPR*, 2020. [2](#)
- [31] Alex Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *ICML*, 2021. [2](#), [4](#), [8](#)
- [32] Bo Pang, Tianyang Zhao, Xu Xie, and Ying Nian Wu. Trajectory prediction with latent belief energy-based model. In *CVPR*, pages 11814–11824, 2021. [1](#), [5](#), [6](#)
- [33] 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. [5](#)

- [34] Vadim Popov, Ivan Vovk, Vladimir Gogoryan, Tasnima Sadekova, and Mikhail Kudinov. Grad-tts: A diffusion probabilistic model for text-to-speech. In *ICML*, 2021. 2
- [35] Alexandre Robicquet, Amir Sadeghian, Alexandre Alahi, and Silvio Savarese. Learning social etiquette: Human trajectory understanding in crowded scenes. In *ECCV*, pages 549–565, 2016. 5
- [36] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015. 4
- [37] Amir Sadeghian, Vineet Kosaraju, Ali Sadeghian, Noriaki Hirose, Hamid Rezatofighi, and Silvio Savarese. Sophie: An attentive gan for predicting paths compliant to social and physical constraints. In *CVPR*, pages 1349–1358, 2019. 1, 2, 6
- [38] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *ECCV*, 2020. 1, 2, 4, 5, 6, 7
- [39] Robin San-Roman, Eliya Nachmani, and Lior Wolf. Noise estimation for generative diffusion models. *arXiv preprint arXiv:2104.02600*, 2021. 8
- [40] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, pages 2256–2265, 2015. 2, 4
- [41] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *NeurIPS*, pages 3483–3491, 2015. 2
- [42] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*, 2021. 8
- [43] Hao Sun, Zhiqun Zhao, and Zhihai He. Reciprocal learning networks for human trajectory prediction. In *CVPR*, 2020. 1, 2
- [44] Jianhua Sun, Qinhong Jiang, and Cewu Lu. Recursive social behavior graph for trajectory prediction. In *CVPR*, June 2020. 2
- [45] Jianhua Sun, Yuxuan Li, Hao-Shu Fang, and Cewu Lu. Three steps to multimodal trajectory prediction: Modality clustering, classification and synthesis. In *ICCV*, 2021. 1, 5, 6
- [46] Charlie Tang and Russ R Salakhutdinov. Multiple futures prediction. In *NeurIPS*, pages 15398–15408, 2019. 1, 2
- [47] Anirudh Vemula, Katharina Muelling, and Jean Oh. Social attention: Modeling attention in human crowds. In *ICRA*, pages 1–7, 2018. 2
- [48] Daniel Watson, Jonathan Ho, Mohammad Norouzi, and William Chan. Learning to efficiently sample from diffusion probabilistic models. *arXiv preprint arXiv:2106.03802*, 2021. 8
- [49] Cunjun Yu, Xiao Ma, Jiawei Ren, Haiyu Zhao, and Shuai Yi. Spatio-temporal graph transformer networks for pedestrian trajectory prediction. In *ECCV*, August 2020. 1, 6
- [50] Cunjun Yu, Xiao Ma, Jiawei Ren, Haiyu Zhao, and Shuai Yi. Spatio-temporal graph transformer networks for pedestrian trajectory prediction. In *ECCV*, 2020. 2
- [51] 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. 2
- [52] Hang Zhao, Jiyang Gao, Tian Lan, Chen Sun, Benjamin Sapp, Balakrishnan Varadarajan, Yue Shen, Yi Shen, Yuning Chai, Cordelia Schmid, et al. Tnt: Target-driven trajectory prediction. 2020. 2
- [53] He Zhao and Richard P Wildes. Where are you heading? dynamic trajectory prediction with expert goal examples. In *ICCV*, pages 7629–7638, 2021. 1, 2, 5, 6
- [54] 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. 1, 2
- [55] Linqi Zhou, Yilun Du, and Jiajun Wu. 3d shape generation and completion through point-voxel diffusion. In *ICCV*, 2021. 2