

# Human Joint Kinematics Diffusion-Refinement for Stochastic Motion Prediction

Dong Wei<sup>1</sup>, Huaijiang Sun<sup>1\*</sup>, Bin Li<sup>2</sup>, Jianfeng Lu<sup>1</sup>, Weiqing Li<sup>1</sup>, Xiaoning Sun<sup>1</sup>, Shengxiang Hu<sup>1</sup>

<sup>1</sup>School of Computer Science and Engineering, Nanjing University of Science and Technology, Nanjing, China

<sup>2</sup>Tianjin AiForward Science and Technology Co., Ltd., Tianjin, China

{csdwei, sunhuaijiang, lujf, li\_weiqing, sunxiaoning, hushengxiang}@njust.edu.cn, libin@aiforward.com

## Abstract

Stochastic human motion prediction aims to forecast multiple plausible future motions given a single pose sequence from the past. Most previous works focus on designing elaborate losses to improve the accuracy, while the diversity is typically characterized by randomly sampling a set of latent variables from the latent prior, which is then decoded into possible motions. This joint training of sampling and decoding, however, suffers from posterior collapse as the learned latent variables tend to be ignored by a strong decoder, leading to limited diversity. Alternatively, inspired by the diffusion process in nonequilibrium thermodynamics, we propose MotionDiff, a diffusion probabilistic model to treat the kinematics of human joints as heated particles, which will diffuse from original states to a noise distribution. This process not only offers a natural way to obtain the “whitened” latents without any trainable parameters, but also introduces a new noise in each diffusion step, both of which facilitate more diverse motions. Human motion prediction is then regarded as the reverse diffusion process that converts the noise distribution into realistic future motions conditioned on the observed sequence. Specifically, MotionDiff consists of two parts: a spatial-temporal transformer-based diffusion network to generate diverse yet plausible motions, and a flexible refinement network to further enable geometric losses and align with the ground truth. Experimental results on two datasets demonstrate that our model yields the competitive performance in terms of both diversity and accuracy.

## Introduction

Human Motion Prediction (HMP) has received increasing attention due to its broad applications such as human-robot interaction (Bajcsy et al. 2021), autonomous driving (Paden et al. 2016) and animation production (Park et al. 2019). The ability to perform such predictions allows robots to understand the future plans of human beings, which is critical to cooperate safely and reasonably with people. While encouraging results have been achieved in previous works (Aksan, Kaufmann, and Hilliges 2019; Mao et al. 2019; Zhong et al. 2022), they neglect the fact that uncertainty and stochasticity are intrinsic properties of human motions. Given a single past observation, predicting multiple possible future se-

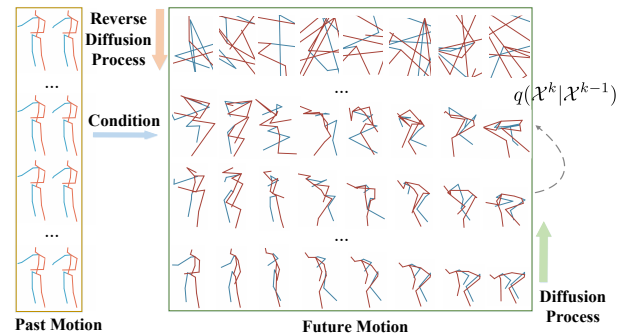


Figure 1: Visualization of the diffusion and reverse process of MotionDiff. For diffusion process, a new noise is gradually incorporated until the kinematic information is completely destroyed. In contrast, the reverse diffusion process recovers the desired realistic future motion from noisy distribution conditioned on the observation via a Markov chain.

quences rather than only one output is gaining in popularity. The latter, i.e., deterministic HMP, which is mostly based on recurrent neural network or graph convolutional network, cannot capture such stochastic behaviors. How to generate accurate human motion predictions and at the same time fully consider the diversity remains a challenging problem.

Recently, deep generative networks have made significant progress in modeling the multi-modal data distribution (Barsoom, Kender, and Liu 2018; Yan et al. 2018; Aliakbarian et al. 2020), such as Generative Adversarial Network (GAN) and Variational AutoEncoder (VAE). Most of them obtain diversity by randomly sampling a set of latent variables from the latent prior, which requires additional neural networks for training (i.e., the discriminator in GAN or the sampling encoder in VAE). This process, however, will bring about training instability or posterior collapse when jointly trained with a powerful decoder (McCarthy et al. 2020). Unfortunately, in the particular case of human motion prediction, a sufficiently high-capacity decoder is indispensable to keep the predictions physically plausible. As a consequence, such decoder tends to model the conditional density directly, giving the network possibility to learn to ignore the stochastic latent variables, and thus limiting the diversity of future motions. To increase the diversity, recent progress on stochastic

\*Corresponding author.

human motion prediction (Aliakbarian et al. 2020; Yuan and Kitani 2020; Zhang, Black, and Tang 2021; Ma et al. 2022) add constraints such as stochastic conditioning schemes or new losses, to force the model to take the noise into account. While these methods indeed yield high diversity, they still suffer from the above inherent limitation. An ideal solution would be directly starting from the latents without training to circumvent this problem, and using appropriate constraints to explore more plausible motions.

In this paper, we propose such a solution named MotionDiff based on Denoising Diffusion Probabilistic Models (DDPM). As a future human pose sequence is composed of a set of 3D joint locations that satisfy the kinematic constraints, we regard these locations as particles in a thermodynamics system in contact with a heat bath. In this light, the particles evolve stochastically in the way that they progressively diffuse from the original states (i.e., kinematics of human joints) to a noise distribution (i.e., chaotic positions). This offers an alternative way to obtain the “whitened” latents without any training process, which naturally avoids posterior collapse. Meanwhile, contrary to previous methods that require extra sampling encoders to obtain diversity, a unique strength of MotionDiff is that it is inherently diverse because the diffusion process is implemented by incorporating a new noise to human motion at each time step. Our high-level idea is to learn the reverse diffusion process, which recovers the target realistic pose sequences from the noisy distribution conditioned on the observed past motion (see Figure 1). This process can be formulated as a Markov chain, and allows us to use a simple mean squared error loss function to optimize the variational lower bound. Nonetheless, directly extending the diffusion model to stochastic human motion prediction results in two key challenges that arise from the following observations: First, since the kinematic information between local joint coordinates has been completely destroyed in the diffusion process, and a certain number of steps in the reverse diffusion process is required, it is necessary to devise an expressive yet efficient decoder to construct such relations; Second, as we do not explicitly guide the future motion generation with any loss, MotionDiff produces realistic predictions that are totally different from the ground truth, which makes the quantitative evaluation challenging (Lugmayr et al. 2022).

To this end, we elaborately design a transformer-based architecture as the core decoder of MotionDiff to tackle the first problem. Instead of performing simple pose embedding (Bouazizi et al. 2022), we devise a spatial transformer module to encode joint embedding, in which local relationships between the 3D joints in each frame can be better investigated. Following, we capture the global dependencies across frames by a temporal transformer module. This separable architecture differs from the autoregressive model (Aksan et al. 2021) that interleaves spatial and temporal modeling with tremendous computations. For the second issue, we further employ a deterministic feed-forward HMP model to refine multiple pose sequences generated from our diffusion network with the help of the observed sequence. Our diffusion-refinement paradigm enables explicit geometric losses, and enjoys a significant approximation to

the ground truth while still keeping diverse and realistic. The contributions of our work are summarized as follows:

- We propose a novel stochastic human motion prediction framework with human joint kinematics diffusion-refinement, which incorporates a new noise at each diffusion step to get inherent diversity.
- We design a spatial-temporal transformer-based architecture for the proposed framework to encode local kinematic information in each frame as well as global temporal dependencies across frames.
- Extensive experiments show that our model achieves state-of-the-art performance on both Human3.6M and HumanEva-I datasets.

## Related Work

### Deterministic Human Motion Prediction

Given the observed past motion, deterministic HMP aims at producing only one output, and thus can be regarded as a regression task. Most existing methods (Fragkiadaki et al. 2015a; Martinez, Black, and Romero 2017; Liu et al. 2022) exploit Recurrent Neural Networks (RNN) to address this problem due to its superiority in modeling sequential data. However, these methods usually suffer from limitations of first-frame discontinuity and error accumulation, especially for long-term prediction. Recent works (Mao et al. 2019; Li et al. 2020; Dang et al. 2021; Sun et al. 2022) propose Graph Convolutional Networks (GCN) to model the joint dependencies of human motion. Motivated by the significant success of Transformer (Vaswani et al. 2017), (Cai et al. 2020) adapt it on the discrete cosine transform coefficients extracted from the observed motion. To learn more desired representations, (Aksan et al. 2021) propose to aggregate spatial and temporal information directly from the data by leveraging the recursive nature of human motion. However, this autoregressive and computationally heavy design is not appropriate for the decoder of MotionDiff because diffusion model is non-autoregressive and requires a certain number of reverse diffusion steps (i.e., decoding), which are valuable to generate high quality predictions. Therefore, we present an efficient architecture that separates spatial and temporal information like (Sofianos et al. 2021; Zheng et al. 2021).

### Stochastic Human Motion Prediction

Due to the diversity of human behaviors, many stochastic HMP methods are proposed to model the multi-modal data distribution. These methods are mainly based on deep generative models (Walker et al. 2017; Kundu, Gor, and Babu 2019; Mao, Liu, and Salzmann 2021; Ma et al. 2022), such as GAN (Goodfellow et al. 2014) and VAE (Kingma and Welling 2013). For GAN-based methods, (Barsoum, Kender, and Liu 2018) develop a HP-GAN framework that models the diversity by combining a random vector with the embedding state at the test time; (Kundu, Gor, and Babu 2019) exploit the discriminator to regress the random vector for diversity. However, these methods involve complex adversarial learning between the generator and discriminator, resulting in instable training. For VAE-based methods,

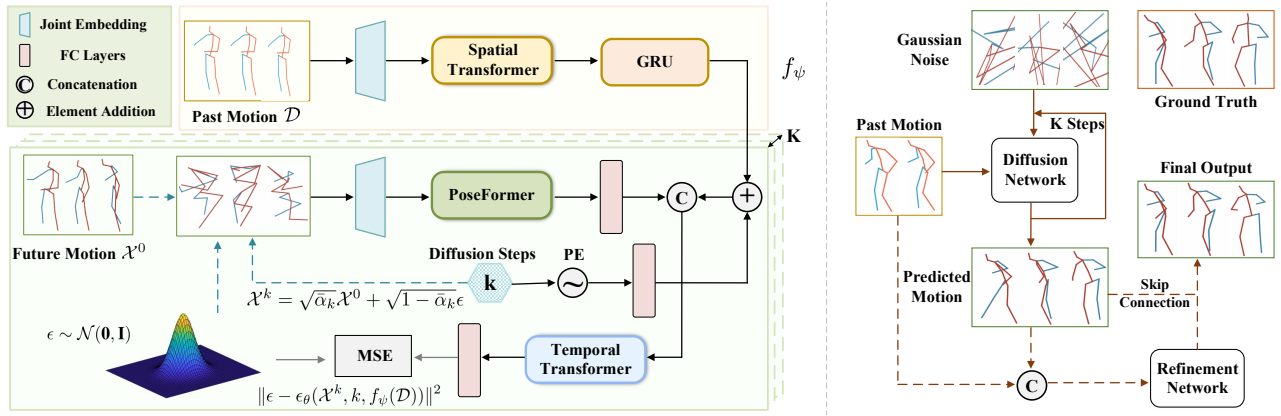


Figure 2: Overview of MotionDiff for stochastic human motion prediction. Left: the architecture of the diffusion network. The model  $\epsilon_\theta$  is fed a state embedding of historical information  $f_\psi(\mathcal{D})$ , diffusion step  $k$ , and a motion sequence  $\mathcal{X}^k$  in a noising step  $k$ . Right: the diffusion-refinement paradigm. It gradually removes Gaussian noise through the diffusion network ( $K$  steps) to get clean future motions, and then refines the predictions to align with the ground truth. Arrows indicate different stages: blue dotted for diffusion, black solid for training, red solid for sampling, and red dotted for refinement.

although such likelihood methods can have a good estimation of the data, they require additional networks to sample a set of latent variables and fail to sample some minor modes. To alleviate this problem, (Aliakbarian et al. 2020) take the noise into account with a mix-and-match perturbation strategy, while (Yuan and Kitani 2020) propose Diversifying Latent Flows (DLow) with new losses to produce diverse samples from a pretrained VAE model, both of which still suffer from the inherent limitation of posterior collapse, and thus limited diversity of predictions. In contrast, we adopt a more powerful diffusion model, and investigate the diversity by adding a new noise to human motion at each diffusion step.

### Denoising Diffusion Probabilistic Models

Denoising Diffusion Probabilistic Models (DDPM) (Sohl-Dickstein et al. 2015; Ho, Jain, and Abbeel 2020) have achieved significant success recently in various applications, such as image generation (Dhariwal and Nichol 2021; Lugmayr et al. 2022), audio synthesis (Chen et al. 2021; Kong et al. 2021) and trajectory prediction (Gu et al. 2022). The diffusion models leverage a Markov chain to progressively convert the noise distribution into the data distribution. In this paper, we draw inspiration from diffusion and its reverse process, to naturally avoid posterior collapse and therefore generate truly diverse and plausible future motions. As far as we know, this is the first work to introduce the diffusion model into stochastic human motion prediction. Instead of UNet backbone (Ronneberger, Fischer, and Brox 2015) for conventional diffusion models, we design a spatial-temporal transformer-based architecture as the core decoder, and further refine the generated samples from our diffusion network, enlightened by DLow and (Lyu et al. 2022).

### Proposed Approach

In this section, we formulate stochastic human motion prediction as the reverse diffusion process. We then introduce the objective function to train the model using variational

inference and sampling process. Finally, we describe the detailed network architecture shown in Figure 2.

### Formulation

We represent the observed historical sequence as  $\mathcal{D} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$  with  $T$  frames, and the predicted sequence as  $\mathcal{X} = \{\mathbf{x}_{T+1}, \mathbf{x}_{T+2}, \dots, \mathbf{x}_{T+f}\}$  with  $f$  frames, where  $\mathbf{x}_t \in \mathbb{R}^{3 \times J}$  is the 3D coordinates at timestamp  $t$  and  $J$  is the number of body joints. Our goal is to employ diffusion model to predict possible and realistic future motions. Intuitively, the diffusion process progressively incorporates Gaussian noise to the clean motion until it is totally transformed into a whitened latent state, while the reverse process starts from this state and gradually recovers the desired future motions conditioned on the historical sequence.

**Diffusion Process.** As illustrated in Figure 1, with time going by, the target future motion  $\mathcal{X}^0$  with kinematic constraints will progressively diffuse into the next chaotic positions  $\mathcal{X}^k$ , and finally converge into a whitened noise  $\mathcal{X}^K$ , where  $K$  is the maximum number of diffusion steps. Unlike conventional stochastic HMP methods that require additional sampling encoder to obtain diversity, our key observation is that the inherent diversity of the proposed model relies on the diffusion process ( $\mathcal{X}^0, \mathcal{X}^1, \dots, \mathcal{X}^K$ ) that is defined as a fixed (rather than trainable) posterior distribution  $q(\mathcal{X}^{1:K} | \mathcal{X}^0)$ . More precisely, we formulate this process as:

$$q(\mathcal{X}^{1:K} | \mathcal{X}^0) = \prod_{k=1}^K q(\mathcal{X}^k | \mathcal{X}^{k-1}), \quad (1)$$

$$q(\mathcal{X}^k | \mathcal{X}^{k-1}) = \mathcal{N}(\mathcal{X}^k; \sqrt{1 - \beta_k} \mathcal{X}^{k-1}, \beta_k \mathbf{I}),$$

where  $\beta_1, \beta_2, \dots, \beta_K$  are pre-determined variance schedulers. Let  $\alpha_k = 1 - \beta_k$  and  $\bar{\alpha}_k = \prod_{s=1}^k \alpha_s$ , the diffusion process can be represented for any step  $k$  in a closed form:

$$q(\mathcal{X}^k | \mathcal{X}^0) = \mathcal{N}(\mathcal{X}^k; \sqrt{\bar{\alpha}_k} \mathcal{X}^0, (1 - \bar{\alpha}_k) \mathbf{I}). \quad (2)$$

This great property indicates that  $\mathcal{X}^0$  will be converted into isotropic Gaussian  $\mathcal{X}^K \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  when gradually adding a new noise with sufficiently large  $K$ .

**Reverse Diffusion Process.** In the opposite direction, given the historical sequence, the future motion generation procedure can be considered as a reverse dynamics of the above diffusion process  $(\mathcal{X}^K, \mathcal{X}^{K-1}, \dots, \mathcal{X}^0)$  with a parameterized Markov chain, starting from the white noise  $\mathcal{X}^K \sim p(\mathcal{X}^K)$ . Assume that the state embedding  $\mathcal{C}$  is derived by using neural networks to encode the historical sequence  $\mathcal{D}$ , i.e.,  $\mathcal{C} = f_\psi(\mathcal{D})$ , we then formulate the reverse diffusion process as a conditional Markov chain as follows:

$$p_\theta(\mathcal{X}^{0:K-1} | \mathcal{X}^K, \mathcal{C}) = \prod_{k=1}^K p_\theta(\mathcal{X}^{k-1} | \mathcal{X}^k, \mathcal{C}),$$

$$p_\theta(\mathcal{X}^{k-1} | \mathcal{X}^k, \mathcal{C}) = \mathcal{N}(\mathcal{X}^{k-1}; \boldsymbol{\mu}_\theta(\mathcal{X}^k, k, \mathcal{C}), \boldsymbol{\Sigma}_\theta(\mathcal{X}^k, k)). \quad (3)$$

$p(\mathcal{X}^K)$  is an initial noise distribution,  $\boldsymbol{\mu}_\theta(\mathcal{X}^k, k)$  is the estimated mean parameterized by  $\theta$ , and  $\boldsymbol{\Sigma}_\theta(\mathcal{X}^k, k) = \sigma^2 \mathbf{I}$  is the variance term of Gaussian transition that empirically set as  $\sigma^2 = \frac{1-\bar{\alpha}_{k-1}}{1-\bar{\alpha}_k} \beta_k$  like (Ho, Jain, and Abbeel 2020).

## Training Objective

We train the diffusion model via variational inference. To be particular, we consider the variational lower bound of log-likelihood (ELBO) to learn the marginal likelihood:

$$\begin{aligned} \mathbb{E} [-\log p_\theta(\mathcal{X}^0)] &\leq \mathbb{E}_q \left[ -\log \frac{p_\theta(\mathcal{X}^{0:K} | \mathcal{C})}{q(\mathcal{X}^{1:K} | \mathcal{X}^0)} \right] \\ &= \mathbb{E}_q \left[ -\log p(\mathcal{X}^K) - \sum_{k=1}^K \log \frac{p_\theta(\mathcal{X}^{k-1} | \mathcal{X}^k, \mathcal{C})}{q(\mathcal{X}^k | \mathcal{X}^{k-1})} \right]. \end{aligned} \quad (4)$$

This loss can be further reduced to:

$$\begin{aligned} L(\psi, \theta) = \mathbb{E}_q \left[ \sum_{k>1} D_{KL}(q(\mathcal{X}^{k-1} | \mathcal{X}^k, \mathcal{X}^0) \| p_\theta(\mathcal{X}^{k-1} | \mathcal{X}^k, \mathcal{C})) \right. \\ \left. + D_{KL}(q(\mathcal{X}^K | \mathcal{X}^0) \| p(\mathcal{X}^K)) - \log p_\theta(\mathcal{X}^0 | \mathcal{X}^1, \mathcal{C}) \right]. \end{aligned} \quad (5)$$

where the first term iteratively performs one reverse diffusion step, and the second term can be ignored due to no learnable parameters. Then, the simplified training objective becomes:

$$L(\psi, \theta) = \mathbb{E}_{k, \mathcal{X}^0, \epsilon} [\|\epsilon - \epsilon_\theta(\mathcal{X}^k, k, f_\psi(\mathcal{D}))\|^2], \quad (6)$$

where the dimension of  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  corresponds to that of the future sequence  $\mathcal{X}^0$ . The neural network  $\epsilon_\theta$  learns to predict the noise  $\epsilon$ , which can be used to denoise the noisy future motions  $\mathcal{X}^k = \sqrt{\bar{\alpha}_k} \mathcal{X}^0 + \sqrt{1 - \bar{\alpha}_k} \epsilon$ .

Remarkably, Average Pairwise Distance (APD) loss or Average Displacement Error (ADE) loss in traditional stochastic human motion prediction methods (Yuan and Kitani 2020; Ma et al. 2022) is not present in (6) because DDPM naturally defines a one-to-one joint mapping between two consecutive motions in the diffusion process as shown in (1).

## Sampling

Given a learned reverse diffusion network  $\epsilon_\theta$ , a historical sequence  $\mathcal{D}$  and its corresponding encoder  $f_\psi$ , we first sample chaotic states  $\mathcal{X}^K$  from  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ , and then progressively generate realistic future motions from  $\mathcal{X}^K$  to  $\mathcal{X}^0$  by the following equation:

$$\mathcal{X}^{k-1} = \frac{1}{\sqrt{\alpha_k}} \left( \mathcal{X}^k - \frac{\beta_k}{\sqrt{1 - \bar{\alpha}_k}} \epsilon_\theta(\mathcal{X}^k, k, f_\psi(\mathcal{D})) \right) + \sqrt{\beta_k} \mathbf{z}, \quad (7)$$

where  $\mathbf{z}$  is a random variable from standard Gaussian.

## Network Architecture

Different from the UNet backbone for conventional diffusion models (Dhariwal and Nichol 2021) in image generation, we design a spatial-temporal transformer-based architecture for stochastic human motion prediction. An overview of the proposed MotionDiff framework is illustrated in Figure 2. The detailed diffusion network (Left) consists of two key parts: an encoder network  $f_\psi$  which learns the state embedding from the historical sequence, and a decoder network  $\epsilon_\theta$  which implements the reverse diffusion process.

Given an observed sequence  $\mathcal{D} \in \mathbb{R}^{T \times J \times 3}$  and its corresponding future sequence  $\mathcal{X}^k = \sqrt{\bar{\alpha}_k} \mathcal{X}^0 + \sqrt{1 - \bar{\alpha}_k} \epsilon \in \mathbb{R}^{f \times J \times 3}$  in a noising step  $k$ , we perform joint embedding that maps 3D coordinates to the hidden dimension  $c$  for each frame  $\mathbf{x} \in \mathbb{R}^{3 \times J}$ , i.e.,  $\mathbf{y} = \mathbf{W}\mathbf{x} + \mathbf{b} \in \mathbb{R}^{c \times J}$ , where  $\mathbf{W} \in \mathbb{R}^{c \times 3}$  and  $\mathbf{b} \in \mathbb{R}^c$  are weights of the layer. The obtained outputs of observed sequence and noisy predicted one  $\mathcal{X}^k$  are then fed into spatial transformer module  $\mathcal{S}$  and PoseFormer, respectively. Note that PoseFormer also contains a spatial transformer module, which shares the same parameters with previous  $\mathcal{S}$ . We directly equip with PoseFormer introduced in (Zheng et al. 2021) for its strong representation ability. In this way, the kinematics between local joints can be sufficiently explored, and  $\mathcal{X}^k$  can learn useful geometrical information from the observed one. This differs from traditional HMP methods (Bouazizi et al. 2022) which flatten  $\mathbf{x}$  into a vector  $\hat{\mathbf{x}}$  and then perform simple pose embedding  $\hat{\mathbf{y}} = \hat{\mathbf{W}}\hat{\mathbf{x}} + \hat{\mathbf{b}} \in \mathbb{R}^{\hat{c} \times 1}$ . Following, the encodings of historical sequence are fed into Gated Recurrent Units (GRU) as the condition of the decoder  $\epsilon_\theta$ . Finally, combined with diffusion step  $k$ , historical information  $f_\psi(\mathcal{D})$  and noisy future sequence  $\mathcal{X}^k$ , the fused feature is fed into temporal transformer module to capture the global dependencies across frames. Compared to (Aksan et al. 2021), another benefit of this separable spatial-temporal design is that it can accelerate the speed of reverse diffusion process.

## Refinement

As discussed in (Lugmayr et al. 2022), the diffusion model are free to predict any motions that semantically align with the training sets because we do not directly involve any loss to supervise the motion prediction. Although these predicted motions are realistic, they might be very different from the ground truth, making the quantitative evaluation challenging. To address this problem, inspired by (Lyu et al. 2022), we use another structure to further refine the predictions as

shown in Figure 2 (right). Notably, this structure only serves as refinement, and can be flexibly substituted with different deterministic feed-forward HMP methods (Mao et al. 2019; Mao, Liu, and Salzmann 2020; Dang et al. 2021). For simplicity, we employ simple yet effective LTD-GCN (Mao et al. 2019) to construct our refinement network.

Different from (Mao et al. 2019) padding the input sequence with replicates of the last observed pose, the inputs of the refinement network are  $N$  future motions  $\mathcal{Y} = \{\mathcal{Y}_i\}_{i=1,2,\dots,N}$  generated by our pretrained diffusion model. With skip connection, the predicted displacement are added to  $\mathcal{Y}$  to derive  $N$  refined possible future motions:  $\mathcal{Z}_i = \mathcal{Y}_i + \gamma\epsilon_\phi(\mathcal{Y}_i, \mathcal{C})$ ,  $i = 1, 2, \dots, N$ , where  $\epsilon_\phi$  is the refinement network, and  $\gamma$  is used to measure the difference  $\sum_{i=1}^N \|\mathcal{Z}_i - \mathcal{Y}_i\|^2$ . We redefine the reconstruction loss between multiple refined sequences  $\mathcal{Z}$  and ground truth sequence  $\mathcal{X}$  to supervise the network  $\epsilon_\phi$ :

$$L_{rec} = \min_i \|\mathcal{Z}_i - \mathcal{X}\|^2. \quad (8)$$

We further use the diversity-promoting loss like (Yuan and Kitani 2020) to explicitly encourage diversity across the generated sequences:

$$L_{div} = \frac{1}{N(N-1)} \sum_{i=1}^N \sum_{j \neq i}^N \exp\left(-\frac{d^2(\mathcal{Z}_i, \mathcal{Z}_j)}{\sigma}\right), \quad (9)$$

where  $d(\cdot, \cdot)$  is Euclidean distance between two motions.

Altogether, our final refined loss is expressed as:

$$L_{refine} = \lambda_1 L_{rec} + \lambda_2 L_{div}, \quad (10)$$

where  $\lambda_1$  and  $\lambda_2$  are hyper-parameters that control the influence of different terms. This refinement design facilitates the use of established geometric losses on motion characteristics, encouraging natural pose sequences.

## Experiments

In this section, we first introduce two datasets, evaluation metrics and implementation details. Then, we assess the proposed model for stochastic human motion prediction in terms of diversity and accuracy. Finally, we carry out ablation study to show the influence of the different components.

### Experimental Setup

**Datasets.** Following (Yuan and Kitani 2020; Ma et al. 2022), we evaluate our model on two public benchmark datasets including Human3.6M (Ionescu et al. 2013) and HumanEva-I (Sigal, Balan, and Black 2010).

*Human3.6M* contains 3.6 million video frames performed by 11 subjects (7 with ground truth). Each subject performs 15 actions and the human motion is recorded at 50Hz. We use 5 subjects (S1, S5, S6, S7, S8) to train the model, and the rest (S9, S11) for evaluation. We consider a 17-joint skeleton for each frame and remove the global translation. We predict 100 future frames (2s) using 25 observed frames (0.5s).

*HumanEva-I* is a relatively small dataset that consists of 3 subjects recorded at 60Hz. Each subject performs 5 actions and is represented by a 15-joint skeleton. We predict 60 future frames (1s) given 15 observed frames (0.25s).

**Evaluation Metrics.** Following (Yuan and Kitani 2020), we measure diversity and accuracy in terms of five metrics. (1) *APD*: Average Pairwise Distance between all pairs of motion samples defined as  $\frac{1}{N(N-1)} \sum_{i=1}^N \sum_{j \neq i}^N \|\mathcal{Z}_i - \mathcal{Z}_j\|_2$ ; (2) *ADE*: Average Displacement Error over the whole sequence between the ground truth and the closest generated motion defined as  $\frac{1}{f} \min_i \|\mathcal{Z}_i - \mathcal{X}\|_2$ ; (3) *FDE*: Final Displacement Error between the last frame of the ground truth and the closest motion’s last frame defined as  $\min_i \|\mathcal{Z}_i[f] - \mathcal{X}[f]\|_2$ ; (4) *MMADE*: the multi-modal version of ADE; (5) *MMFDE*: the multi-modal version of FDE. Note that APD is used for diversity while the others are used for accuracy.

**Implementation Details.** For the diffusion network, we use joint embedding layer to upsample the 3D coordinate of human joints from 3 to 32, and then feed it into transformer where the hidden dimension is set to 512. We employ three fully-connected layers with Leaky ReLU to gradually downsample the temporal transformer output to the predicted future motions, such that 512d-256d-128d-3d. For the refinement network, we use a 12-layers graph convolution network and set the hidden size to 256 in each layer. Following the protocol of (Yuan and Kitani 2020), our diffusion network generates 50 diverse future motions ( $N = 50$ ) given a past motion. We set the variance schedules to be  $\beta_1 = 0.0001$  and  $\beta_K = 0.05$ , where  $\beta_k$  are linearly interpolated ( $1 < k < K$ ). Our code is in Pytorch (Paszke et al. 2017) and we use ADAM (Kingma and Ba 2015) optimizer. We train our diffusion model for 1,000 epochs with a batch size of 64 for both datasets. All the experiments are implemented on an NVIDIA RTX 3080 GPU.

### Comparison with State-of-the-art

We compare our method with two types of baselines. (1) Deterministic motion prediction methods, including ERD (Fragkiadaki et al. 2015b) and acLSTM (Li et al. 2018); (2) Stochastic motion prediction methods, including GAN based methods HP-GAN (Barsoum, Kender, and Liu 2018), DeLiGAN (Gurumurthy, Kiran Sarvadevabhatla, and Venkatesh Babu 2017) as well as VAE based methods Pose-Knows (Walker et al. 2017) and MT-VAE (Yan et al. 2018), BoM (Bhattacharyya, Schiele, and Fritz 2018), GMVAE (Dilokthanakul et al. 2016), DSF (Yuan and Kitani 2019), DLow (Yuan and Kitani 2020), MOJO (Zhang, Black, and Tang 2021) and MOVAE (Ma et al. 2022).

The comparison results on Human3.6M and HumanEva-I datasets are summarized in Table 1. For all stochastic motion prediction baselines, we provide 50 samplings for each historical motion. Overall, our diffusion-refinement model effectively improves on the state-of-the-art. We can observe that our model outperforms other methods with a large margin in terms of the diversity. The reason is that each diffusion step in our model is inherently diverse since it incorporates a new noise from Gaussian distribution. We also achieve comparable performance in terms of the prediction accuracies. In general, stochastic motion prediction methods (e.g., BoM, GMVAE) can achieve better diversity and accuracy compared to deterministic ones (e.g., ERD, acLSTM). Deterministic prediction models tend to generate poor future

Type	Method	Human3.6M					HumanEva-I				
		APD $\uparrow$	ADE $\downarrow$	FDE $\downarrow$	MMADE $\downarrow$	MMFDE $\downarrow$	APD $\uparrow$	ADE $\downarrow$	FDE $\downarrow$	MMADE $\downarrow$	MMFDE $\downarrow$
Deterministic	ERD (ICCV'15)	0	0.722	0.969	0.776	0.995	0	0.382	0.461	0.521	0.595
	acLSTM (ICLR'18)	0	0.789	1.126	0.849	1.139	0	0.429	0.541	0.530	0.608
Stochastic	Pose-Knows(ICCV'17)	6.723	0.461	0.560	0.522	0.569	2.302	0.269	0.296	0.384	0.375
	MT-VAE (ECCV'18)	0.403	0.457	0.595	0.716	0.883	0.021	0.345	0.403	0.518	0.577
	HP-GAN (CVPRW'18)	7.214	0.858	0.867	0.847	0.858	1.139	0.772	0.749	0.776	0.769
	BoM (CVPR'18)	6.265	0.448	0.533	0.514	0.544	2.846	0.271	0.279	0.373	0.351
	GMVAE (arXiv'16)	6.769	0.461	0.555	0.524	0.566	2.443	0.305	0.345	0.408	0.410
	DeLiGAN (CVPR'17)	6.509	0.483	0.534	0.520	0.545	2.177	0.306	0.322	0.385	0.371
	DSF (ICLR'19)	9.330	0.493	0.592	0.550	0.599	4.538	0.273	0.290	0.364	0.340
	DLow (ECCV'20)	11.741	0.425	0.518	<b>0.495</b>	<b>0.531</b>	4.855	0.251	0.268	0.362	0.339
	MOJO (CVPR'21)	12.579	0.412	0.514	0.497	0.538	4.181	0.234	0.244	0.369	0.347
	MOVAE (CVPR'22)	14.240	0.414	0.516	-	-	5.786	<b>0.228</b>	<b>0.236</b>	-	-
Ours		<b>15.353</b>	<b>0.411</b>	<b>0.509</b>	0.508	0.536	<b>5.931</b>	0.232	<b>0.236</b>	<b>0.352</b>	<b>0.320</b>

Table 1: Quantitative results by predicting 50 future motions for each historical motion. The best results are marked in bold.




Figure 3: Qualitative comparison between DLow and ours. Given the observed motion, we show the end poses of ten future predictions. Left: Our model yields motions with more diversity. Right: Red boxes in DLow indicate unreasonable predictions.

motions in long-term horizon ( $>1s$ ) while stochastic predict models can make a trade-off between diversity and accuracy.

Notably, since our model is constructed with two stages, we draw our attention to the comparison with DLow (Yuan and Kitani 2020) that also generates diverse motions with a two-stage design. Our model outperforms DLow on both datasets, especially achieving a very significant improvement of 31% (Human3.6M) / 22% (HumanEva-I) on diversity. This is owing to the advantage of intrinsic diversity that the diffusion process brings in DDPM, rather than directly designing losses from a pretrained VAE model in DLow. Besides, our approach is benefited more on larger dataset (Human3.6M) for diversity, which is in line with the results in stochastic trajectory prediction (Gu et al. 2022).

We further investigate the ability of our method by the qualitative analysis. Figure 3 illustrates ten end poses of future motions generated by DLow and ours. It can be seen that our predictions have more different modes than that of DLow, especially for Human3.6M dataset, which confirms the inherent diversity of our diffusion process. On the other hand, our method yields more realistic future motions than DLow (highlighted by red boxes). This can be ascribed to the high capability of DDPM framework. More visualization results can be found in <https://github.com/csdwei/MotionDiff>.

### Ablation Study

In this subsection, we conducted ablation studies to investigate the effect of individual design component including diffusion network and refinement network. We further explore

detailed architecture parameter combinations.

**Diffusion Network.** The diversity of our method mainly lies in noise addition at each diffusion step. To evaluate the influence of different maximum number of diffusion steps  $K$ , we provide an analysis between  $K$  and diversity as well as prediction error on HumanEva-I in Figure 4 (left). We notice that when  $K$  is small, the diversity is limited while the predictions are closer to the ground truth (except for  $K = 50$ ). As  $K$  increases, ADE, FDE, MMADE and MMFDE all become worse whereas the diversity gets higher. This is reasonable since more diffusion steps can generate more diverse yet realistic samples that are very different from the ground truth while too few steps ( $K < 100$ ) is unable to produce the whitened noise. Hence, our diffusion framework can flexibly make a trade-off between diversity and accuracy by adjusting the maximum diffusion step numbers  $K$ .

Furthermore, we investigate the performance of the reverse diffusion process from  $K$  to 0 as exhibited in Figure 4 (right). Although the diversity is very high when the reverse diffusion step  $k$  is large, the kinematics of human joints have been completely destroyed. As  $k$  decreases, we gradually get realistic future motions with the decline of diversity.

**Refinement Network.** To analyze whether the refinement network causes the approximation to the ground truth, we degrade our diffusion-refinement paradigm into two separable parts. Comparison results on Human3.6M and HumanEva-I datasets are tabulated in Table 2. Note that ‘Refinement’ network actually represents the determinis-

Diffusion	Refinement	Human3.6M					HumanEva-I				
		APD $\uparrow$	ADE $\downarrow$	FDE $\downarrow$	MMADE $\downarrow$	MMFDE $\downarrow$	APD $\uparrow$	ADE $\downarrow$	FDE $\downarrow$	MMADE $\downarrow$	MMFDE $\downarrow$
$\times$	$\checkmark$	0	0.516	0.756	0.627	0.795	0	0.415	0.555	0.509	0.613
$\checkmark$	$\times$	<b>15.534</b>	0.486	0.536	0.551	0.564	<b>6.508</b>	0.273	0.259	0.367	0.335
$\checkmark$	$\checkmark$	15.353	<b>0.411</b>	<b>0.509</b>	<b>0.508</b>	<b>0.536</b>	5.931	<b>0.232</b>	<b>0.236</b>	<b>0.352</b>	<b>0.320</b>

Table 2: Comparison results on two datasets in terms of training the diffusion network and the refinement network.

Metrics	① Spatial Module		② Spatial Trans Dimension			③ Temporal Module		④ Temporal Trans Dimension			⑤ $L_1$ vs. $L_2$	
	w/o Trans	w/ Trans*	16	32*	48	MLP	Trans*	256	512*	1024	$L_1$	$L_2^*$
APD $\uparrow$	15.212	<b>15.353</b>	14.947	15.353	<b>15.761</b>	10.943	<b>15.353</b>	12.4735	<b>15.353</b>	15.248	13.894	<b>15.353</b>
ADE $\downarrow$	0.417	<b>0.411</b>	<b>0.409</b>	0.411	0.424	0.419	<b>0.411</b>	0.4301	<b>0.411</b>	0.415	<b>0.411</b>	<b>0.411</b>
FDE $\downarrow$	0.535	<b>0.509</b>	0.511	<b>0.509</b>	0.543	0.531	<b>0.509</b>	0.5326	<b>0.509</b>	0.516	0.511	<b>0.509</b>
MMADE $\downarrow$	0.519	<b>0.508</b>	<b>0.506</b>	0.508	0.530	0.517	<b>0.508</b>	0.5209	0.508	<b>0.494</b>	<b>0.508</b>	<b>0.508</b>
MMFDE $\downarrow$	0.564	<b>0.536</b>	0.542	<b>0.536</b>	0.560	0.559	<b>0.536</b>	0.5561	0.536	<b>0.511</b>	0.539	<b>0.536</b>

Table 3: Comparison with different network architecture designs on Human3.6M dataset. Trans is abbreviation of Transformer.




Figure 4: Left: The diversity and prediction error of our model trained with different maximum diffusion step numbers  $K$ . Right: The trend of diversity and prediction error at the  $k$ -th reverse diffusion step within  $K$  when  $K = 100$ .

tic method LTD-GCN (Mao et al. 2019). We can see that both diffusion network and diffusion-refinement architecture yield high diversity compared to other baselines, while the latter one significantly approximates to the ground truth at the cost of diversity. We also illustrate the predicted human motions from the diffusion network and diffusion-refinement architecture for Human3.6M dataset in Figure 5. The diffusion strategy without refinement is far from the ground truth (highlighted by red boxes), which again validates the effectiveness of the refinement network. Remarkably, although the prediction accuracies without refinement are relatively low due to the lack of direct guided loss, its generated future motions are still plausible.

**Detailed Architecture.** We perform five groups of ablation studies about detailed architecture of our method. Comparison results on Human3.6M dataset are reported in Table 3. **(1) Effect of spatial transformer.** ‘w/o Trans’ represents that we use a simple linear projection layer to perform pose embedding as mentioned before, while ‘w/ Trans’ represents the combination of joint embedding and spatial transformer. The results clearly demonstrate the advantage of using spatial transformer to expressively model the relationships between joints in each frame. **(2) Dimension of spatial transformer.** The best APD is obtained when  $c = 48$ , while the best FDE, MMFDE are obtained when  $c = 32$ . We select 32 as the default value of  $c$ . **(3) Effect of temporal transformer.** ‘MLP’ means that we employ the efficient and pop-




Figure 5: Visualization of predicted time sequences from our diffusion and diffusion-refinement networks. Diffusion without refinement yields results far from the ground truth, but are still plausible.

ular multiple layer perceptron structure in (Guo et al. 2022). We find temporal transformer module outperforms MLP, especially for diversity. It indicates that temporal transformer is effective for our diffusion-refinement model to capture the global dependencies across frames. **(4) Dimension of temporal transformer.** We observe that the temporal transformer with 512 dimensions results in the best performance, and higher dimension fails to acquire better results. **(5) Using  $L_2$  or  $L_1$ .** (Chen et al. 2021) suggest that substituting the original  $L_2$  distance metric with  $L_1$  offers better training stability. For stochastic human motion prediction, we find that  $L_2$  distance achieves higher diversity as well as accuracy.

## Conclusion

In this paper, we propose a new probabilistic model for stochastic human motion prediction. Our method marries denoising diffusion models with human joint kinematics, and derives the diversity by incorporating noises in diffusion process. By learning a parameterized Markov chain conditioned on the historical sequence to progressively remove noise from the noise distribution, we can generate diverse yet realistic future motions. Besides, we design a spatial-temporal transformer-based architecture to encode local relationships in each frame and global dependencies across frames, which is then fed into the refinement network to further improve the accuracy. Experimental results show that our method is competitive with the state-of-the-art.

## Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (NO. 62176125, 61772272).

## References

- Aksan, E.; Kaufmann, M.; Cao, P.; and Hilliges, O. 2021. A spatio-temporal transformer for 3d human motion prediction. In *Proceedings of the International Conference on 3D Vision*, 565–574. IEEE.
- Aksan, E.; Kaufmann, M.; and Hilliges, O. 2019. Structured prediction helps 3d human motion modelling. In *Proceedings of the IEEE International Conference on Computer Vision*, 7144–7153.
- Aliakbarian, S.; Saleh, F. S.; Salzmann, M.; Petersson, L.; and Gould, S. 2020. A stochastic conditioning scheme for diverse human motion prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5223–5232.
- Bajcsy, A.; Siththaranjan, A.; Tomlin, C. J.; and Dragan, A. D. 2021. Analyzing human models that adapt online. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2754–2760. IEEE.
- Barsoum, E.; Kender, J.; and Liu, Z. 2018. Hp-gan: Probabilistic 3d human motion prediction via gan. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 1418–1427.
- Bhattacharyya, A.; Schiele, B.; and Fritz, M. 2018. Accurate and diverse sampling of sequences based on a “best of many” sample objective. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8485–8493.
- Bouazizi, A.; Holzbock, A.; Kressel, U.; Dietmayer, K.; and Belagiannis, V. 2022. MotionMixer: MLP-based 3D Human Body Pose Forecasting. In *Proceedings of the International Joint Conferences on Artificial Intelligence*.
- Cai, Y.; Huang, L.; Wang, Y.; Cham, T.-J.; Cai, J.; Yuan, J.; Liu, J.; Yang, X.; Zhu, Y.; Shen, X.; et al. 2020. Learning progressive joint propagation for human motion prediction. In *Proceedings of the European Conference on Computer Vision*, 226–242.
- Chen, N.; Zhang, Y.; Zen, H.; Weiss, R. J.; Norouzi, M.; and Chan, W. 2021. WaveGrad: Estimating gradients for waveform generation. In *Proceedings of the International Conference on Learning Representations*.
- Dang, L.; Nie, Y.; Long, C.; Zhang, Q.; and Li, G. 2021. MSR-GCN: Multi-Scale Residual Graph Convolution Networks for Human Motion Prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, 11467–11476.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. In *Proceedings of the Advances in Neural Information Processing Systems*, 8780–8794.
- Dilokthanakul, N.; Mediano, P. A.; Garnelo, M.; Lee, M. C.; Salimbeni, H.; Arulkumaran, K.; and Shanahan, M. 2016. Deep unsupervised clustering with gaussian mixture variational autoencoders. *arXiv preprint arXiv:1611.02648*.
- Fragkiadaki, K.; Levine, S.; Felsen, P.; and Malik, J. 2015a. Recurrent network models for human dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*, 4346–4354.
- Fragkiadaki, K.; Levine, S.; Felsen, P.; and Malik, J. 2015b. Recurrent network models for human dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*, 4346–4354.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Proceedings of the Advances in Neural Information Processing Systems*, 27.
- Gu, T.; Chen, G.; Li, J.; Lin, C.; Rao, Y.; Zhou, J.; and Lu, J. 2022. Stochastic trajectory prediction via motion indeterminacy diffusion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 17113–17122.
- Guo, W.; Du, Y.; Shen, X.; Lepetit, V.; Alameda-Pineda, X.; and Moreno-Noguer, F. 2022. Back to MLP: A simple baseline for human motion prediction. *arXiv preprint arXiv:2207.01567*.
- Gurumurthy, S.; Kiran Sarvadevabhatla, R.; and Venkatesh Babu, R. 2017. Deligan: Generative adversarial networks for diverse and limited data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 166–174.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. In *Proceedings of the Advances in Neural Information Processing Systems*, 6840–6851.
- Ionescu, C.; Papava, D.; Olaru, V.; and Sminchisescu, C. 2013. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7): 1325–1339.
- Kingma, D. P.; and Ba, J. 2015. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations*.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Kong, Z.; Ping, W.; Huang, J.; Zhao, K.; and Catanzaro, B. 2021. DiffWave: A versatile diffusion model for audio synthesis. In *Proceedings of the International Conference on Learning Representations*.
- Kundu, J. N.; Gor, M.; and Babu, R. V. 2019. Bihmp-gan: Bidirectional 3d human motion prediction gan. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 8553–8560.
- Li, M.; Chen, S.; Zhao, Y.; Zhang, Y.; Wang, Y.; and Tian, Q. 2020. Dynamic multiscale graph neural networks for 3d skeleton based human motion prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 214–223.
- Li, X.; Li, H.; Joo, H.; Liu, Y.; and Sheikh, Y. 2018. Structure from recurrent motion: From rigidity to recurrency. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3032–3040.



- Liu, Z.; Wu, S.; Jin, S.; Liu, Q.; Ji, S.; Lu, S.; and Cheng, L. 2022. Investigating pose representations and motion contexts modeling for 3D motion prediction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Lugmayr, A.; Danelljan, M.; Romero, A.; Yu, F.; Timofte, R.; and Van Gool, L. 2022. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 11461–11471.
- Lyu, Z.; Kong, Z.; Xudong, X.; Pan, L.; and Lin, D. 2022. A conditional point diffusion-refinement paradigm for 3D point cloud completion. In *Proceedings of the International Conference on Learning Representations*.
- Ma, H.; Li, J.; Hosseini, R.; Tomizuka, M.; and Choi, C. 2022. Multi-Objective Diverse Human Motion Prediction With Knowledge Distillation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8161–8171.
- Mao, W.; Liu, M.; and Salzmann, M. 2020. History repeats itself: Human motion prediction via motion attention. In *Proceedings of the European Conference on Computer Vision*, 474–489. Springer.
- Mao, W.; Liu, M.; and Salzmann, M. 2021. Generating smooth pose sequences for diverse human motion prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, 13309–13318.
- Mao, W.; Liu, M.; Salzmann, M.; and Li, H. 2019. Learning trajectory dependencies for human motion prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, 9489–9497.
- Martinez, J.; Black, M. J.; and Romero, J. 2017. On human motion prediction using recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2891–2900.
- McCarthy, A. D.; Li, X.; Gu, J.; and Dong, N. 2020. Addressing posterior collapse with mutual information for improved variational neural machine translation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 8512–8525.
- Paden, B.; Čáp, M.; Yong, S. Z.; Yershov, D.; and Frazzoli, E. 2016. A survey of motion planning and control techniques for self-driving urban vehicles. *IEEE Transactions on Intelligent Vehicles*, 1(1): 33–55.
- Park, S.; Ryu, H.; Lee, S.; Lee, S.; and Lee, J. 2019. Learning predict-and-simulate policies from unorganized human motion data. *ACM Transactions on Graphics*, 38(6): 1–11.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in pytorch. In *Proceedings of the Advances in Neural Information Processing Systems Workshops*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention*, 234–241. Springer.
- Sigal, L.; Balan, A. O.; and Black, M. J. 2010. Humaneva: Synchronized video and motion capture dataset and baseline algorithm for evaluation of articulated human motion. *International Journal of Computer Vision*, 87(1): 4–27.
- Sofianos, T.; Sampieri, A.; Franco, L.; and Galasso, F. 2021. Space-time-separable graph convolutional network for pose forecasting. In *Proceedings of the IEEE International Conference on Computer Vision*, 11209–11218.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the International Conference on Machine Learning*, 2256–2265. PMLR.
- Sun, X.; Cui, Q.; Sun, H.; Li, B.; Li, W.; and Lu, J. 2022. Overlooked Poses Actually Make Sense: Distilling Privileged Knowledge for Human Motion Prediction. In *Proceedings of the European Conference on Computer Vision*, 678–694.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Proceedings of the Advances in Neural Information Processing Systems*.
- Walker, J.; Marino, K.; Gupta, A.; and Hebert, M. 2017. The pose knows: Video forecasting by generating pose futures. In *Proceedings of the IEEE International Conference on Computer Vision*, 3332–3341.
- Yan, X.; Rastogi, A.; Villegas, R.; Sunkavalli, K.; Shechtman, E.; Hadap, S.; Yumer, E.; and Lee, H. 2018. Mt-vae: Learning motion transformations to generate multimodal human dynamics. In *Proceedings of the European Conference on Computer Vision*, 265–281.
- Yuan, Y.; and Kitani, K. 2020. Dlow: Diversifying latent flows for diverse human motion prediction. In *Proceedings of the European Conference on Computer Vision*, 346–364.
- Yuan, Y.; and Kitani, K. M. 2019. Diverse Trajectory Forecasting with Determinantal Point Processes. In *Proceedings of the International Conference on Learning Representations*.
- Zhang, Y.; Black, M. J.; and Tang, S. 2021. We are more than our joints: Predicting how 3d bodies move. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3372–3382.
- Zheng, C.; Zhu, S.; Mendieta, M.; Yang, T.; Chen, C.; and Ding, Z. 2021. 3d human pose estimation with spatial and temporal transformers. In *Proceedings of the IEEE International Conference on Computer Vision*, 11656–11665.
- Zhong, C.; Hu, L.; Zhang, Z.; Ye, Y.; and Xia, S. 2022. Spatio-temporal gating-adjacency GCN for human motion prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6447–6456.