

Diverse Human Motion Prediction Guided by Multi-Level Spatial-Temporal Anchors

Sirui Xu[✉], Yu-Xiong Wang*, and Liang-Yan Gui*

University of Illinois at Urbana-Champaign
{siruiXu2, yxw, lgui}@illinois.edu

Abstract. Predicting diverse human motions given a sequence of historical poses has received increasing attention. Despite rapid progress, existing work captures the multi-modal nature of human motions primarily through likelihood-based sampling, where the mode collapse has been widely observed. In this paper, we propose a simple yet effective approach that disentangles randomly sampled codes with a *deterministic learnable component named anchors* to promote sample precision and diversity. Anchors are further factorized into spatial anchors and temporal anchors, which provide attractively *interpretable* control over spatial-temporal disparity. In principle, our spatial-temporal anchor-based sampling (STARS) can be applied to different motion predictors. Here we propose an interaction-enhanced spatial-temporal graph convolutional network (IE-STGCN) that encodes prior knowledge of human motions (*e.g.*, spatial locality), and incorporate the anchors into it. Extensive experiments demonstrate that our approach outperforms state of the art in both stochastic and deterministic prediction, suggesting it as a *unified* framework for modeling human motions. Our code and pretrained models are available at <https://github.com/Sirui-Xu/STARS>.

Keywords: Stochastic Human Motion Prediction, Generative Models, Graph Neural Networks

1 Introduction

Predicting the evolution of the surrounding physical world over time is an essential aspect of human intelligence. For example, in a seamless interaction, a robot is supposed to have some notion of how people move or act in the near future, conditioned on a series of historical movements. Human motion prediction has thus been widely used in computer vision and robotics, such as autonomous driving [54], character animation [62], robot navigation [59], motion tracking [48], and human-robot interaction [7, 34, 35, 38]. Owing to deep learning techniques, there has been significant progress over the past few years in modeling and predicting motions. Despite notable successes, forecasting human motions, especially over longer time horizons (*i.e.*, up to several seconds), is fundamentally challenging, because of the difficulty of modeling multi-modal motion dynamics

* Yu-Xiong Wang and Liang-Yan Gui contributed equally to this work.

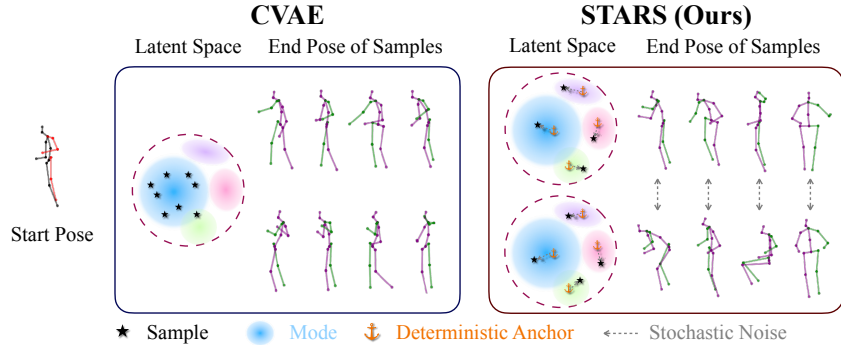


Fig. 1: Our Spatial-Temporal AnchoR-based Sampling (STARS) is able to *capture multiple modes*, thus facilitating diverse human motion prediction. **Left:** with the traditional generative models such as conditional variational autoencoders (CVAEs), the predicted motions are often concentrated in the major mode with less diversity (illustrated with 8 samples). **Right:** STARS is able to cover more modes, where motions in the same mode have similar characteristics but vary widely across modes. Here, we use 4 anchors to pinpoint different modes. With each anchor, we sample noise and generate 2 similar motions with slight variation in each mode

and uncertainty of human conscious movements. Learning such uncertainty can, for example, help reduce the search space in motion tracking problems.

As a powerful tool, deep generative models are thus introduced for this purpose, where random codes from a prior distribution are employed to capture the multi-modal distribution of future human motion. However, current motion capture datasets are typically constructed in a way that there is *only a single ground truth future sequence* for each single historical sequence [30, 60], which makes it difficult for generators to model the underlying multi-modal densities of future motion distribution. Indeed, in practice, generators tend to ignore differences in random codes and simply produce similar predictions. This is known as *mode collapse* – the samples are concentrated in the major mode, as depicted with a representative example in Fig. 1, which has been widely observed [72]. Recent work has alleviated this problem by explicitly promoting diversity in sampling using post-hoc diversity mappings [72], or through sequentially generating different body parts [51] to achieve combinatorial diversity. These techniques, however, induce additional modeling complexity, without guaranteeing that the diversity modeling accurately covers multiple plausible modes of human motions.

To this end, we propose a simple yet effective strategy – Multi-Level Spatial-Temporal AnchoR-Based Sampling (STARS) – with the *key insight* that future motions are not completely random or independent of each other; they share some deterministic properties in line with physical laws and human body constraints, and continue trends of historical movements. For example, we may expect changes in velocity or direction to be shared deterministically among some future motions, whereas they might differ in the magnitude stochastically. Based on this observation, we disentangle latent codes in the generative model into a

stochastic component (noise) and a *deterministic learnable* component named *anchors*. With this disentanglement, the diversity of predictions is jointly affected by random noise as well as anchors that are learned to be specialized for certain modes of future motion. In contrast, the diversity from traditional generative models is determined by solely independent noise, as depicted in Fig. 1. Now, on the one hand, random noise only accounts for modeling the uncertainty *within* the mode identified by the anchor, which reduces the burden of having to model the entire future diversity. On the other hand, the model can better capture deterministic states of multiple modes by directly optimizing the anchors, thereby reducing the modeling complexity.

Naturally, human motions exhibit variation in the spatial and temporal domains, and these two types of variation are comparatively independent. Inspired by this, we propose a further decomposition to *factorize anchors into spatial and temporal anchors*. Specifically, our designed spatial anchors capture future motion variation at the spatial level, but remain constant at the temporal level, and vice versa. Another appealing property of our approach is that, by introducing straightforward linear interpolation of spatial-temporal anchors, we achieve flexible and seamless control over the predictions (Fig. 6 and Fig. 7). Unlike low-level controls that combine motions of different body parts [51, 72], our work enables manipulation of future motions in the *native space* and *time*, which is an under-explored problem. Additionally, we propose a multi-level mechanism for spatial-temporal anchors to capture multi-scale modes of future motions.

As a key advantage, spatial-temporal anchors are compatible with any motion predictor. Here, we introduce an **Interaction-Enhanced Spatial-Temporal Graph Covolutional Network** (IE-STGCN). This model encodes the spatial locality of human motion and achieves state-of-the-art performance as a motion predictor.

Our contributions can be summarized as follows. (1) We propose a novel anchor-based generative model that formulates sampling as learning deterministic anchors with likelihood sampling to better capture the multiple modes of human motions. (2) We propose a multi-level spatial-temporal decomposition of anchors for interpretable control over future motions. (3) We develop a spatial-temporal graph neural network with interaction enhancement to incorporate our anchor-based sampling. (4) We demonstrate that our approach, as a *unified* framework for modeling human motions, significantly outperforms state-of-the-art models in both diverse and deterministic human motion prediction.

2 Related Work

Deterministic Motion Prediction. Existing work on deterministic human motion forecasting predicts a single future motion based on a sequence of past poses [1, 6, 15, 42, 49], or video frames [11, 71, 74], or under the constraints of the scene context [8, 12, 25], by using recurrent neural networks (RNNs) [63], temporal convolutional networks (TCNs) [3], and graph neural networks (GNNs) [33] for sequence modeling. Common early trends involve the use of RNNs [20–22, 31, 65], which are limited in long-term temporal encoding due to error ac-

cumulation [18, 53] and training difficulty [55]. Some recent attempts exploit GNNs [16, 50] to encode poses from the spatial level, but such work still relies on RNNs [41], CNNs [14, 39, 40], or feed-forward networks [52] for temporal modeling. Recently, spatial-temporal graph convolutional networks (STGCNs) [61, 67, 69] are proposed to jointly encode the spatial and temporal correlations with spatial-temporal graphs. Continuing this effort, we propose IE-STGCN, which additionally encodes inductive biases such as spatial locality into STGCNs.

Stochastic Motion Prediction. Stochastic human motion prediction is an emerging trend with the development of deep generative models such as variational autoencoders (VAEs) [32], generative adversarial networks (GANs) [19], and normalizing flows (NFs) [58]. Most existing work [2, 4, 26, 37, 43, 64, 66, 73] produces various predictions from a set of codes independently sampled from a given distribution. As depicted in DLow [72], such likelihood-based sampling cannot produce enough diversity, as many samples are merely perturbations in the major mode. To overcome the issue, DLow employs a two-stage framework, using post-hoc mappings to shape the latent samples to improve the diversity. GSPS [51] generates different body parts in a sequential manner to achieve combinatorial diversity. Nevertheless, their explicit promotion of diversity induces additional complexity but does not directly enhance multi-mode capture. We introduce anchors that are comparatively easy to optimize, to locate deterministic components of motion modes and impose sample diversity.

Controllable Motion Prediction. Controllable motion prediction has been explored in computer graphics for virtual character generation [27, 28, 44]. In the prediction task, DLow [72] and GSPS [51] propose to control the predicted motion by separating upper and lower body parts, fixing one part while controlling the diversity of the other. In this paper, through the use of spatial-temporal anchors, we propose different but more natural controllability in native space and time. By varying and interpolating the spatial and temporal anchors, we achieve high-level control over the spatial and temporal variation, respectively.

Learnable Anchors. Our anchor-based sampling, *i.e.*, sampling with deterministic learnable codes, is inspired by work on leveraging predefined primitives and learnable codes for applications such as trajectory prediction [10, 13, 36, 46, 57], object detection [9, 45], human pose estimation [68], and video representation learning [24]. Anchors usually refer to the *hypothesis* of predictions, such as box candidates with different shapes and locations in object detection [45]. In a similar spirit, anchors in the context of human motion prediction indicate assumptions about future movements. The difference is that the anchors here are *not hand-crafted or predefined* primitives; instead, they are latent codes learned from the data. In the meantime, we endow anchors with explainability *i.e.*, to describe the multi-level spatial-temporal variation of future motions.

3 Methodology

Problem Formulation. We denote the input motion sequence of length T_h as $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{T_h}]^T$, where the 3D coordinates of V joints are used to describe




Fig. 2: Comparison of generative models without and with anchor-based sampling. Anchors and network parameters are jointly optimized. (a) Conventional generative model with only stochastic noise; (b) Generative model with deterministic anchor process: an anchor with Gaussian noise corresponds to a prediction; (c) Spatial-temporal compositional anchors: any pair of combined spatial and temporal anchors corresponds to a prediction; (d) Multi-level spatial-temporal anchors: anchors at different levels are combined for encoding multi-scale modes

each pose $\mathbf{x}_i \in \mathbb{R}^{V \times C^{(0)}}$. Here, we have $C^{(0)} = 3$. The K output sequences of length T_p are denoted as $\hat{\mathbf{Y}}_1, \hat{\mathbf{Y}}_2, \dots, \hat{\mathbf{Y}}_K$. We have access to a *single* ground truth future motion of length T_p as \mathbf{Y} . *Our objectives* are: (1) one of the K predictions is as close to the ground truth as possible; and (2) the K sequences are as diverse as possible, yet representing realistic future motions.

In this section, we first briefly review deep generative models, describe how they draw samples to generate multiple futures, and discuss their limitations (Sec. 3.1). We then detail our insights on STARS including anchor-based sampling and multi-level spatial-temporal anchors (Sec. 3.1 and Fig. 2). To model the human motion, we design an IE-STGCN and incorporate our spatial-temporal anchors into it (Sec. 3.2), as illustrated in Fig. 3.

3.1 Multi-Level Spatial-Temporal Anchor-Based Sampling

Preliminaries: Deep Generative Models. There is a large body of work on the generation of multiple hypotheses with deep generative models, most of which learn a parametric probability distribution function explicitly or implicitly. Let $p(\mathbf{Y}|\mathbf{X})$ denote the distribution of the future human motion \mathbf{Y} conditioned on the past sequence \mathbf{X} . With a latent variable $\mathbf{z} \in \mathcal{Z}$, the distribution can be reparameterized as $p(\mathbf{Y}|\mathbf{X}) = \int p(\mathbf{Y}|\mathbf{X}, \mathbf{z})p(\mathbf{z})d\mathbf{z}$, where $p(\mathbf{z})$ is often a Gaussian prior distribution. To generate a future motion sequence $\hat{\mathbf{Y}}$, \mathbf{z} is drawn from the given distribution $p(\mathbf{z})$, and then a deterministic generator $\mathcal{G} : \mathcal{Z} \times \mathcal{X} \rightarrow \mathcal{Y}$ is used for mapping, as illustrated in Fig. 2(a):

$$\mathbf{z} \sim p(\mathbf{z}), \quad \hat{\mathbf{Y}} = \mathcal{G}(\mathbf{z}, \mathbf{X}), \quad (1)$$

where \mathcal{G} is a deep neural network parameterized by θ . The goal of generative modeling is to make the distribution $p_\theta(\hat{\mathbf{Y}}|\mathbf{X})$ derived from the generator \mathcal{G} close to the actual distribution $p(\mathbf{Y}|\mathbf{X})$.

To generate K diverse motion predictions, traditional approaches first independently sample a set of latent codes $Z = \{\mathbf{z}_1, \dots, \mathbf{z}_K\}$ from a prior dis-

tribution $p(\mathbf{z})$. Although in theory, generative models are capable of covering different modes, they are not guaranteed to locate all the modes precisely, and mode collapse has been widely observed [70, 72].

Anchor-Based Sampling. To address this problem, we propose a simple yet effective sampling strategy. Our intuition is that the diversity in future motions could be characterized by: (1) deterministic component – across different actions performed by different subjects, there exist correlated or shareable changes in velocity, direction, movement patterns, etc., which naturally emerge and can be directly learned from data; and (2) stochastic component – given an action carried out by a subject, the magnitude of the changes exists which is stochastic.

Therefore, we disentangle the code in the latent space of the generative model into a *stochastic* component sampled from $p(\mathbf{z})$, and a *deterministic* component represented by a set of K learnable parameters called *anchors* $\mathcal{A} = \{\mathbf{a}_k\}_{k=1}^K$. Deterministic anchors are expected to identify as many modes as possible, which is achieved through carefully designed optimizations, while stochastic noise further specifies motion variation within certain modes. With this latent code disentanglement, we denote the new multi-modal distribution as

$$p_\theta(\hat{\mathbf{Y}}|\mathbf{X}, \mathcal{A}) = \frac{1}{K} \sum_{k=1}^K \int p_\theta(\hat{\mathbf{Y}}|\mathbf{X}, \mathbf{z}, \mathbf{a}_k) p(\mathbf{z}) d\mathbf{z}. \quad (2)$$

Consequently, as illustrated in Fig. 2(b), suppose we select the k -th learned anchor $\mathbf{a}_k \in \mathcal{A}$, along with the randomly sampled noise $\mathbf{z} \in Z$, we can generate the prediction $\hat{\mathbf{Y}}_k$ as,

$$\mathbf{z} \sim p(\mathbf{z}), \quad \hat{\mathbf{Y}}_k = \mathcal{G}(\mathbf{a}_k, \mathbf{z}, \mathbf{X}). \quad (3)$$

We can produce a total of K predictions if using each anchor once, though all anchors are not limited to being used or used only once. To incorporate anchors into the network, we find it effective to make simple additions between selected anchors and latent features, as shown in Fig. 3.

Spatial-Temporal Compositional Anchors. We observe that the diversity of future motions can be roughly divided into two types, namely *spatial variation* and *temporal variation*, which are relatively independent. This sheds light on a feasible further decomposition of the K anchors into two types of learnable codes: spatial anchors $\mathcal{A}_s = \{\mathbf{a}_i^s\}_{i=1}^{K_s}$ and temporal anchors $\mathcal{A}_t = \{\mathbf{a}_j^t\}_{j=1}^{K_t}$, where $K = K_s \times K_t$. With this decomposition, we still can yield a total of $K_s \times K_t$ compositional anchors through *each pair of spatial-temporal anchors*. Note that the temporal anchors here, in fact, control the frequency variation of future motion sequences, since our temporal features are in the frequency domain, as we will demonstrate in Sec. 3.2. To be more specific, conceptually, all spatial anchors are set to be identical in the temporal dimension but characterize the variation of motion in the spatial dimension, taking control of the movement trends and directions. Meanwhile, all temporal anchors remain unchanged in the spatial dimension but differ in the temporal dimension, producing disparities in frequency to affect the movement speed.




Fig. 3: Overview of our STARS w/ IE-STGCN framework. We combine the multi-level spatial-temporal anchors, the sampled noise, with the backbone IE-STGCN. To generate one of the predictions given a past motion, we draw noise \mathbf{z}_k , and add the selected spatial-temporal anchors to the latent feature at each level

To produce $\hat{\mathbf{Y}}_k$, as depicted in Fig. 2(c), we sample \mathbf{z} and select i -th spatial anchor \mathbf{a}_i^s and j -th temporal anchor \mathbf{a}_j^t ,

$$\mathbf{z} \sim p(\mathbf{z}), \hat{\mathbf{Y}}_k = \mathcal{G}(\mathbf{a}_i^s + \mathbf{a}_j^t, \mathbf{z}, \mathbf{X}), \quad (4)$$

where $\mathbf{a}_i^s + \mathbf{a}_j^t$ is a spatial-temporal compositional anchor corresponding to an original anchor \mathbf{a}_k . Furthermore, motion control over spatial and temporal variation can be customized through these spatial-temporal anchors. For example, we can produce future motions with similar trends by fixing the spatial anchors while varying or interpolating the temporal anchors, as shown in Sec. 4.3.

Multi-Level Spatial-Temporal Anchors. To further learn and capture multi-scale modes of future motions, we propose a multi-level mechanism to extend the spatial-temporal anchors. As an illustration, Fig. 2(d) shows a simple two-level case for this design. We introduce two different spatial-temporal anchor sets, $\{\mathcal{A}_t^{(1)}, \mathcal{A}_s^{(1)}\}$ and $\{\mathcal{A}_t^{(2)}, \mathcal{A}_s^{(2)}\}$, and assign them sequentially to different network parts $\mathcal{G}^{(1)}, \mathcal{G}^{(2)}$. Suppose (i, j) is a spatial-temporal index corresponding to the 1D index k , we can generate $\hat{\mathbf{Y}}_k$ through a two-level process as

$$\mathbf{z} \sim p(\mathbf{z}), \hat{\mathbf{Y}}_k = \mathcal{G}^{(2)}(\mathbf{a}_i^{s_2} + \mathbf{a}_j^{t_2}, \mathbf{z}, \mathcal{G}^{(1)}(\mathbf{a}_i^{s_1} + \mathbf{a}_j^{t_1}, \mathbf{X})), \quad (5)$$

where $a_i^{s_1} \in \mathcal{A}_s^{(1)}, a_j^{t_1} \in \mathcal{A}_t^{(1)}, a_i^{s_2} \in \mathcal{A}_s^{(2)}, a_j^{t_2} \in \mathcal{A}_t^{(2)}$. As a principled way, anchors can be applied at more levels to encode richer assumptions about future motions.

Training. During training, the model uses *each spatial-temporal anchor* explicitly to generate K future motions for each past motion sequence. The loss functions are mostly adopted as proposed in [51], which we summarize into three categories: (1) reconstruction losses, which optimize *the best predictions* under different definitions among K generated motions, and thus optimize anchors to their own nearest modes; (2) a diversity-promoting loss that explicitly promotes pairwise distances in predictions, avoiding that anchors collapse to the same; and (3) motion constraint losses that encourage output movements to be realistic. All anchors are directly learned from the data via gradient descent. In the forward pass, we explicitly take *every* anchor $\mathbf{a}_i \in \mathcal{A} = \{\mathbf{a}_k\}_{k=1}^K$ as an additional input to

the network and produce a total of K outputs. In the backward pass, each anchor is optimized *separately* based on its corresponding outputs and losses, while the backbone network is updated based on the *fused* losses from all outputs. This separate backward pass is automatically done via PyTorch [56]. Please refer to the supplementary material for more details.

3.2 Interaction-Enhanced Spatial-Temporal Graph Convolutional Network

In principle, our proposed anchor-based sampling permits flexible network architecture. Here, to incorporate our multi-level spatial-temporal anchors, we naturally represent motion sequences as spatial-temporal graphs (to be precise, spatial-frequency graphs), instead of the widely used spatial graphs [51, 52]. Our approach builds upon the Discrete Cosine Transform (DCT) [51, 52] to transform the motion into the frequency domain. Specifically, given a past motion $\mathbf{X}_{1:T_h} \in \mathbb{R}^{T_h \times V \times C^{(0)}}$, where each pose has V joints, we first replicate the last pose T_p times to get $\mathbf{X}_{1:T_h+T_p} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{T_h}, \mathbf{x}_{T_h}, \dots, \mathbf{x}_{T_h}]^T$. With predefined M basis $\mathbf{C} \in \mathbb{R}^{M \times (T_h+T_p)}$ for DCT, the motion is transformed as

$$\tilde{\mathbf{X}} = \mathbf{C}\mathbf{X}_{1:T_h+T_p}. \quad (6)$$

We formulate $\tilde{\mathbf{X}} \in \mathbb{R}^{M \times V \times C^{(0)}}$ in the 0-th layer and latent features in any l -th graph layer as spatial-temporal graphs $(\mathcal{V}^{(l)}, \mathcal{E}^{(l)})$ with $M \times V$ nodes. We specify the node i by 2D index (f_i, v_i) for joint v_i with frequency f_i component. The edge $(i, j) \in \mathcal{E}^{(l)}$ associated with the interaction between node i and node j is represented by $\mathbf{Adj}^{(l)}[i][j]$, where the adjacency matrix $\mathbf{Adj}^{(l)} \in \mathbb{R}^{MV \times MV}$ is learnable. We bottleneck spatial-temporal interactions as [61], by factorizing the adjacency matrix into the product of low-rank spatial and temporal matrices $\mathbf{Adj}^{(l)} = \mathbf{Adj}_s^{(l)} \mathbf{Adj}_f^{(l)}$. The spatial adjacency matrix $\mathbf{Adj}_s^{(l)} \in \mathbb{R}^{MV \times MV}$ connects only nodes with the same frequency. And the frequency adjacency matrix $\mathbf{Adj}_f^{(l)} \in \mathbb{R}^{MV \times MV}$ is merely responsible for the interplay between the nodes representing the same joint.

The spatial-temporal graph can be conveniently encoded by a graph convolutional network (GCN). Given a set of trainable weights $\mathbf{W}^{(l)} \in \mathbb{R}^{C^{(l)} \times C^{(l+1)}}$ and activation function $\sigma(\cdot)$, such as ReLU, a spatial-temporal graph convolutional layer projects the input from $C^{(l)}$ to $C^{(l+1)}$ dimensions by

$$\mathbf{H}_k^{(l+1)} = \sigma(\mathbf{Adj}^{(l)} \mathbf{H}_k^{(l)} \mathbf{W}^{(l)}) = \sigma(\mathbf{Adj}_s^{(l)} \mathbf{Adj}_f^{(l)} \mathbf{H}_k^{(l)} \mathbf{W}^{(l)}), \quad (7)$$

where $\mathbf{H}_k^{(l)} \in \mathbb{R}^{MV \times C^{(l)}}$ denotes the latent feature of the prediction $\hat{\mathbf{Y}}_k$ at l -th layer. The backbone consists of multiple graph convolutional layers. After generating predicted DCT coefficients $\tilde{\mathbf{Y}}_k \in \mathbb{R}^{M \times V \times C^{(L)}}$ reshaped from $\mathbf{H}_k^{(L)}$, where $C^{(L)} = 3$, we recover $\hat{\mathbf{Y}}_k$ via Inverse DCT (IDCT) as

$$\hat{\mathbf{Y}}_k = (\mathbf{C}^T \tilde{\mathbf{Y}}_k)_{T_h+1:T_h+T_p}, \quad (8)$$

where the last T_p frames of the recovered sequence represent future poses.

Conceptually, interactions between spatial-temporal nodes should be relatively invariant across layers, and different interactions should not be equally important. For example, we would expect constraints and dependencies between “left arm” and “left forearm,” while the movements of “head” and “left forearm” are relatively independent. We consider it redundant to construct a *complete* spatial-temporal graph for each layer *independently*. Therefore, we introduce cross-layer interaction sharing to share parameters between graphs in different layers, and spatial interaction pruning to prune the complete graph.

Cross-Layer Interaction Sharing. Much care has been taken into employing learnable interactions between spatial nodes across all graph layers [52, 61, 67]. We consider the spatial relationship to be relatively unchanged. Empirically, we find that sharing the adjacency matrix at intervals of one layer is effective. As shown in Fig. 3, we set $\mathbf{Adj}_s^{(4)} = \mathbf{Adj}_s^{(6)} = \mathbf{Adj}_s^{(8)}$ and $\mathbf{Adj}_s^{(5)} = \mathbf{Adj}_s^{(7)}$.

Spatial Interaction Pruning. To emphasize the physical relationships and constraints between spatial joints, we prune the spatial connections $\widehat{\mathbf{Adj}}_s^{(l)} = \mathbf{M}_s \odot \mathbf{Adj}_s^{(l)}$ in every graph layer l using a predefined mask \mathbf{M}_s , where \odot is an element-wise product. Inspired by [47], we emphasize spatial locality based on skeletal connections and mirror symmetry tendencies. We denote our proposed predefined mask matrix as

$$\mathbf{M}_s[i][j] = \begin{cases} 1, & v_i \text{ and } v_j \text{ are physically connected, } f_i = f_j \\ 1, & v_i \text{ and } v_j \text{ are mirror-symmetric, } f_i = f_j \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

Finally, our architecture consists of four original **STGCNs** without spatial pruning and four **Pruned STGCNs**, as illustrated in Fig. 3. Please refer to the supplementary material for more information of the architecture.

4 Experiments

4.1 Experimental Setup for Diverse Prediction

Datasets. We perform evaluation on two motion capture datasets, Human3.6M [30] and HumanEva-I [60]. Human3.6M consists of 11 subjects and 3.6 million frames at 50 Hz. Following [51, 72], we use a 17-joint skeleton representation and train our model to predict 100 future frames given 25 past frames without global translation. We train on five subjects (S1, S5, S6, S7, and S8) and test on two subjects (S9 and S11). HumanEva-I contains 3 subjects recorded at 60 Hz. Following [51, 72], the pose is represented by 15 joints. We use the official train/test split [60]. The model forecasts 60 future frames given 15 past frames.

Metrics. For a fair comparison, we measure the diversity and accuracy of the predictions according to the evaluation metrics in [2, 51, 70, 72]. (1) **Average Pairwise Distance (APD)**: average ℓ_2 distance between all prediction pairs, defined as $\frac{1}{K(K-1)} \sum_{i=1}^K \sum_{j \neq i}^K \|\hat{\mathbf{Y}}_i - \hat{\mathbf{Y}}_j\|_2$. (2) **Average Displacement Error**

(**ADE**): average ℓ_2 distance over the time between the ground truth and the closest prediction, computed as $\frac{1}{T_p} \min_k \|\hat{\mathbf{Y}}_k - \mathbf{Y}\|_2$. (3) **Final Displacement Error (FDE)**: ℓ_2 distance of the last frame between the ground truth and the closest prediction, defined as $\min_k \|\hat{\mathbf{Y}}_k[T_p] - \mathbf{Y}[T_p]\|_2$. To measure the ability to produce multi-modal predictions, we also report multi-modal versions of ADE and FDE. We define the multi-modal ground truth [70] as $\{\mathbf{Y}_n\}_{n=1}^N$, which is clustered based on historical pose distances, representing possible multi-modal future motions. The detail of multi-modal ground truth is in the supplementary material. (4) **Multi-Modal ADE (MMADE)**: the average displacement error between the predictions and the multi-modal ground truth, denoted as $\frac{1}{NT_p} \sum_{n=1}^N \min_k \|\hat{\mathbf{Y}}_k - \mathbf{Y}_n\|_2$. (5) **Multi-Modal FDE (MMFDE)**: the final displacement error between the predictions and the multi-modal ground truth, denoted as $\frac{1}{N} \sum_{n=1}^N \min_k \|\hat{\mathbf{Y}}_k[T_p] - \mathbf{Y}_n[T_p]\|_2$. All metrics here are in *meters*.

Baselines. To evaluate our stochastic motion prediction method, we consider two types of baselines: (1) Stochastic methods, including CVAE-based methods, **Pose-Knows** [64] and **MT-VAE** [66], as well as CGAN-based methods, **HP-GAN** [4]; (2) Diversity promoting methods, including **Best-of-Many** [5], **GM-VAE** [17], **DeLiGAN** [23], **DSF** [70], **DLow** [72], **MOJO** [75], and **GSPS** [51].

Implementation Details. The backbone consists of 8 GCN layers. We perform spatial pruning on 4 GCN layers (denoted as ‘Pruned’). The remaining 4 layers are not pruned. In each layer, we use batch normalization [29] and residual connections. We add K spatial-temporal compositional anchors at layers 4 and 6, and random sampling at the layer 5. Here, $K = 50$ unless otherwise specified. For Human3.6M, the model is trained for 500 epochs, with a batch size of 16 and 5000 training instances per epoch. For HumanEva-I, the model is trained for 500 epochs, with a batch size of 16 and 2000 training instances per epoch. Additional implementation details are provided in the supplementary material.

4.2 Quantitative Results and Ablation of Diverse Prediction

We compare our method with the baselines in Table 1 on Human3.6M and HumanEva-I. We produce one prediction using *each spatial-temporal anchor* for a total of 50 predictions, which is consistent with the literature [51, 72]. For all metrics, our method consistently outperforms all baselines on both datasets. Methods such as GMVAE [17] and DeLiGAN [23] have relatively low accuracy (ADE, FDE, MMADE, and MMFDE) and diversity (APD), since they still follow a pure random sampling. Methods such as DSF [70], DLow [72] and GSPS [51] explicitly promote diversity by introducing assumptions in the latent codes or directly in the generation process. Instead, we propose to use anchors to locate diverse modes directly learned from the data, which is more effective.

Effectiveness of multi-level spatial-temporal anchors. As shown in Table 2, compared with not using spatial-temporal decoupling (I), using it (II and III) leads to relatively lower diversity, but facilitates mode capture and results in higher accuracy on Human3.6M. Applying the multi-level mechanism (IV) improves diversity, but sacrifices a little accuracy on Human3.6M. By contrast,

Table 1: Quantitative results on Human3.6M and HumanEva-I for $K = 50$. Our model significantly outperforms all stochastic prediction baselines on all metrics. The results of baselines are reported from [51, 72, 75]

Method	Human3.6M [30]					HumanEva-I [60]				
	APD \uparrow	ADE \downarrow	FDE \downarrow	MMADE \downarrow	MMFDE \downarrow	APD \uparrow	ADE \downarrow	FDE \downarrow	MMADE \downarrow	MMFDE \downarrow
Pose-Knows [64]	6.723	0.461	0.560	0.522	0.569	2.308	0.269	0.296	0.384	0.375
MT-VAE [66]	0.403	0.457	0.595	0.716	0.883	0.021	0.345	0.403	0.518	0.577
HP-GAN [4]	7.214	0.858	0.867	0.847	0.858	1.139	0.772	0.749	0.776	0.769
BoM [5]	6.265	0.448	0.533	0.514	0.544	2.846	0.271	0.279	0.373	0.351
GMVAE [17]	6.769	0.461	0.555	0.524	0.566	2.443	0.305	0.345	0.408	0.410
DeLiGAN [23]	6.509	0.483	0.534	0.520	0.545	2.177	0.306	0.322	0.385	0.371
DSF [70]	9.330	0.493	0.592	0.550	0.599	4.538	0.273	0.290	0.364	0.340
DLow [72]	11.741	0.425	0.518	0.495	0.531	4.855	0.251	0.268	0.362	0.339
MOJO [75]	12.579	0.412	0.514	0.497	0.538	4.181	0.234	0.244	0.369	0.347
GSPS [51]	14.757	0.389	0.496	0.476	0.525	5.825	0.233	0.244	0.343	0.331
STARS (Ours)	15.884	0.358	0.445	0.442	0.471	6.031	0.217	0.241	0.328	0.321

Table 2: Ablation study on Human3.6M and HumanEva-I for $K = 50$. We compare the following 4 cases: (I) 50 original anchors; (II) 2 temporal anchors and 25 spatial anchors; (III) 50 spatial-temporal compositional anchors from 5 temporal anchors and 10 spatial anchors; (IV) 50 spatial-temporal compositional anchors for both levels

# of anchors	Human3.6M [30]					HumanEva-I [60]				
	APD \uparrow	ADE \downarrow	FDE \downarrow	MMADE \downarrow	MMFDE \downarrow	APD \uparrow	ADE \downarrow	FDE \downarrow	MMADE \downarrow	MMFDE \downarrow
(I) 50	16.974	0.363	0.447	0.444	0.473	7.786	0.221	0.249	0.327	0.321
(II) 2×25	16.303	0.356	0.442	0.440	0.468	5.254	0.224	0.253	0.337	0.331
(III) 5×10	13.681	0.355	0.442	0.439	0.467	6.199	0.226	0.255	0.334	0.330
(IV) $(5 \times 10) \times 2$	15.884	0.358	0.445	0.442	0.471	6.031	0.217	0.241	0.328	0.321

we observe improvements in both diversity and accuracy on HumanEva-I. The results suggest that there is an intrinsic trade-off between diversity and accuracy. Higher diversity indicates that the model has a better chance of covering multiple modes. However, when the diversity exceeds a certain level, the trade-off between diversity and accuracy becomes noticeable.

Impact of number of anchors and samples. We investigate the effect of two important hyperparameters on the model, *i.e.*, the number of anchors and the number of samples. As illustrated in Fig. 4(a), we fix the number of samples to 50 and compare the results when the number of anchors varies within 0, 5, 10, 25, 50. The results show that more anchors enable the model to better capture the major modes (ADE, FDE) and also other modes (MMADE, MMFDE). In Fig. 4(b), we vary the sample size to be 10, 20, 50, 100 and keep the number of anchors the same as the number of samples. The results show that the larger the number of samples is, the easier it is for a sample to approach the ground truth.

Generalizability of anchor-based sampling. In Table 3, we demonstrate that our anchor-based sampling is *model-agnostic* and can be inserted as a *plug-and-play* module into different motion predictors. Concretely, we apply our anchor-based sampling to the baseline method GSPS [51], which also achieves




Fig. 4: Ablation study on Human3.6M. We report ADE, MMADE, FDE, and MMFDE, comparing settings with different numbers of anchors and samples

Table 3: Ablation study on Human3.6M for $K = 100$. We demonstrate the generalizability of our anchor-based sampling. For a fair comparison, we add single-level anchor-based sampling to GPS [51] and IE-STGCN, without changing any other design and without using spatial-temporal decomposition. We observe that our anchor-based sampling mechanism consistently improves diversity and accuracy for both approaches. Meanwhile, our backbone is more lightweight but performs better

Backbone	Parameter	random sampling					single-level anchor-based sampling (Ours)				
		APD \uparrow	ADE \downarrow	FDE \downarrow	MMADE \downarrow	MMFDE \downarrow	APD \uparrow	ADE \downarrow	FDE \downarrow	MMADE \downarrow	MMFDE \downarrow
GPS [51]	1.30M	14.751	0.369	0.462	0.455	0.491	19.022	0.364	0.443	0.443	0.465
IE-STGCN (Ours)	0.29M	11.701	0.352	0.445	0.443	0.476	14.554	0.344	0.411	0.423	0.436

consistent improvements under every metric, with improvements in terms of diversity and multi-modal accuracy being particularly evident. For simplicity, this evaluation uses simple single-level anchors, but the improvements are pronounced. We would also like to emphasize that the total number of parameters in our IE-STGCN predictor is only **22%** of that in GPS.

4.3 Qualitative Results of Diverse Prediction

We visualize the start pose, the end pose of the ground truth future motions, and the end pose of 10 motion samples in Fig. 5. The qualitative comparisons support the ADE results in Table 1 that our best predicted samples are closer to the ground truth. In Fig. 6 and Fig. 7, we provide the predicted sequences sampled every ten frames. As mentioned before, our spatial-temporal anchors provide a new kind of control over spatial-temporal aspects. With the same temporal anchor, the motion frequencies are similar, but the motion patterns are different. Conversely, if we control the spatial anchors to be the same, the motion trends are similar, but the speed might be different. We show a smooth control through the linear interpolation of spatial-temporal anchors in Fig. 7. The new interpolated anchors produce some interesting and valid pose sequences. And smooth changes in spatial trend and temporal velocity can be observed.

4.4 Effectiveness on Deterministic Prediction

Our model can be easily extended to deterministic prediction by specifying $K = 1$. Without diverse sampling, we retrain two deterministic prediction model




Fig. 5: Visualization of end poses on Human3.6M. We show the historical poses in red and black skeletons, and the predicted end poses with purple and green. As highlighted by the red and blue dashed boxes, the best predictions of our method are closer to the ground truth than the state-of-the-art baseline GSPS [51]




Fig. 6: Visualization of controllable motion prediction on Human3.6M and HumanEva-I. We control different trends and speeds of motions by controlling spatial and temporal anchors. For example, the third and fourth rows have similar motion trends, but the motion in the third row is faster

variants: IE-STGCN-Short dedicated to short-term prediction and IE-STGCN-Long for long-term prediction. We use different settings for deterministic prediction, following existing work [16, 52, 61] and for fair comparisons. Here, we evaluate on Human3.6M and use the 22-joint skeleton representations. Given a 400 ms historical motion sequence, the model generates a 400 ms motion for short-term prediction and a 1000 ms motion for long-term prediction. We use five subjects (S1, S6, S7, S8 and S9) for training and a subject 5 (S5) for testing. We compare our two model variants with recent state-of-the-art deterministic prediction baselines: **LTD** [52], **STS-GCN** [61], and **MSR-GCN** [16]. We evaluate this by reporting Mean Per Joint Position Error (MPJPE) [30] in *millimeter* at each time step, defined as $\frac{1}{V} \sum_{i=1}^V \|\hat{\mathbf{y}}_t[i] - \mathbf{y}_t[i]\|_2$, where $\hat{\mathbf{y}}_t[i]$ and $\mathbf{y}_t[i]$ are produced and ground truth 3D positions of the i -th joint at time t . Table 4 includes short-term (80~400 ms) and long-term (560~1000 ms) comparisons, showing




Fig. 7: Linear interpolation of anchors. We seamlessly control different trends and speeds of future motions by linear interpolation of spatial and temporal anchors. Specifically, given two anchors \mathbf{a}_1 and \mathbf{a}_2 and a coefficient α , we produce predictions from the interpolated anchor formulated as $(1 - \alpha)\mathbf{a}_1 + \alpha\mathbf{a}_2$

Table 4: Quantitative Results on Human3.6M for $K = 1$. Both our long-term and short-term deterministic models significantly outperform all deterministic baselines

Method	Short-term prediction				Long-term prediction			
	80	160	320	400	560	720	880	1000
LTD [52]	11.2	23.3	47.9	59.3	79.9	94.3	106.1	113.3
STS-GCN [61]	13.5	27.7	54.4	65.8	85.0	98.3	108.9	117.0
MSR-GCN [16]	11.3	24.3	50.8	61.9	80.0	93.8	105.5	112.9
IE-STGCN-Short (Ours)	9.7	21.2	44.5	55.5	77.1	91.1	102.6	110.1
IE-STGCN-Long (Ours)	10.0	21.8	45.7	56.9	75.8	89.3	100.8	108.4

that our models outperform the baseline models on both short-term and long-term horizons. Additional experimental results and implementation details of our two deterministic prediction models are provided in the supplementary material.

5 Conclusion

In this paper, we present a simple yet effective approach, STARS, to predict multiple plausible and diverse future motions. And our spatial-temporal anchors enable novel controllable motion prediction. To incorporate our spatial-temporal anchors, we propose a novel motion predictor IE-STGCN. Extensive experiments on Human3.6M and HumanEva-I show the state-of-the-art performance of our unified approach for both diverse and deterministic motion predictions. In the future, we will consider human-scene interaction and investigate the integration of our predictor into human-robot interaction systems.

Acknowledgement. This work was supported in part by NSF Grant 2106825, the Jump ARCHES endowment through the Health Care Engineering Systems Center, the New Frontiers Initiative, the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign through the NCSA Fellows program, and the IBM-Illinois Discovery Accelerator Institute.

References

1. Aksan, E., Kaufmann, M., Hilliges, O.: Structured prediction helps 3D human motion modelling. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7144–7153 (2019) [3](#)
2. Aliakbarian, S., Saleh, F.S., Salzmann, M., Petersson, L., Gould, S.: A stochastic conditioning scheme for diverse human motion prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5223–5232 (2020) [4](#), [9](#)
3. Bai, S., Kolter, J.Z., Koltun, V.: An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. arXiv preprint arXiv:1803.01271 (2018) [3](#)
4. Barsoum, E., Kender, J.R., Liu, Z.: HP-GAN: Probabilistic 3D Human Motion Prediction via GAN. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 1418–1427 (2018) [4](#), [10](#), [11](#)
5. Bhattacharyya, A., Schiele, B., Fritz, M.: 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. pp. 8485–8493 (2018) [10](#), [11](#)
6. Bütetage, J., Black, M.J., Kragic, D., Kjellstrom, H.: Deep representation learning for human motion prediction and classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6158–6166 (2017) [3](#)
7. Bütetage, J., Kjellström, H., Kragic, D.: Anticipating many futures: Online human motion prediction and generation for human-robot interaction. In: IEEE International Conference on Robotics and Automation. pp. 4563–4570 (2018) [1](#)
8. Cao, Z., Gao, H., Mangalam, K., Cai, Q.Z., Vo, M., Malik, J.: Long-term human motion prediction with scene context. In: European Conference on Computer Vision. pp. 387–404. Springer (2020) [3](#)
9. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: European Conference on Computer Vision. pp. 213–229. Springer (2020) [4](#)
10. Chai, Y., Sapp, B., Bansal, M., Anguelov, D.: Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. arXiv preprint arXiv:1910.05449 (2019) [4](#)
11. Chao, Y.W., Yang, J., Price, B., Cohen, S., Deng, J.: Forecasting human dynamics from static images. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 548–556 (2017) [3](#)
12. Corona, E., Pumarola, A., Alenya, G., Moreno-Noguer, F.: Context-aware human motion prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6992–7001 (2020) [3](#)
13. Cui, H., Radosavljevic, V., Chou, F.C., Lin, T.H., Nguyen, T., Huang, T.K., Schneider, J., Djuric, N.: Multimodal trajectory predictions for autonomous driving using deep convolutional networks. In: International Conference on Robotics and Automation. pp. 2090–2096 (2019) [4](#)
14. Cui, Q., Sun, H.: Towards accurate 3D human motion prediction from incomplete observations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4801–4810 (2021) [4](#)
15. Cui, Q., Sun, H., Yang, F.: Learning dynamic relationships for 3D human motion prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6519–6527 (2020) [3](#)

16. Dang, L., Nie, Y., Long, C., Zhang, Q., Li, G.: MSR-GCN: Multi-scale residual graph convolution networks for human motion prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11467–11476 (2021) [4](#), [13](#), [14](#)
17. Dilokthanakul, N., Mediano, P.A., Garnelo, M., Lee, M.C., Salimbeni, H., Arulkumaran, K., Shanahan, M.: Deep unsupervised clustering with Gaussian mixture variational autoencoders. arXiv preprint arXiv:1611.02648 (2016) [10](#), [11](#)
18. Fragkiadaki, K., Levine, S., Felsen, P., Malik, J.: Recurrent network models for human dynamics. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4346–4354 (2015) [4](#)
19. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. *Advances in Neural Information Processing Systems* **27** (2014) [4](#)
20. Gui, L.Y., Wang, Y.X., Liang, X., Moura, J.M.F.: Adversarial geometry-aware human motion prediction. In: European Conference on Computer Vision. pp. 786–803 (2018) [3](#)
21. Gui, L.Y., Wang, Y.X., Ramanan, D., Moura, J.M.F.: Few-shot human motion prediction via meta-learning. In: European Conference on Computer Vision. pp. 432–450 (2018) [3](#)
22. Gui, L.Y., Zhang, K., Wang, Y.X., Liang, X., Moura, J.M.F., Veloso, M.: Teaching robots to predict human motion. In: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 562–567 (2018) [3](#)
23. Gurumurthy, S., Kiran Sarvadevabhatla, R., Venkatesh Babu, R.: DeLiGAN : Generative adversarial networks for diverse and limited data. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 166–174 (2017) [10](#), [11](#)
24. Han, T., Xie, W., Zisserman, A.: Memory-augmented dense predictive coding for video representation learning. In: European Conference on Computer Vision. pp. 312–329 (2020) [4](#)
25. Hassan, M., Ceylan, D., Villegas, R., Saito, J., Yang, J., Zhou, Y., Black, M.: Stochastic scene-aware motion prediction. In: Proceedings of the International Conference on Computer Vision. pp. 11374–11384 (2021) [3](#)
26. Hernandez, A., Gall, J., Moreno-Noguer, F.: Human motion prediction via spatio-temporal inpainting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7134–7143 (2019) [4](#)
27. Holden, D., Komura, T., Saito, J.: Phase-functioned neural networks for character control. *ACM Transactions on Graphics* **36**, 1 – 13 (2017) [4](#)
28. Holden, D., Saito, J., Komura, T.: A deep learning framework for character motion synthesis and editing. *ACM Transactions on Graphics* **35**, 1 – 11 (2016) [4](#)
29. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: International Conference on Machine Learning. pp. 448–456. PMLR (2015) [10](#)
30. Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: 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 (2013) [2](#), [9](#), [11](#), [13](#)
31. Jain, A., Zamir, A.R., Savarese, S., Saxena, A.: Structural-RNN: Deep learning on spatio-temporal graphs. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5308–5317 (2016) [3](#)
32. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014) [4](#)

33. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907 (2016) [3](#)
34. Koppula, H.S., Saxena, A.: Anticipating human activities using object affordances for reactive robotic response. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **38**(1), 14–29 (2016) [1](#)
35. Koppula, H.S., Saxena, A.: Anticipating human activities for reactive robotic response. *IEEE/RSJ International Conference on Intelligent Robots and Systems* pp. 2071–2071 (2013) [1](#)
36. Kothari, P., Siffringer, B., Alahi, A.: Interpretable social anchors for human trajectory forecasting in crowds. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 15551–15561 (2021) [4](#)
37. Kundu, J.N., Gor, M., Babu, R.V.: BiHMP-GAN: Bidirectional 3D human motion prediction GAN. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 33, pp. 8553–8560 (2019) [4](#)
38. Lasota, P.A., Shah, J.A.: A multiple-predictor approach to human motion prediction. In: *IEEE International Conference on Robotics and Automation*. pp. 2300–2307 (2017) [1](#)
39. Lebailly, T., Kiciroglu, S., Salzmann, M., Fua, P., Wang, W.: Motion prediction using temporal inception module. In: *Proceedings of the Asian Conference on Computer Vision* (2020) [4](#)
40. Li, C., Zhang, Z., Lee, W.S., Lee, G.H.: Convolutional sequence to sequence model for human dynamics. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5226–5234 (2018) [4](#)
41. Li, M., Chen, S., Zhao, Y., Zhang, Y., Wang, Y., Tian, Q.: Dynamic multiscale graph neural networks for 3D skeleton based human motion prediction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 214–223 (2020) [4](#)
42. Li, X., Li, H., Joo, H., Liu, Y., Sheikh, Y.: Structure from recurrent motion: From rigidity to recurrency. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3032–3040 (2018) [3](#)
43. Lin, X., Amer, M.R.: Human motion modeling using DVGANs. arXiv preprint arXiv:1804.10652 (2018) [4](#)
44. Ling, H.Y., Zinno, F., Cheng, G., Van De Panne, M.: Character controllers using motion VAEs. *ACM Transactions on Graphics* **39**(4), 40–1 (2020) [4](#)
45. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: SSD: Single shot multibox detector. In: *European conference on computer vision*. pp. 21–37. Springer (2016) [4](#)
46. Liu, Y., Zhang, J., Fang, L., Jiang, Q., Zhou, B.: Multimodal motion prediction with stacked transformers. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7577–7586 (2021) [4](#)
47. Liu, Z., Su, P., Wu, S., Shen, X., Chen, H., Hao, Y., Wang, M.: Motion prediction using trajectory cues. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 13299–13308 (2021) [9](#)
48. Luber, M., Stork, J.A., Tipaldi, G.D., Arras, K.O.: People tracking with human motion predictions from social forces. In: *IEEE International conference on robotics and automation*. pp. 464–469 (2010) [1](#)
49. Lyu, K., Liu, Z., Wu, S., Chen, H., Zhang, X., Yin, Y.: Learning human motion prediction via stochastic differential equations. In: *Proceedings of ACM International Conference on Multimedia*. pp. 4976–4984 (2021) [3](#)

50. Mao, W., Liu, M., Salzmann, M.: History repeats itself: Human motion prediction via motion attention. In: European Conference on Computer Vision. pp. 474–489 (2020) [4](#)
51. Mao, W., Liu, M., Salzmann, M.: Generating smooth pose sequences for diverse human motion prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13309–13318 (2021) [2](#), [3](#), [4](#), [7](#), [8](#), [9](#), [10](#), [11](#), [12](#), [13](#)
52. Mao, W., Liu, M., Salzmann, M., Li, H.: Learning trajectory dependencies for human motion prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9489–9497 (2019) [4](#), [8](#), [9](#), [13](#), [14](#)
53. Martinez, J., Black, M.J., Romero, J.: On human motion prediction using recurrent neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2891–2900 (2017) [4](#)
54. Paden, B., Cáp, M., Yong, S.Z., Yershov, D.S., Frazzoli, E.: A survey of motion planning and control techniques for self-driving urban vehicles. *IEEE Transactions on Intelligent Vehicles* **1**, 33–55 (2016) [1](#)
55. Pascanu, R., Mikolov, T., Bengio, Y.: On the difficulty of training recurrent neural networks. In: International Conference on Machine Learning. pp. 1310–1318. PMLR (2013) [4](#)
56. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: PyTorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems* **32** (2019) [8](#)
57. Phan-Minh, T., Grigore, E.C., Boulton, F.A., Beijbom, O., Wolff, E.M.: CoverNet: Multimodal behavior prediction using trajectory sets. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14074–14083 (2020) [4](#)
58. Rezende, D., Mohamed, S.: Variational inference with normalizing flows. In: International Conference on Machine Learning. pp. 1530–1538. PMLR (2015) [4](#)
59. Rudenko, A., Palmieri, L., Arras, K.O.: Joint long-term prediction of human motion using a planning-based social force approach. In: IEEE International Conference on Robotics and Automation. pp. 4571–4577 (2018) [1](#)
60. Sigal, L., Balan, A.O., Black, M.J.: 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 (2010) [2](#), [9](#), [11](#)
61. Sofianos, T., Sampieri, A., Franco, L., Galasso, F.: Space-Time-Separable Graph Convolutional Network for pose forecasting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11209–11218 (2021) [4](#), [8](#), [9](#), [13](#), [14](#)
62. Starke, S., Zhao, Y., Zinno, F., Komura, T.: Neural animation layering for synthesizing martial arts movements. *ACM Transactions on Graphics* **40**, 1 – 16 (2021) [1](#)
63. Sutskever, I., Martens, J., Hinton, G.: Generating text with recurrent neural networks. International Conference on Machine Learning pp. 1017–1024 (2011) [3](#)
64. Walker, J., Marino, K., Gupta, A., Hebert, M.: The pose knows: Video forecasting by generating pose futures. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3332–3341 (2017) [4](#), [10](#), [11](#)
65. Wang, B., Adeli, E., Chiu, H.k., Huang, D.A., Niebles, J.C.: Imitation learning for human pose prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7124–7133 (2019) [3](#)

66. Yan, X., Rastogi, A., Villegas, R., Sunkavalli, K., Shechtman, E., Hadap, S., Yumer, E., Lee, H.: MT-VAE: Learning motion transformations to generate multimodal human dynamics. In: European Conference on Computer Vision. pp. 276–293 (2018) [4](#), [10](#), [11](#)
67. Yan, Z., Zhai, D.H., Xia, Y.: DMS-GCN: Dynamic mutiscale spatiotemporal graph convolutional networks for human motion prediction. arXiv preprint arXiv:2112.10365 (2021) [4](#), [9](#)
68. Yang, Y., Ramanan, D.: Articulated pose estimation with flexible mixtures-of-parts. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition pp. 1385–1392 (2011) [4](#)
69. Yu, B., Yin, H., Zhu, Z.: Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. arXiv preprint arXiv:1709.04875 (2017) [4](#)
70. Yuan, Y., Kitani, K.: Diverse trajectory forecasting with determinantal point processes. arXiv preprint arXiv:1907.04967 (2019) [6](#), [9](#), [10](#), [11](#)
71. Yuan, Y., Kitani, K.: Ego-pose estimation and forecasting as real-time PD control. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 10082–10092 (2019) [3](#)
72. Yuan, Y., Kitani, K.: DLow: Diversifying latent flows for diverse human motion prediction. In: European Conference on Computer Vision. pp. 346–364 (2020) [2](#), [3](#), [4](#), [6](#), [9](#), [10](#), [11](#)
73. Yuan, Y., Kitani, K.: Residual force control for agile human behavior imitation and extended motion synthesis. Advances in Neural Information Processing Systems **33**, 21763–21774 (2020) [4](#)
74. Zhang, J.Y., Felsen, P., Kanazawa, A., Malik, J.: Predicting 3D human dynamics from video. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7114–7123 (2019) [3](#)
75. Zhang, Y., Black, M.J., Tang, S.: We are more than our joints: Predicting how 3D bodies move. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3372–3382 (2021) [10](#), [11](#)