

Auxiliary Tasks Benefit 3D Skeleton-based Human Motion Prediction

Chenxin Xu^{1,2}, Robby T. Tan², Yuhong Tan¹, Siheng Chen^{1,3*}, Xinchao Wang², Yanfeng Wang^{3,1}

¹Shanghai Jiao Tong University, ²National University of Singapore, ³Shanghai AI Laboratory

{cxcwakaka, tyheeee, sihengc, wangyanfeng}@sjtu.edu.cn, {robby.tan, xinchao}@nus.edu.sg

Abstract

Exploring spatial-temporal dependencies from observed motions is one of the core challenges of human motion prediction. Previous methods mainly focus on dedicated network structures to model the spatial and temporal dependencies. This paper considers a new direction by introducing a model learning framework with auxiliary tasks. In our auxiliary tasks, partial body joints' coordinates are corrupted by either masking or adding noise and the goal is to recover corrupted coordinates depending on the rest coordinates. To work with auxiliary tasks, we propose a novel auxiliary-adapted transformer, which can handle incomplete, corrupted motion data and achieve coordinate recovery via capturing spatial-temporal dependencies. Through auxiliary tasks, the auxiliary-adapted transformer is promoted to capture more comprehensive spatial-temporal dependencies among body joints' coordinates, leading to better feature learning. Extensive experimental results have shown that our method outperforms state-of-the-art methods by remarkable margins of 7.2%, 3.7%, and 9.4% in terms of 3D mean per joint position error (MPJPE) on the Human3.6M, CMU Mocap, and 3DPW datasets, respectively. We also demonstrate that our method is more robust under data missing cases and noisy data cases. Code is available at <https://github.com/MediaBrain-SJTU/AuxFormer>.

1. Introduction

3D skeleton-based human motion prediction aims to forecast future human motions based on past observations, which has a wide range of applications, such as human-machine interaction [23, 21, 62] and autonomous driving [34, 9, 38]. One of the main challenges of this problem is extracting spatial-temporal dependencies among observed motions to make feature representations more informative. These dependencies arise due to the complex interactions between different joints and the temporal dynamics of motion. Therefore, developing effective methods to capture these dependencies

*Corresponding author.

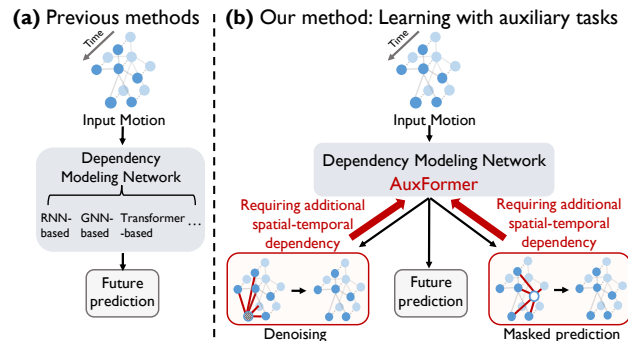


Figure 1: Compared to previous methods with various operation designs, we propose a new direction: an auxiliary learning framework that incorporates auxiliary tasks, including denoising and masking prediction. These auxiliary tasks impose additional requirements through recovery, forcing the network to exploit spatial-temporal dependencies more comprehensively.

is crucial for a more accurate human motion prediction.

Existing methods proposed dedicated network structures to model the spatial and temporal dependencies. A few methods use RNN [48, 22] and TCN [11] structures to model temporal dependencies, but neglect the spatial ones. To learn spatial dependencies between joints, [47] proposes a GCN network with learnable weights where nodes are body joints. Following the GCN design, DMGNN [37] and MSR-GCN [12] further build multiscale body graphs to model local-global spatial features. PGBIG [43] additionally proposes temporal graph convolutions to extract spatial-temporal features. SPGSN [36] proposes a graph scattering network to further model temporal dependencies from multiple graph spectrum bands. Besides GCN-based methods, transformer architectures [1, 5] are also used to model pair-wise spatial-temporal dependencies. The structures of models to capture spatial-temporal dependencies have been extensively studied. This raises a natural question: can we enhance spatial-temporal dependency modeling from other perspectives?

In this paper, besides network structures, we further consider an orthogonal approach: proposing a new auxiliary model learning framework by adding auxiliary tasks to promote better learning of spatial-temporal dependencies. The proposed framework jointly learns the primary motion pre-

diction task along with additional auxiliary tasks, with sharing the same dependency modeling network. The core idea of the auxiliary tasks is to corrupt partial observed coordinates and set a goal to recover corrupted coordinates using correlated normal coordinates according to their spatial-temporal dependencies. The goals of the auxiliary tasks are highly correlated with the primary prediction task, as they both require the network to model spatial-temporal dependencies effectively. Therefore, through the additional requirements imposed by the auxiliary tasks, the dependency modeling network is forced to learn more effective and comprehensive spatial-temporal dependencies. Our learning framework complements existing methods by further emphasizing the effective learning of the network structure.

To be specific, we introduce two kinds of auxiliary tasks: a denoising task and a masked feature prediction task. The denoising task randomly adds noise into joint coordinates at different timestamps in the input motion, and the aim is to recover the original input motion. The masked feature prediction task randomly masks joint coordinates at different timestamps, and the goal is to predict the masked joint positions. Comparing to previous popular methods based on masking/denoising autoencoding like masked autoencoder (MAE) [25] and denoising autoencoder (DAE) [56], which are mainly designed for model pre-training, we treat the denoising and masking prediction as auxiliary tasks to aid the primary fully-supervised motion prediction task and jointly learn all the tasks together. Moreover, previous methods using masking/denoising strategies mostly focus on data types of images [25, 45], videos [16, 53], languages [14, 50, 67] and point clouds [49, 59], but rarely on motion sequences, especially human motions. Our work fills this gap and utilizes the strategies to promote more effective spatial-temporal dependencies learning in human motion prediction.

To cooperate with auxiliary tasks in the learning framework, the dependency modeling network structure faces two demands. First, the network has to learn spatial-temporal dependencies between the corrupted coordinate and the normal coordinate on a coordinate-wise basis to enable recovery. Second, the network has to be adaptive to incomplete motions, caused by the masking prediction task. Thus, we specifically propose an auxiliary-adapted transformer network to meet both two demands. To model the coordinate-wise dependency, we consider each coordinate as one individual feature in our network and use spatial-temporal attention that models spatial-temporal dependencies between two coordinates’ features. To be adaptive to incomplete data, we add tokens into the masked coordinates’ feature to inform the network that the data is missing, and incorporate a mask-aware design into spatial-temporal attention that enables arbitrary incomplete inputs. Integrating the above learning framework and network design, we name our method *AuxFormer*.

We conduct experiments on both short-term and long-

term motion prediction on large-scale datasets: Human3.6M [28], CMUMocap and 3DPW [57]. Our method significantly outperforms state-of-the-art methods in terms of mean per joint position error (MPJPE). We also show our method is more robust under data missing and noisy cases. The main contributions of our work are summarized here:

- We propose a new auxiliary learning framework for human motion prediction to jointly learn the prediction task with two auxiliary tasks: denoising and masking prediction. Through auxiliary tasks, the model network is forced to learn more comprehensive spatial-temporal dependencies.
- We propose an auxiliary-adapted transformer to cooperate with the learning framework. The auxiliary-adapted transformer models coordinate-wise spatial-temporal dependencies and is adaptive to incomplete motion data.
- We conduct experiments to verify that our method significantly outperforms existing works by 7.2%/3.7%/9.4%. We also show our method is more robust under data missing cases and noisy data cases.

2. Related Work

Human Motion Prediction. Early methods for human motion prediction captured shallow temporal dynamics using state prediction [33, 54]. Later, RNN-based models like ERD [18], Pose-VAE [58], Structural-RNN [30], [22], Res-sup [48] and AHMR [42] improved temporal dependency modeling. [24, 35] use spatial convolutions and predict whole sequences without accumulation. More recent approaches use graph convolutional networks [32, 8, 63, 66, 26, 31] to represent human poses as graphs and capture spatial dependencies. Researchers build fully-connected graphs [47, 11], multi-scale graphs [37, 12], and semi-constrained graphs [41] to encode skeletal connections and prior knowledge. [52, 43] extend graph convolution to temporal dimension to model both temporal and spatial dependencies. [36] proposes graph scattering networks to decompose pose features into richer graph spectrum bands. [10] proposes a multi-task graph convolutional network for the incomplete observations cases. EqMotion [64] proposes an equivariant motion prediction model to achieve motion equivariance and interaction invariance based on graphs. Besides GCN-based methods, transformer structures [1, 5] also are adapted to model pairwise spatial-temporal dependencies. Unlike most previous methods that focus on dedicated network structures, we adopt a novel strategy by introducing an auxiliary learning framework to efficiently train a network for better performance.

Denoising and Masked Autoencoding. Both denoising and masked autoencoding aim to learn representative features by reconstructing corrupted data. Denoising autoencoding, first proposed by [56], makes learned representations robust to partial corruption. Different corruptions are further

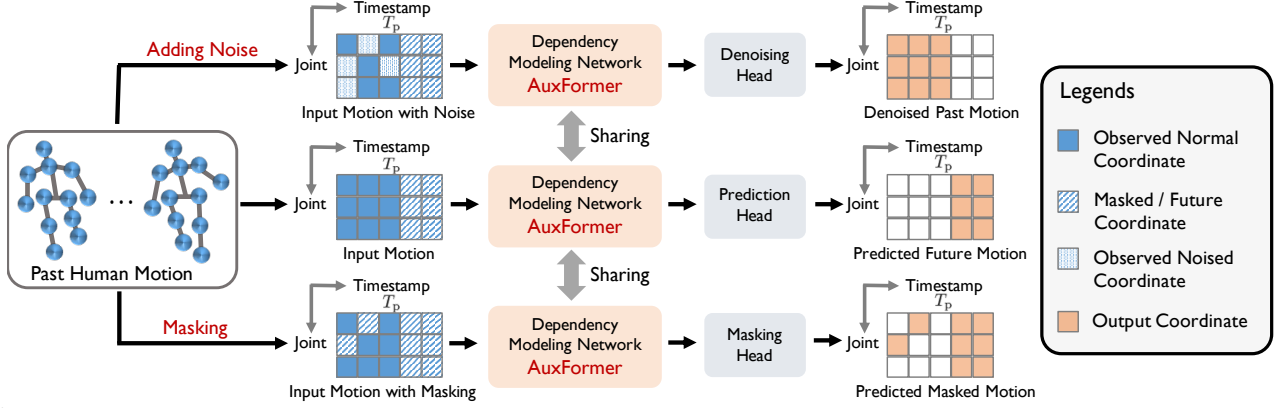


Figure 2: The learning framework architecture. The framework consists of three tasks: the primary future prediction task (middle branch), the auxiliary denoising task (upper branch), and the auxiliary masking prediction task (lower branch). These three tasks share the same dependency modeling network and use different heads.

proposed like gray colorization [70], pixel masking [45], and channel splitting [71]. Masked autoencoding applies region masking as a corruption and is a special form of denoising autoencoding. Motivated by large language pre-training models like BERT [14] and GPT [50, 51], masked image prediction methods are proposed for vision model pre-training with different reconstruction targets like image pixels [7, 15, 61, 2, 40], discrete tokens [4, 73, 44], and deep features [3, 60]. [25] presents the masked autoencoder (MAE) to accelerate model pre-training by masking a high proportion of input images and developing an asymmetric encoder-decoder architecture. This idea is used in other data types such as point clouds [49, 69, 59, 68] and videos [16, 53, 19]. Our method focuses on motion sequence data and incorporates masking/denoising as auxiliary tasks to assist the motion prediction task, different from previous pre-training methods.

Multi-Task Learning and Auxiliary Learning. Multi-task learning [72, 6] aims to improve tasks’ performance by sharing information between multiple relevant tasks. Similarly, auxiliary learning uses multiple auxiliary tasks to assist the primary task, but it only focuses on improving the performance of the primary task. Auxiliary learning has been used in various tasks like semantic segmentation [65], scene understanding [39], image and video captioning [27, 20], speech recognition [55], view synthesis [17], and reinforcement learning [29] based on images, videos, and languages. However, auxiliary learning has not been extensively applied in 3D human motion prediction. To the best of our knowledge, this paper is the first to introduce auxiliary learning to 3D human motion prediction.

3. Prediction Framework with Auxiliary Tasks

3.1. Problem Formulation

The 3D skeleton-based human motion task aims to predict future human motions given past motions. Mathemati-

cally, let $\mathbf{X}^t = [\mathbf{x}_1^t, \mathbf{x}_2^t, \dots, \mathbf{x}_j^t] \in \mathbb{R}^{J \times 3}$ denotes the human pose consisting 3D coordinates of M body joints at timestamp t , where $\mathbf{x}_j^t \in \mathbb{R}^3$ is the j th joint’s coordinate. Let $\mathbb{X}^- = [\mathbf{X}^1, \mathbf{X}^2, \dots, \mathbf{X}^{T_p}] \in \mathbb{R}^{T_p \times J \times 3}$ be the past motion and $\mathbb{X}^+ = [\mathbf{X}^{T_p+1}, \mathbf{X}^{T_p+2}, \dots, \mathbf{X}^{T_p+T_f}] \in \mathbb{R}^{T_f \times J \times 3}$ be the future motion. T_p and T_f are the length of past and future motions. $T = T_p + T_f$ is the total length of motion. To facilitate subsequent explanation, we pad the past motion \mathbb{X}^- to $\mathbb{X} \in \mathbb{R}^{T \times J \times 3}$ with zeros on future timestamps. Our goal is to learn a prediction model $\mathcal{F}_{\text{pred}}(\cdot)$ so that the predicted future motions $\hat{\mathbb{X}}^+ = \mathcal{F}_{\text{pred}}(\mathbb{X})$ are as close to the ground-truth future motions \mathbb{X}^+ as possible. To achieve an accurate prediction, one of the keys is capturing comprehensive spatial-temporal dependencies in the input motion. We are going to capture spatial-temporal dependencies via learning with auxiliary tasks.

3.2. Framework Architecture

Here we propose the model learning framework with multiple auxiliary tasks, which is sketched in Figure 2. The core idea of the framework is to learn the primary prediction task simultaneously with extra-designed auxiliary tasks that raise additional requirements to force the model to capture more comprehensive spatial-temporal dependencies. We introduce two additional auxiliary tasks: masking prediction and denoising. To be specific, taking the past human motion $\mathbb{X} \in \mathbb{R}^{T \times J \times 3}$ as the input, in the masking prediction task, each 3D coordinate \mathbf{x}_j^t in the past motion has a probability of p_m to be masked by setting the value to zeros. The goal of the masking prediction task is to recover the masked coordinates from the remaining observed coordinates. In the denoising task, every 3D coordinate \mathbf{x}_j^t has a probability of p_n to be noisy by adding a Gaussian noise $\epsilon \sim \mathcal{N}(0, \sigma)$, where σ is the deviation. The goal of the denoising task is to erase the noise according to the remaining noiseless coordinates.

To combine the two auxiliary tasks, we propose a joint learning framework that enables the primary future predic-

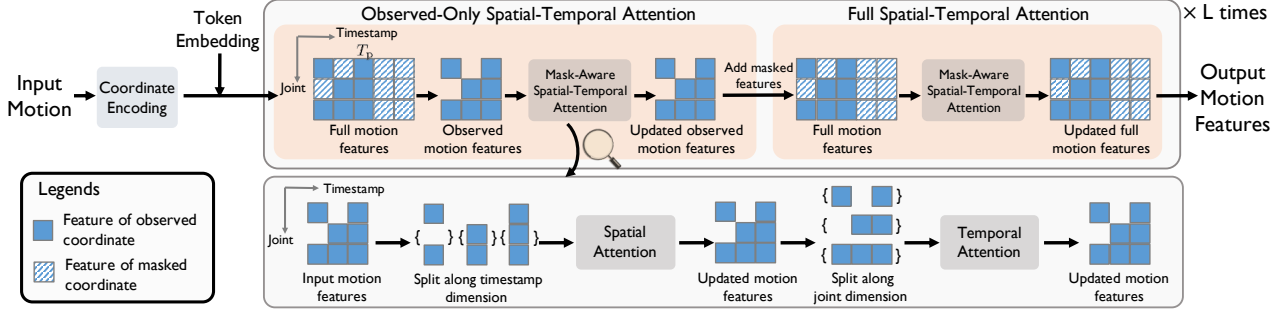


Figure 3: Auxiliary-adapted transformer (AuxFormer) as the dependency modeling network. We take the input of the masked prediction task as an example. The network first encodes coordinates and adds joint, timestamp, and masking information with token embedding. Then the network iteratively employs observed-only and full spatial-temporal attention to update the observed and whole motion features, respectively. Both observed-only and full spatial-temporal attention mechanisms use the same mask-aware spatial-temporal attention mechanism, which incorporates spatial and temporal attention for spatial and temporal dependency modeling.

tion task and the two auxiliary tasks to be learned simultaneously in an end-to-end manner. All three tasks share the same dependency modeling network but use different prediction heads. Mathematically, the mask of masking prediction task $\mathbf{M}_M \in \mathbb{R}^{T \times J}$ is defined as

$$\mathbf{M}_M(t, j) = \begin{cases} 1 & \text{if } t \leq T_p \text{ and } \mathbf{x}_j^t \text{ is unmasked,} \\ 0 & \text{otherwise.} \end{cases}$$

Here we treat the future coordinates also be the "masked" coordinates since it is unknown. Similarly, we also define the mask of normal prediction task \mathbf{M}_P and the mask of denoising task \mathbf{M}_D :

$$\mathbf{M}_P(t, j) = \mathbf{M}_D(t, j) = \begin{cases} 1 & \text{if } t \leq T_p, \\ 0 & \text{otherwise.} \end{cases}$$

Let \mathbb{X} , \mathbb{X}_M and \mathbb{X}_D be the original input motion, input motion with masking and input motion with noise, the joint learning framework is formulated by,

$$\begin{aligned} \mathbb{H} &= \mathcal{F}(\mathbb{X}, \mathbf{M}_P), \widehat{\mathbb{X}} = \mathcal{P}_{\text{pred}}(\mathbb{H}), \\ \mathbb{H}_M &= \mathcal{F}(\mathbb{X}_M, \mathbf{M}_M), \widehat{\mathbb{X}}_M = \mathcal{P}_{\text{mask}}(\mathbb{H}_M), \\ \mathbb{H}_D &= \mathcal{F}(\mathbb{X}_D, \mathbf{M}_D), \widehat{\mathbb{X}}_D = \mathcal{P}_{\text{denoise}}(\mathbb{H}_D), \end{aligned} \quad (1)$$

where $\mathcal{F}(\cdot)$ denotes the auxiliary-adapted transformer (AuxFormer) as the dependency modeling network. \mathbb{H} , \mathbb{H}_M , $\mathbb{H}_D \in \mathbb{R}^{T \times J \times F}$ are motion features of future prediction task, mask prediction task and denoising task, respectively. F is the dimension of the coordinate feature. $\mathcal{P}_{\text{pred}}(\cdot)$, $\mathcal{P}_{\text{mask}}(\cdot)$, and $\mathcal{P}_{\text{denoise}}(\cdot)$ are three prediction heads, which are simple linear or MLP functions that map the F -dimensional space to the 3D coordinate space, and $\widehat{\mathbb{X}}$, $\widehat{\mathbb{X}}_M$, $\widehat{\mathbb{X}}_D \in \mathbb{R}^{T \times J \times 3}$ are the output sequences of the three tasks, which contains predicted future motions, predicted masked motions and denoised past motions, respectively.

Note that i) both the masking prediction task and the denoising task require the network to recover corrupted coordinates using their correlated coordinates according to their spatial-temporal dependencies. Through the two tasks, the dependency modeling framework is given extra force to-

ward a more comprehensive spatial-temporal dependency modeling; ii) the dependency modeling network contains most of the model parameters and the prediction head is lightweight. Thus introducing extra auxiliary tasks into the learning framework only slightly increase the model size; iii) in the inference time we only perform the primary future prediction task by taking the future timestamps of the future prediction sequence as the predicted output $\widehat{\mathbb{X}}^+ = \widehat{\mathbb{X}}_{[T_p+1:T_f]}$.

Compared to [10] that combines motion repairing with the motion prediction task, our framework has two major differences. First, [10] aims to predict motions from incomplete observations; while we consider the classical setting of motion prediction. Second, the motion repairing task in [10] is a necessary, non-splittable part of the model and will be included in the inference phase; while our auxiliary tasks serve as adjunctive, parallel branches to enhance model training and will not be involved in the inference phase.

3.3. Loss Function

For different tasks, we supervise different parts of the prediction sequences to satisfy the task demands. To be specific, for the prediction/masking/denoising task, we supervise the future/masked/past coordinates. Mathematically, assuming the masking set is $\mathcal{M} = \{(j, t) | \mathbf{M}(j, t) = 0\}$, given the output sequence $\widehat{\mathbb{X}}$, $\widehat{\mathbb{X}}_M$, $\widehat{\mathbb{X}}_D$ of three tasks, the overall loss function is formulated by

$$\mathcal{L}_{\text{prediction}} = \frac{1}{T_f J} \sum_{t=T_p+1}^{T_f} \sum_{j=1}^J \|\widehat{\mathbf{x}}_j^t - \mathbf{x}_j^t\|^2,$$

$$\mathcal{L}_{\text{mask}} = \frac{1}{|\mathcal{M}|} \sum_{(j,t) \in \mathcal{M}} \|\widehat{\mathbf{x}}_{M,j}^t - \mathbf{x}_j^t\|^2,$$

$$\mathcal{L}_{\text{denoise}} = \frac{1}{T_p J} \sum_{t=1}^{T_p} \sum_{j=1}^J \|\widehat{\mathbf{x}}_{D,j}^t - \mathbf{x}_j^t\|^2,$$

$$\mathcal{L} = \mathcal{L}_{\text{prediction}} + \alpha_1 \mathcal{L}_{\text{mask}} + \alpha_2 \mathcal{L}_{\text{denoise}}.$$

We use the average ℓ_2 distance between different targets and predictions. α_1 and α_2 are weight hyperparameters.

4. Auxiliary-Adapted Transformer

This section introduces an auxiliary-adapted transformer network to implement the dependency modeling network $\mathcal{F}(\cdot)$ in eq.(1). To work with auxiliary tasks, the network design faces two demands. First, it must learn spatial-temporal dependencies between corrupted and normal coordinates on a coordinate-wise basis to enable recovery. Second, the network must be adaptable to incomplete motions caused by the masking prediction task. To the first demand, we separately encode each 3D coordinate into an individual feature and use spatial-temporal attention to model coordinate-wise spatial-temporal dependencies. To the second demand, we add a masked token into the masked coordinate’s feature to indicate it is masked and incorporate a mask-aware design into the spatial-temporal attention mechanism that enables inputs with arbitrary missing positions.

4.1. Overview

Mathematically, given the observed motion of an arbitrary task $\mathbb{X} \in \mathbb{R}^{T \times J \times 3}$ and its corresponding mask $\mathbf{M} \in \mathbb{R}^{T \times J}$ as the input, our auxiliary-adapted transformer works as,

$$\mathbb{E} = \mathcal{F}_{\text{CE}}(\mathbb{X}), \quad (2a)$$

$$\mathbb{H} = \mathcal{F}_{\text{TE}}(\mathbb{E}, \mathbf{M}), \quad (2b)$$

$$\mathbb{H}' = \mathcal{F}_{\text{OSTA}}^{(l)}(\mathbb{H}, \mathbf{M}), \quad (2c)$$

$$\mathbb{H}'' = \mathcal{F}_{\text{FSTA}}^{(l)}(\mathbb{H}', \mathbf{1}), \quad (2d)$$

where $\mathbb{E}, \mathbb{H}, \mathbb{H}', \mathbb{H}'' \in \mathbb{R}^{T \times J \times F}$ are motion features, where \mathbb{E} and \mathbb{H} are initial and embedded features, while \mathbb{H}' and \mathbb{H}'' are updated features. F is the feature dimension and $\mathbf{1} \in \mathbb{R}^{T \times J}$ is a mask matrix with all one values..

Step (2a) uses a coordinate encoder $\mathcal{F}_{\text{CE}}(\cdot)$ to obtain the initial motion features. Step (2b) embeds joint, timestamp, and masked information into the initial coordinate features to obtain the embedded coordinate features using a token embedding function $\mathcal{F}_{\text{TE}}(\cdot)$. Steps (2c) apply observed-only spatial-temporal attention $\mathcal{F}_{\text{OSTA}}^{(l)}(\cdot)$ to update the observed coordinate features. Steps (2d) apply full spatial-temporal attention $\mathcal{F}_{\text{FSTA}}^{(l)}(\cdot)$ to update the whole coordinate features. Step (2c) and (2d) are repeated iteratively L times. The observed-only and full spatial-temporal attention have the same mask-aware spatial temporal attention mechanism with the only difference of the input mask. The whole auxiliary-adapted transformer network is sketched in Figure 3.

Note that: i) instead of solely using full spatial-temporal attention, we additionally incorporate observed-only spatial-temporal attention. The intuition is modeling spatial-temporal dependencies inner observed coordinates provides more correlations to improve the ability to infer spatial-temporal dependencies between observed and masked coordinates; ii) instead of performing observed-only and full spatial-temporal attention separately for L times, we employ

an iterative approach because interaction with masked coordinate features can also enhance observed features. We evaluate the effects of these design choices in Section 5.4.

4.2. Structure Details

We now elaborate the details of each step.

Coordinate encoding. Coordinate encoding aims to map each 3D coordinate to a high-dimensional embedding space for subsequent feature learning. Here we simply use a linear layer to implement the coordinate encoder $\mathcal{F}_{\text{CE}}(\cdot)$.

Token embedding. In token embedding, we encode the joint, timestamp and masking information. To encode joint information, we consider dictionary learning to convert the joint number j into a learnable code $\mathbf{w}_{J,j} \in \mathbb{R}^F$, which represents the j th element of the joint embedding dictionary $\mathbf{W}_J \in \mathbb{R}^{J \times F}$. Similarly, we convert the timestamp t into a learnable code $\mathbf{w}_{T,t} \in \mathbb{R}^F$. For the masked or future coordinates, we add a learnable masked token $\mathbf{w}_M \in \mathbb{R}^F$. Mathematically, the embedded coordinate feature for the j th joint of the t th timestamp is

$$\mathbf{h}_j^t = \begin{cases} \mathbf{e}_j^t + \mathbf{w}_{T,t} + \mathbf{w}_{J,j} & \text{if } \mathbf{M}(t, j) = 1, \\ \mathbf{w}_M + \mathbf{w}_{T,t} + \mathbf{w}_{J,j} & \text{if } \mathbf{M}(t, j) = 0. \end{cases}$$

\mathbf{e}_j^t is the initial feature of the t th timestamp of the j th joint. **Mask-aware spatial-temporal attention.** The mask-aware spatial-temporal attention is used to model spatial-temporal dependencies among assigned features based on the input mask, which implements $\mathcal{F}_{\text{OSTA}}^{(l)}(\cdot)$ and $\mathcal{F}_{\text{FSTA}}^{(l)}(\cdot)$ in eq.(2c) and eq.(2d) respectively. As the spatial relationship between joints and the temporal relation between timestamps have different patterns, we use spatial and temporal attention mechanisms to model the spatial and temporal dependencies separately. Spatial attention considers features of the same timestamp, while temporal attention considers features of the same joint.

The spatial attention simultaneously operates features at different timestamps. For the feature $\mathbf{H}^t \in \mathbb{R}^{J \times D}$ of t th timestamp, we first compute its spatial attention matrix $\mathbf{A}_s \in \mathbb{R}^{J \times J}$ according to the input mask,

$$\mathbf{A}_s(j_1, j_2) = \begin{cases} 1 & \text{if } \mathbf{M}(t, j_1) = 1 \text{ and } \mathbf{M}(t, j_2) = 1 \\ 0 & \text{otherwise,} \end{cases}$$

where $\mathbf{A}_s(j_1, j_2)$ indicates whether performing attention between joint j_1 and j_2 . Then the spatial attention is

$$\begin{aligned} \mathbf{Q}^t, \mathbf{K}^t, \mathbf{V}^t &= f_{\text{QKV}}(\mathbf{H}^t), \\ \mathbf{O}^t &= [\text{Softmax}(\frac{\mathbf{Q}^t \mathbf{K}^{t \top}}{\sqrt{F}}) \cdot \mathbf{A}_s] \mathbf{V}^t, \\ \mathbf{H}^t &\leftarrow \mathbf{H}^t + \text{FFN}(\parallel_{i=1}^H \mathbf{O}_i^t), \end{aligned}$$

where $f_{\text{QKV}}(\cdot)$ are linear operations, $\text{Softmax}(\cdot)$ represents the softmax function. We use a regular multi-head feed-forward operation that produces totally H heads $\{\mathbf{O}_i^t\}$ in

Table 1: Comparisons of short-term prediction on Human3.6M. Results at 80ms, 160ms, 320ms and 400ms in the future are shown. **Bold/underline** font represent the best/second best result.

Motion	Walking				Eating				Smoking				Discussion			
	80ms	160ms	320ms	400ms	80ms	160ms	320ms	400ms	80ms	160ms	320ms	400ms	80ms	160ms	320ms	400ms
millisecond	35.4	57.3	76.3	87.7	34.5	63.4	124.6	142.5	38.0	69.3	115.0	126.7	36.1	69.1	130.5	157.1
Res-sup. [48]	29.4	50.8	76.0	81.5	16.8	30.6	56.9	68.7	23.0	42.6	70.1	82.7	32.9	61.2	90.9	96.2
Traj-GCN [47]	12.3	23.0	39.8	46.1	8.4	16.9	33.2	40.7	7.9	16.2	31.9	38.9	12.5	27.4	58.5	71.7
DMGNN [37]	17.3	30.7	54.6	65.2	11.0	21.4	36.2	43.9	9.0	17.6	32.1	40.3	17.3	34.8	61.0	69.8
MSRGCN [12]	12.2	22.7	38.6	45.2	8.4	17.1	33.0	40.4	8.0	16.3	31.3	38.2	12.0	26.8	57.1	69.7
PGBIG [43]	10.2	19.8	34.5	40.3	7.0	15.1	30.6	38.1	6.6	14.1	28.2	34.7	10.0	23.8	53.6	66.7
SPGSN [36]	10.1	19.4	34.8	41.5	7.1	14.9	30.5	37.9	6.7	13.8	28.0	34.6	10.4	23.8	53.6	67.1
AuxFormer (Ours)	8.9	16.9	30.1	36.1	6.4	14.0	28.8	35.9	5.7	11.4	22.1	27.9	8.6	18.8	38.8	49.2
Motion	Directions				Greeting				Phoning				Posing			
	80ms	160ms	320ms	400ms	80ms	160ms	320ms	400ms	80ms	160ms	320ms	400ms	80ms	160ms	320ms	400ms
millisecond	35.4	57.3	76.3	87.7	34.5	63.4	124.6	142.5	38.0	69.3	115.0	126.7	36.1	69.1	130.5	157.1
Res-sup. [48]	9.0	19.9	43.4	53.7	18.7	38.7	77.7	93.4	10.2	21.0	42.5	52.3	13.7	29.9	66.6	84.1
Traj-GCN [47]	13.1	24.6	64.7	81.9	23.3	50.3	107.3	132.1	12.5	25.8	48.1	58.3	15.3	29.3	71.5	96.7
DMGNN [37]	8.6	19.7	43.3	53.8	16.5	37.0	77.3	93.4	10.1	20.7	41.5	51.3	12.8	29.4	67.0	85.0
MSRGCN [12]	7.2	17.6	40.9	51.5	15.2	34.1	71.6	87.1	8.3	18.3	38.7	48.4	10.7	25.7	60.0	76.6
PGBIG [43]	7.4	17.2	39.8	50.3	14.6	32.6	70.6	86.4	8.7	18.3	38.7	48.5	10.7	25.3	59.9	76.5
SPGSN [36]	12.8	28.6	61.0	74.4	9.3	19.4	42.3	53.6	14.2	27.7	56.8	70.7	8.8	18.9	41.5	52.7
AuxFormer (Ours)	6.8	17.0	40.3	51.6	13.5	31.3	69.2	85.4	7.9	17.3	37.4	47.2	8.8	19.1	39.2	51.0
Motion	Purchases				Sitting				Sittingdown				Takingphoto			
	80ms	160ms	320ms	400ms	80ms	160ms	320ms	400ms	80ms	160ms	320ms	400ms	80ms	160ms	320ms	400ms
millisecond	36.3	60.3	86.5	95.9	42.6	81.4	134.7	151.8	47.3	86.0	145.8	168.9	26.1	47.6	81.4	94.7
Res-sup. [48]	15.6	32.8	65.7	79.3	10.6	21.9	46.3	57.9	16.1	31.1	61.5	75.5	9.9	20.9	45.0	56.6
Traj-GCN [47]	21.4	38.7	75.7	92.7	11.9	25.1	44.6	50.2	15.0	32.9	77.1	93.0	13.6	29.0	46.0	58.8
DMGNN [37]	14.8	32.4	66.1	79.6	10.5	22.0	46.3	57.8	16.1	31.6	62.5	76.8	9.9	21.0	44.6	56.3
MSRGCN [12]	12.5	28.7	60.1	73.3	8.8	19.2	42.4	53.8	13.9	27.9	57.4	71.5	8.4	18.9	42.0	53.3
PGBIG [43]	12.8	28.6	61.0	74.4	9.3	19.4	42.3	53.6	14.2	27.7	56.8	70.7	8.8	18.9	41.5	52.7
SPGSN [36]	11.9	28.0	61.8	76.3	8.7	19.0	42.1	53.3	13.5	27.6	57.7	72.2	8.2	18.4	41.5	53.0
AuxFormer (Ours)	11.9	28.0	61.8	76.3	8.7	19.0	42.1	53.3	13.5	27.6	57.7	72.2	8.2	18.4	41.5	53.0
Motion	Walking				Walking Dog				Walking Together				Average			
	80ms	160ms	320ms	400ms	80ms	160ms	320ms	400ms	80ms	160ms	320ms	400ms	80ms	160ms	320ms	400ms
millisecond	30.6	57.8	106.2	121.5	64.2	102.1	141.1	164.4	26.8	50.1	80.2	92.2	34.7	62.0	101.1	115.5
Res-sup. [48]	11.4	24.0	50.1	61.5	23.4	46.2	83.5	96.0	10.5	21.0	38.5	45.2	12.7	26.1	52.3	63.5
Traj-GCN [47]	12.2	24.2	59.6	77.5	47.1	93.3	160.1	171.2	14.3	26.7	50.1	63.2	17.0	33.6	65.9	79.7
DMGNN [37]	10.7	23.1	48.3	59.2	20.7	42.9	80.4	93.3	10.6	20.9	37.4	43.9	12.1	25.6	51.6	62.9
MSRGCN [12]	8.9	20.1	43.6	54.3	18.8	39.3	73.7	86.4	8.7	18.6	34.4	41.0	10.3	22.7	47.4	58.5
PGBIG [43]	9.2	19.8	43.1	54.1	17.8	37.2	71.7	84.9	8.9	18.2	33.8	40.9	10.4	22.3	47.1	58.3
SPGSN [36]	8.2	18.5	41.2	52.2	17.1	36.5	70.4	83.0	7.8	15.9	30.2	37.0	9.5	20.6	43.4	54.1
AuxFormer (Ours)	8.2	18.5	41.2	52.2	17.1	36.5	70.4	83.0	7.8	15.9	30.2	37.0	9.5	20.6	43.4	54.1

Table 2: Comparisons of long-term prediction on 7 representative actions on Human3.6M. Results at 560ms and 1000ms in the future are shown. **Bold/underline** font represent the best/second best result.

Motion	Walking		Smoking		Discussion		Greeting		Posing		Walking Dog		Walking Together		Average	
	560ms	1000ms	560ms	1000ms	560ms	1000ms	560ms	1000ms	560ms	1000ms	560ms	1000ms	560ms	1000ms	560ms	1000ms
millisecond	81.7	100.7	94.8	137.4	121.3	161.7	156.3	184.3	165.7	236.8	173.6	202.3	94.5	110.5	129.2	165.0
Res-Sup. [48]	54.1	59.8	50.7	72.6	91.6	121.5	115.4	148.8	114.5	173.0	111.9	148.9	55.0	65.6	81.6	114.3
Traj-GCN [47]	71.4	85.8	50.9	72.2	81.9	138.3	144.5	170.5	125.5	188.2	183.2	210.2	70.5	86.9	93.6	127.6
DMGNN [37]	52.7	63.0	49.5	71.6	88.6	117.6	116.3	147.2	116.3	174.3	111.9	148.2	52.9	65.9	81.1	114.2
MSRGCN [12]	48.1	56.4	46.5	69.5	87.1	118.2	110.2	143.5	106.1	164.8	104.7	139.8	51.9	64.3	76.9	110.3
PGBIG [43]	46.9	53.6	46.7	68.6	89.7	118.6	111.0	143.2	110.3	165.4	102.4	138.0	49.8	60.9	77.4	109.6
SPGSN [36]	43.8	52.0	42.0	63.0	77.6	102.3	110.5	141.6	91.6	137.1	103.3	133.3	47.3	58.8	75.3	107.0
AuxFormer (Ours)	43.8	52.0	42.0	63.0	77.6	102.3	110.5	141.6	91.6	137.1	103.3	133.3	47.3	58.8	75.3	107.0

parallel and use a feedforward function $\text{FFN}(\cdot)$ to obtain the output of spatial attention. \parallel denotes concatenation.

Similarly, temporal attention simultaneously operates features of different joints. For the j th joint's feature $\mathbf{H}_j \in \mathbb{R}^{T \times D}$, we compute its temporal attention matrix $\mathbf{A}_t \in \mathbb{R}^{T \times T}$ by

$$\mathbf{A}_t(t_1, t_2) = \begin{cases} 1 & \text{if } \mathbf{M}(j, t_1) = 1, \text{ and } \mathbf{M}(j, t_2) = 1, \\ 0 & \text{otherwise.} \end{cases}$$

The temporal attention is

$$\begin{aligned} \mathbf{Q}_j, \mathbf{K}_j, \mathbf{V}_j &= f_{\text{QKV}}(\mathbf{H}_j), \\ \mathbf{O}_j &= [\text{Softmax}(\frac{\mathbf{Q}_j \mathbf{K}_j^\top}{\sqrt{D}}) \cdot \mathbf{A}_t] \mathbf{V}_j, \\ \mathbf{H}_j &\leftarrow \mathbf{H}_j + \text{FFN}(\parallel_{i=1}^H \mathbf{O}_{j,i}). \end{aligned}$$

Compared to the previous spatial-temporal transformer structure used in [1], our auxiliary-adapted transformer has two advantages: adaptability to missing data cases, which the previous method cannot handle, and modeling of global spatial-temporal dependencies between arbitrary timestamps,

while previous method only models temporal dependencies within a time window.

5. Experiment

5.1. Datasets

Human3.6M [28] has 7 subjects performing 15 types of actions with 22 body joints. The sequences are downsampled temporally by 2 and converted to 3D coordinates, excluding global rotation and translation of the pose. Subjects S5 and S11 are reserved for testing and validation, respectively, while the remaining subjects are used for training.

CMU-MoCap has 8 human action types including 38 body joints which are also converted to 3D coordinates. The global rotations and translations of the poses are excluded. Following [13, 47], we keep 25 joints and divide the training and testing datasets.

3DPW [57] is a dataset for human pose prediction containing indoor and outdoor activities. They are represented in the 3D space and each subject has 23 joints. The frame rate of the motions is 30Hz.

Table 3: Prediction MPJPEs of methods on CMU Mocap for both short-term and long-term prediction across 7 actions, as well as their average prediction results all on actions. **Bold/underline** font represent the best/second best result.

Motion millisecond	Basketball					Basketball Signal					Jumping					Running				
	80	160	320	400	1000	80	160	320	400	1000	80	160	320	400	1000	80	160	320	400	1000
Res-sup. [48]	15.45	26.88	43.51	49.23	88.73	20.17	32.98	42.75	44.65	60.57	26.85	48.07	93.50	108.90	162.84	25.76	48.91	88.19	100.80	158.19
DMGNN [37]	15.57	28.72	59.01	73.05	138.62	5.03	9.28	20.21	26.23	52.04	31.97	54.32	96.66	119.92	224.63	17.42	26.82	38.27	40.08	46.40
Traj-GCN [47]	11.68	21.26	40.99	50.78	97.99	3.33	6.25	13.58	17.98	54.00	17.18	32.37	60.12	72.55	127.41	14.53	24.20	37.44	41.10	51.73
MSR-GCN [12]	10.28	18.94	37.68	47.03	86.96	3.03	5.68	12.35	16.26	47.91	14.99	28.66	55.86	69.05	124.79	12.84	20.42	30.58	34.42	48.03
STSGCN [52]	12.56	23.04	41.92	50.33	94.17	4.72	6.69	14.53	17.88	49.52	17.52	31.48	58.74	72.06	127.40	16.70	27.58	36.15	36.42	55.34
PGBIG [43]	<u>9.53</u>	<u>17.53</u>	35.32	44.23	84.14	2.71	4.88	10.77	14.63	50.19	<u>13.93</u>	<u>27.78</u>	55.80	69.01	125.60	12.69	23.18	38.31	42.24	51.71
SPGSN [36]	10.24	18.54	38.22	48.68	89.58	2.91	5.25	11.31	15.01	47.31	14.93	28.16	56.72	71.16	125.20	<u>10.75</u>	<u>16.67</u>	<u>26.07</u>	<u>30.08</u>	52.92
AuxFormer (Ours)	9.35	17.06	<u>35.46</u>	<u>45.50</u>	80.77	2.62	4.79	10.57	14.20	48.19	12.79	25.52	53.37	67.28	124.34	9.98	15.78	25.31	28.81	41.64
Motion millisecond	Soccer					Walking					Washing Window					Average				
	80	160	320	400	1000	80	160	320	400	1000	80	160	320	400	1000	80	160	320	400	1000
Res-sup. [48]	17.75	31.30	52.55	61.40	107.37	44.35	76.66	126.83	151.43	194.33	22.84	44.71	86.78	104.68	202.73	24.74	44.21	76.30	88.73	139.30
DMGNN [37]	14.86	25.29	52.21	65.42	111.90	9.57	15.53	26.03	30.37	67.01	7.93	14.68	33.34	44.24	82.84	14.07	24.44	45.90	55.45	104.33
Traj-GCN [47]	13.33	24.00	43.77	53.20	108.26	6.62	10.74	17.40	20.35	34.41	5.96	11.62	24.77	31.63	66.85	9.94	18.02	33.55	40.95	81.85
MSR-GCN [12]	10.92	19.50	37.05	46.38	99.32	6.31	10.30	17.64	21.12	39.70	5.49	11.07	25.05	32.51	71.30	8.72	15.83	30.57	38.10	79.01
STSGCN [52]	13.49	25.24	39.87	51.58	109.63	7.18	10.99	17.84	22.61	44.12	6.79	12.10	24.92	36.66	69.48	10.80	18.19	31.18	41.05	81.76
PGBIG [43]	11.09	20.62	39.48	48.72	99.98	6.23	10.34	16.84	19.76	33.92	4.63	9.16	20.87	27.34	65.69	8.20	15.41	30.13	37.27	76.69
SPGSN [36]	10.86	18.99	35.05	45.16	99.51	6.32	<u>10.21</u>	<u>16.34</u>	<u>20.19</u>	34.83	4.86	9.44	<u>21.50</u>	<u>28.37</u>	65.08	8.30	14.80	28.64	36.96	77.82
AuxFormer (Ours)	10.01	18.21	<u>36.31</u>	<u>45.79</u>	95.98	5.76	9.16	15.69	18.80	34.81	<u>4.69</u>	<u>9.39</u>	21.87	28.83	72.90	7.54	13.78	27.95	35.39	76.32

Table 4: The average prediction MPJPEs on 3DPW dataset. **Bold/underline** font represent the best/second best result.

millisecond	Average MPJPE					
	100	200	400	600	800	1000
Res-sup. [48]	102.28	113.24	173.94	191.47	201.39	210.58
CSM [35]	57.83	71.53	124.01	155.16	174.87	187.06
Traj-GCN [47]	16.28	35.62	67.46	90.36	106.79	117.84
DMGNN [37]	17.80	37.11	70.38	94.12	109.67	123.93
HisRep [46]	15.88	35.14	66.82	93.55	107.63	114.75
MSR-GCN [12]	15.70	33.48	65.02	93.81	108.15	116.31
STSGCN [52]	18.32	37.79	67.51	92.75	106.65	112.22
PGBIG [43]	17.66	35.32	67.83	<u>89.60</u>	<u>102.59</u>	<u>109.41</u>
SPGSN [36]	<u>15.39</u>	<u>32.91</u>	<u>64.54</u>	91.62	103.98	111.05
AuxFormer (Ours)	14.21	30.04	58.50	89.45	100.78	107.45

5.2. Model and Experimental Settings

Implementation details. We set the feature dimension F as 64/96/80 in Human3.6M/CMU Mocap/3DPW dataset. For the short/long-term motion prediction, the number of attention layers L is 3/4 in the dependency modeling network. We use 8 heads in the multi-head attention. We use a mask ratio of 50% and a noisy ratio of 30%. The loss weight α_1, α_2 are both set to 1. The model is trained for 50 epochs with a batch size of 16. We use the Adam optimizer to train the model on a single NVIDIA RTX-3090 GPU. We scale the input motion by 1/100 and rescale the output motion. To obtain a more generalized evaluation with lower test bias, we use all the clips in the 5th subject of H3.6M and the test folder of CMU Mocap, instead of testing on a few samples like in [22, 35, 37]. For more implementation details, please refer to the supplementary material.

Evaluation Metrics. We use the Mean Per Joint Position Error (MPJPE), which calculates the average ℓ_2 distance between predicted joints and target ones at each prediction timestamp in 3D Euclidean space. Compared to previous mean angle error (MAE) [37, 48], the MPJPE reflects larger degrees of freedom of human poses and covers larger ranges of errors for a clearer comparison.

5.3. Main Results

Short-term prediction. Short-term motion prediction aims to predict the poses within 400 milliseconds. Table 1 presents




Figure 4: Visualization results of different methods on H3.6M.

the MPJPEs of our method and many previous methods on all actions at multiple prediction timestamps on the H3.6M dataset. We see that i) our method achieves superior performance at most of the timestamps and very close performance to the best results on other timestamps; ii) our method achieves significantly lower MPJPEs on average by 8.7%/7.6%/7.9%/7.2% at 80/160/320/400ms.

Besides, Table 3 reports the comparison of our method and many previous methods on the CMU Mocap dataset. We show our method achieves more effective prediction and has a significant improvement by 9.2%/6.9%/3.3%/3.7% at 80/160/320/400ms. In Table 4, we present the average MPJPEs across all the test samples at different prediction timestamps on the 3DPW dataset. Compared to the state-of-the-art methods, our method greatly reduces the MPJPE by 7.7%/8.7%/9.4% at 100/200/400ms.

Long-term prediction. Long-term motion prediction aims to predict the poses over 400 milliseconds. Table 2 shows prediction MPJPEs of various methods at 560 and 1000 ms. Due to the space limit, we show the results of seven representative actions and the average results of all actions. More detailed tables are provided in the supplementary material. We see that our method achieves a more accurate prediction on most of the actions and our method has a lower MPJPE by 2.1%/2.4% at 560/1000ms compared to state-of-the-art methods. Table 3 and 4 also report the comparison of our method with previous methods on the CMU Mocap and 3DPW datasets. We see that our method also significantly outperforms the state-of-the-art method on the CMU Mocap

Table 5: Ablation study on auxiliary tasks on H3.6M.

Pred	Mask	Denoise	80ms	160ms	320ms	400ms	560ms	1000ms
✓			10.2	21.8	45.7	57.0	79.7	110.5
✓	✓		9.7	21.0	44.0	54.8	77.7	108.2
✓		✓	9.6	20.9	44.2	55.0	76.6	108.1
✓	✓	✓	9.5	20.6	43.4	54.1	75.3	107.0

Table 6: Ablation study on auxiliary tasks on CMU Mocap.

Pred	Mask	Denoise	80ms	160ms	320ms	400ms	1000ms
✓			8.57	15.48	30.40	38.35	75.94
✓	✓		8.25	14.78	29.13	36.46	74.13
✓		✓	8.14	14.57	28.98	36.60	72.95
✓	✓	✓	7.89	14.27	28.36	35.60	71.23




Figure 5: Effect of different masking ratios and noise ratios in the model training on H3.6M.

and the 3DPW dataset.

Qualitative results. Figure 4 shows an example of the predicted poses of different methods. We can see our prediction result is more accurate than two representative baselines, especially on the motion of two arms.

5.4. Ablation Study and Model Analysis

Effect of auxiliary tasks. We investigate the impact of two auxiliary tasks in our learning framework: masking prediction (Mask) and denoising (Denoise). Table 5 and 6 show the results on the H3.6M and CMU Mocap datasets, respectively. "Pred" denotes the future prediction task. We observe that i) adding the masking prediction and denoising tasks individually leads to a significant improvement in the model performance, demonstrating the effectiveness of these auxiliary tasks; ii) combining the two auxiliary tasks further improves the model's performance.

Effect of masking & noise ratios. We explore the effect of different masking ratios and noise ratios during model training. Figure 5 shows the results. The noise deviation is set to 50. We see that i) setting masking and noising ratios that are either too low or too high results in suboptimal performance, as the auxiliary tasks become either too easy or too difficult, rendering them unsuitable for model learning; ii) a moderate masking and noising ratio yielded the best results, with an appropriate masking ratio range of 0.3 to 0.7 and an appropriate noising ratio range of 0.2 to 0.6.

Ablation on model structure. We explore the effect of designs in our dependency modeling network structure. Table

Table 7: Ablation study on network structure design on H3.6M.

Ablation	80ms	160ms	320ms	400ms	AVG
w/o $\mathcal{F}_{OSTA}(\cdot)$	9.7	21.6	46.0	57.3	33.7
Separate attention	9.5	20.7	44.1	55.1	32.4
Iterative attention (Ours)	9.5	20.6	43.4	54.1	31.9




Figure 6: Comparison of model performance under missing data case and noisy data case on H3.6M.

7 shows the results on H3.6M dataset. "w/o $\mathcal{F}_{OSTA}(\cdot)$ " denotes not use the observed-only spatial-temporal attention. "Separate attention" denotes repeating $\mathcal{F}_{OSTA}(\cdot)$ for L times followed by $\mathcal{F}_{FSTA}(\cdot)$ repeated for L times. We can see that adding the observed-only spatial-temporal attention and using an iterative design on both attentions result in more effective motion prediction.

Robustness analysis under incomplete & noisy data cases.

We evaluate our method's robustness under two scenarios that frequently arise in real-world applications: incomplete and noisy data. For incomplete data, we randomly mask input motion coordinates during the testing phase, with various ratios. To enable the baseline methods to handle missing values in the input motion, we apply linear interpolation and extrapolation to fill in the missing values before feeding the motion sequence into the model. Figure 6 (a) shows the MPJPE at 400ms of our method and state-of-the-art baseline methods under different missing data ratios. We see that i) our method achieves the best prediction performance under all masking ratios; ii) with the increase of the masking ratio, the performance gap between our method and baseline methods widens, showing that our method is more robust with missing data than previous methods.

In the noisy data cases, we randomly add noises to different ratios of input motion coordinates in the test phase. Each coordinate's noise is sampled from Gaussian distribution $\mathcal{N}(0, 50)$. Figure 6 (b) shows the comparison of MPJPE at 400ms under different noisy data ratios. We see that i) our proposed method achieves the best prediction performance under all noisy data ratios; ii) previous methods suffer greatly even when the noisy ratio is small, such as 10% and 20%, while our method's performance only experiences a slight drop in these cases and still maintains promising performance, showing better robustness.

6. Conclusion

In this work, we propose AuxFormer, a novel auxiliary model learning framework with an auxiliary-adaptive transformer network, to promote more comprehensive spatial-temporal dependency learning for 3D human motion prediction. The learning framework jointly learns the primary future prediction task with additional auxiliary tasks. To cooperate with the learning framework, we further propose an auxiliary-adapted transformer to model coordinate-wise spatial-temporal dependencies and be adaptive to missing data. We evaluate our method on three human motion prediction datasets and show our method achieves state-of-the-art prediction performance and equips with higher robustness.

Acknowledgements

This research is supported by NSFC under Grant 62171276, the Science and Technology Commission of Shanghai Municipal under Grant 21511100900 and 22DZ2229005, and the National Research Foundation, Singapore under its Medium Sized Center for Advanced Robotics Technology Innovation. Robby T. Tan's work is supported by MOE AcRF Tier 1, A-0009455-01-00.

References

- [1] Emre Aksan, Manuel Kaufmann, Peng Cao, and Otmar Hilliges. A spatio-temporal transformer for 3d human motion prediction. In *2021 International Conference on 3D Vision (3DV)*, pages 565–574. IEEE, 2021. 1, 2, 6
- [2] Roman Bachmann, David Mizrahi, Andrei Atanov, and Amir Zamir. Multima: Multi-modal multi-task masked autoencoders. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVII*, pages 348–367. Springer, 2022. 3
- [3] Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatao Gu, and Michael Auli. Data2vec: A general framework for self-supervised learning in speech, vision and language. In *International Conference on Machine Learning*, pages 1298–1312. PMLR, 2022. 3
- [4] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Bert: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. 3
- [5] Yujun Cai, Lin Huang, Yiwei Wang, Tat-Jen Cham, Jianfei Cai, Junsong Yuan, Jun Liu, Xu Yang, Yiheng Zhu, Xiaohui Shen, et al. Learning progressive joint propagation for human motion prediction. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VII 16*, pages 226–242. Springer, 2020. 1, 2
- [6] Rich Caruana. *Multitask learning*. Springer, 1998. 3
- [7] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pretraining from pixels. In *International conference on machine learning*, pages 1691–1703. PMLR, 2020. 3
- [8] Ming Chen, Zhewei Wei, Zengfeng Huang, Bolin Ding, and Yaliang Li. Simple and deep graph convolutional networks. In *International conference on machine learning*, pages 1725–1735. PMLR, 2020. 2
- [9] Siheng Chen, Baoan Liu, Chen Feng, Carlos Vallespi-Gonzalez, and Carl Wellington. 3d point cloud processing and learning for autonomous driving: Impacting map creation, localization, and perception. *IEEE Signal Processing Magazine*, 38(1):68–86, 2020. 1
- [10] Qiongjie Cui and Huaijiang Sun. Towards accurate 3d human motion prediction from incomplete observations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4801–4810, 2021. 2, 4
- [11] Qiongjie Cui, Huaijiang Sun, and Fei Yang. Learning dynamic relationships for 3d human motion prediction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6519–6527, 2020. 1, 2
- [12] Lingwei Dang, Yongwei Nie, Chengjiang Long, Qing Zhang, and Guiqing Li. Msr-gcn: Multi-scale residual graph convolution networks for human motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11467–11476, 2021. 1, 2, 6, 7
- [13] Lingwei Dang, Yongwei Nie, Chengjiang Long, Qing Zhang, and Guiqing Li. Msr-gcn: Multi-scale residual graph convolution networks for human motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11467–11476, October 2021. 6
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2, 3
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3
- [16] Christoph Feichtenhofer, Haoqi Fan, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. *arXiv preprint arXiv:2205.09113*, 2022. 2, 3
- [17] John Flynn, Ivan Neulander, James Philbin, and Noah Snavely. Deepstereo: Learning to predict new views from the world's imagery. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5515–5524, 2016. 3
- [18] Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. Recurrent network models for human dynamics. In *Proceedings of the IEEE international conference on computer vision*, pages 4346–4354, 2015. 2
- [19] Valentin Gabeur, Arsha Nagrani, Chen Sun, Karteek Alahari, and Cordelia Schmid. Masking modalities for cross-modal video retrieval. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1766–1775, 2022. 3
- [20] Lianli Gao, Yu Lei, Pengpeng Zeng, Jingkuan Song, Meng Wang, and Heng Tao Shen. Hierarchical representation network with auxiliary tasks for video captioning and video question answering. *IEEE Transactions on Image Processing*, 31:202–215, 2021. 3

- [21] Dominic Gorecky, Mathias Schmitt, Matthias Loskyll, and Detlef Zühlke. Human-machine-interaction in the industry 4.0 era. In *2014 12th IEEE international conference on industrial informatics (INDIN)*, pages 289–294. Ieee, 2014. [1](#)
- [22] Liang-Yan Gui, Yu-Xiong Wang, Xiaodan Liang, and José MF Moura. Adversarial geometry-aware human motion prediction. In *Proceedings of the european conference on computer vision (ECCV)*, pages 786–803, 2018. [1](#), [2](#), [7](#)
- [23] Liang-Yan Gui, Kevin Zhang, Yu-Xiong Wang, Xiaodan Liang, José MF Moura, and Manuela Veloso. Teaching robots to predict human motion. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 562–567. IEEE, 2018. [1](#)
- [24] Xiao Guo and Jongmoo Choi. Human motion prediction via learning local structure representations and temporal dependencies. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 2580–2587, 2019. [2](#)
- [25] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. [2](#), [3](#)
- [26] Duc N.M Hoang, Shiwei Liu, Radu Marculescu, and Zhangyang Wang. Revisiting pruning at initialization through the lens of ramanujan graph. In *International Conference on Learning Representations*, 2023. [2](#)
- [27] Mehrdad Hosseinzadeh and Yang Wang. Image change captioning by learning from an auxiliary task. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2725–2734, 2021. [3](#)
- [28] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. 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](#), [6](#)
- [29] Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*, 2016. [3](#)
- [30] Ashesh Jain, Amir R Zamir, Silvio Savarese, and Ashutosh Saxena. Structural-rnn: Deep learning on spatio-temporal graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5308–5317, 2016. [2](#)
- [31] Ajay Kumar Jaiswal, Shiwei Liu, Tianlong Chen, Ying Ding, and Zhangyang Wang. Graph laddling: Shockingly simple parallel gnn training without intermediate communication. In *International Conference on Machine Learning*, 2023. [2](#)
- [32] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. [2](#)
- [33] Andreas M Lehrmann, Peter V Gehler, and Sebastian Nowozin. Efficient nonlinear markov models for human motion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1314–1321, 2014. [2](#)
- [34] Jesse Levinson, Jake Askeland, Jan Becker, Jennifer Dolson, David Held, Soeren Kammel, J Zico Kolter, Dirk Langer, Oliver Pink, Vaughan Pratt, et al. Towards fully autonomous driving: Systems and algorithms. In *2011 IEEE intelligent vehicles symposium (IV)*, pages 163–168. IEEE, 2011. [1](#)
- [35] Chen Li, Zhen Zhang, Wee Sun Lee, and Gim Hee Lee. Convolutional sequence to sequence model for human dynamics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5226–5234, 2018. [2](#), [7](#)
- [36] Maosen Li, Siheng Chen, Zijing Zhang, Lingxi Xie, Qi Tian, and Ya Zhang. Skeleton-parted graph scattering networks for 3d human motion prediction. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VI*, pages 18–36. Springer, 2022. [1](#), [2](#), [6](#), [7](#)
- [37] Maosen Li, Siheng Chen, Yangheng Zhao, Ya Zhang, Yanfeng Wang, and Qi Tian. 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*, pages 214–223, 2020. [1](#), [2](#), [6](#), [7](#)
- [38] Ming Liang, Bin Yang, Rui Hu, Yun Chen, Renjie Liao, Song Feng, and Raquel Urtasun. Learning lane graph representations for motion forecasting. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 541–556. Springer, 2020. [1](#)
- [39] Lukas Liebel and Marco Körner. Auxiliary tasks in multi-task learning. *arXiv preprint arXiv:1805.06334*, 2018. [3](#)
- [40] Songhua Liu, Kai Wang, Xingyi Yang, Jingwen Ye, and Xinchao Wang. Dataset distillation via factorization. In *Advances in neural information processing systems*, 2022. [3](#)
- [41] Zhenguang Liu, Pengxiang Su, Shuang Wu, Xuanjing Shen, Haipeng Chen, Yanbin Hao, and Meng Wang. Motion prediction using trajectory cues. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 13299–13308, 2021. [2](#)
- [42] Zhenguang Liu, Shuang Wu, Shuyuan Jin, Shouling Ji, Qi Liu, Shijian Lu, and Li Cheng. Investigating pose representations and motion contexts modeling for 3d motion prediction. *IEEE transactions on pattern analysis and machine intelligence*, 45(1):681–697, 2022. [2](#)
- [43] Tiezheng Ma, Yongwei Nie, Chengjiang Long, Qing Zhang, and Guiqing Li. Progressively generating better initial guesses towards next stages for high-quality human motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6437–6446, 2022. [1](#), [2](#), [6](#), [7](#)
- [44] Xinyin Ma, Gongfan Fang, and Xinchao Wang. LLM-Pruner: On the Structural Pruning of Large Language Models. *arXiv preprint arXiv:2305.11627*, 2023. [3](#)
- [45] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens Van Der Maaten. Exploring the limits of weakly supervised pretraining. In *Proceedings of the European conference on computer vision (ECCV)*, pages 181–196, 2018. [2](#), [3](#)
- [46] Wei Mao, Miaomiao Liu, and Mathieu Salzmann. History repeats itself: Human motion prediction via motion attention. In *Computer Vision—ECCV 2020: 16th European Conference*,

- Glasgow, UK, August 23–28, 2020, *Proceedings, Part XIV 16*, pages 474–489. Springer, 2020. 7
- [47] Wei Mao, Miaomiao Liu, Mathieu Salzmann, and Hongdong Li. Learning trajectory dependencies for human motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9489–9497, 2019. 1, 2, 6, 7
- [48] Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2891–2900, 2017. 1, 2, 6, 7
- [49] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*, pages 604–621. Springer, 2022. 2, 3
- [50] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018. 2, 3
- [51] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. 3
- [52] Theodoros Sofianos, Alessio Sampieri, Luca Franco, and Fabio Galasso. Space-time-separable graph convolutional network for pose forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11209–11218, 2021. 2, 7
- [53] Hao Tan, Jie Lei, Thomas Wolf, and Mohit Bansal. Vim-pac: Video pre-training via masked token prediction and contrastive learning. *arXiv preprint arXiv:2106.11250*, 2021. 2, 3
- [54] Graham W Taylor and Geoffrey E Hinton. Factored conditional restricted boltzmann machines for modeling motion style. In *Proceedings of the 26th annual international conference on machine learning*, pages 1025–1032, 2009. 2
- [55] Shubham Toshniwal, Hao Tang, Liang Lu, and Karen Livescu. Multitask learning with low-level auxiliary tasks for encoder-decoder based speech recognition. *arXiv preprint arXiv:1704.01631*, 2017. 3
- [56] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103, 2008. 2
- [57] Timo Von Marcard, Roberto Henschel, Michael J Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *Proceedings of the European conference on computer vision (ECCV)*, pages 601–617, 2018. 2, 6
- [58] Jacob Walker, Kenneth Marino, Abhinav Gupta, and Martial Hebert. The pose knows: Video forecasting by generating pose futures. In *Proceedings of the IEEE international conference on computer vision*, pages 3332–3341, 2017. 2
- [59] Guangming Wang, Xinrui Wu, Zhe Liu, and Hesheng Wang. Pwclo-net: Deep lidar odometry in 3d point clouds using hierarchical embedding mask optimization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15910–15919, 2021. 2, 3
- [60] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14668–14678, 2022. 3
- [61] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simsim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9653–9663, 2022. 3
- [62] Chenxin Xu, Siheng Chen, Maosen Li, and Ya Zhang. Invariant teacher and equivariant student for unsupervised 3d human pose estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 3013–3021, 2021. 1
- [63] Chenxin Xu, Maosen Li, Zhenyang Ni, Ya Zhang, and Siheng Chen. Groupnet: Multiscale hypergraph neural networks for trajectory prediction with relational reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6498–6507, 2022. 2
- [64] Chenxin Xu, Robby T Tan, Yuhong Tan, Siheng Chen, Yu Guang Wang, Xinchao Wang, and Yanfeng Wang. Eqmotion: Equivariant multi-agent motion prediction with invariant interaction reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1410–1420, 2023. 2
- [65] Lian Xu, Wanli Ouyang, Mohammed Bennamoun, Farid Boussaid, Ferdous Sohel, and Dan Xu. Leveraging auxiliary tasks with affinity learning for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6984–6993, 2021. 3
- [66] Yiding Yang, Jiayan Qiu, Mingli Song, Dacheng Tao, and Xinchao Wang. Distilling knowledge from graph convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [67] Tao Yu, Zhihe Lu, Xin Jin, Zhibo Chen, and Xinchao Wang. Task residual for tuning vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 2
- [68] Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen Lu. Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19313–19322, 2022. 3
- [69] Farhad Ghazvinian Zanjani, David Anssari Moin, Frank Claessen, Teo Cherici, Sarah Parinussa, Arash Pourtaherian, Svitlana Zinger, and Peter HN de With. Mask-mcnet: Instance segmentation in 3d point cloud of intra-oral scans. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part V 22*, pages 128–136. Springer, 2019. 3
- [70] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, Octo-*

ber 11-14, 2016, *Proceedings, Part III 14*, pages 649–666. Springer, 2016. 3

- [71] Richard Zhang, Phillip Isola, and Alexei A Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1058–1067, 2017. 3
- [72] Yu Zhang and Qiang Yang. A survey on multi-task learning. *IEEE Transactions on Knowledge and Data Engineering*, 34(12):5586–5609, 2021. 3
- [73] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021. 3