

Joint-Relation Transformer for Multi-Person Motion Prediction

Qingyao Xu^{1*}, Weibo Mao^{1*}, Jingze Gong¹, Chenxin Xu¹,
Siheng Chen^{1,2†}, Weidi Xie^{1,2}, Ya Zhang^{1,2}, Yanfeng Wang^{1,2†}

¹Shanghai Jiao Tong University ²Shanghai AI Laboratory

{xuqingyao, kirino.mao, xcxwakaka, sihengc, weidi, ya-zhang, wangyanfeng}@sjtu.edu.cn
autubot@163.com

Abstract

Multi-person motion prediction is a challenging problem due to the dependency of motion on both individual past movements and interactions with other people. Transformer-based methods have shown promising results on this task, but they miss the explicit relation representation between joints, such as skeleton structure and pairwise distance, which is crucial for accurate interaction modeling. In this paper, we propose the Joint-Relation Transformer, which utilizes relation information to enhance interaction modeling and improve future motion prediction. Our relation information contains the relative distance and the intra-/inter-person physical constraints. To fuse relation and joint information, we design a novel joint-relation fusion layer with relation-aware attention to update both features. Additionally, we supervise the relation information by forecasting future distance. Experiments show that our method achieves a 13.4% improvement of 900ms VIM on 3DPW-SoMoF/RC and 17.8%/12.0% improvement of 3s MPJPE on CMU-Mpcap/MuPoTS-3D dataset. *Code is available at <https://github.com/MediaBrain-SJTU/JRTransformer>.*

1. Introduction

Multi-person motion prediction aims to predict the future positions of skeleton joints for multiple individuals based on their historical movements. Compared to traditional single-person motion prediction [11, 7, 4, 24, 48, 47, 27], multi-person motion prediction is more practical as people are mostly associated with a group and interacting with each other. It is also more challenging because sophisticated interactions across different individuals need to be considered. The related methods are playing significant roles in a wide range of practical applications, including autonomous driving [15, 38, 10, 51], surveillance systems [13, 42, 52] and healthcare monitoring [43]. They also pave a path to better human-robot interaction [6, 12].

Previous works on multi-person motion prediction gen-

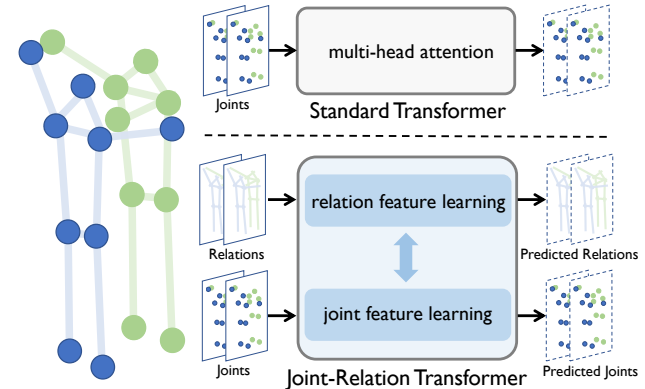


Figure 1. System comparison between the standard Transformer and the proposed Joint-Relation Transformer.

erally adopt two types of architectures, including graph neural networks (GNNs) and Transformer. The former one, such as TRiPOD [3], models the multi-person interaction via a graph structure, however, it suffers from the inherent oversmoothing problem in GNNs, thus can only afford shallow models with limited learning capacities; while the latter one, such as MRT [46] and SoMoFormer [40], treat temporal sequence or skeleton joints as a sequence input, and learn to establish the relations between them via self-attention mechanism. Compared with GNN-based methods, Transformer has shown strong learning ability, thus becoming a *default* backbone for multi-person motion prediction.

While encouraging results are shown in the Transformer-based methods, those previous works only implicitly learn the inter-joint relation through the attention mechanism and lack the explicit awareness of skeleton structure. To address this issue, we propose Joint-Relation Transformer, a two-stream Transformer architecture for multi-person motion prediction. Instead of only updating the features for body joints, Joint-Relation Transformer uses two streams to achieve feature learning for both joints and relations. In specific, one stream encodes the sequence of skeleton joints, containing the historical movement information of the joints in world coordinates; and the other one explicitly takes the joints-to-joints relation as input, including the relative distance between two joints and physical constraints, such as inter-body constraints and intra-body skeleton connections.

To effectively fuse and update features of joints and re-

*Equal contribution. †Corresponding author.

lation branches, we design a novel relation-aware attention to update joints' features with the incorporation of relation features. During the update procedure, attention scores between skeleton joints are calculated from two sources: the similarity score between two joints' features and the additional relation score based on the relation feature between these two joints. This approach enhances the model's ability to distinguish between similar joint features belonging to distinct persons, leading to more accurate joint updates. Moreover, this design allows for increased granularity in attention allocation, enabling the concentration on the most pertinent joint features for each person, consequently enhancing prediction performance.

To train our proposed Joint-Relation Transformer, in addition to inferring joint positions, we also supervise the model by predicting future inter-joint distances, which contains the future relationship between two joints. This relative distance supervision is translation and rotation invariant of the input sequence, adhering to invariance properties of multi-person interactions. As to evaluate the effectiveness of our method, we conduct experiments on four multi-person motion prediction datasets: 3DPW-SoMoF [3], 3DPW-SoMoF/RC, CMU-Mocap [1], and MuPoTS-3D [32]. The quantitative results show we outperform the previous methods and achieve state-of-the-art performance on most datasets. The qualitative results verify the reasonable attention allocation and vivid predictions.

To summarise, in this paper, we make the following contributions: (i) We propose the Joint-Relation Transformer for multi-person motion prediction. We innovatively introduce the relation information, which explicitly builds the relationship between joints of the inter-/intra-body; (ii) We design a relation-aware attention module to update the joint information with the corporation of explicit relation information, increasing granularity in attention allocation and enhancing prediction performance; (iii) We further supervise the relation information between two joints with the future relative distance to better capture the interaction information hidden in the distance variation; (iv) We perform our experiments on several common datasets and our proposed method outperforms most state-of-the-art methods. We also conduct thorough ablation study to show the importance of the proposed relation information and relation-aware attention in the task of multi-person motion estimation.

2. Related Work

Single person motion prediction. To forecast motions, traditional methods typically employ Hidden Markov Models [23], Gaussian-process [45] and other methods with hand-crafted features [49], etc. With the development of deep learning, many RNN-based networks [11, 20, 14, 36, 16, 7, 30, 26] have been developed and achieve great success. For instance, ERD [11] uses an RNN architecture with

nonlinear encoder and decoder networks to predict human motion, while TP-RNN [7] captures the latent hierarchical structure of human poses at different time scales via RNN for a better prediction. Despite their success, RNN-based networks suffer from error accumulation problems inherent to their architecture. Some works [4, 24, 48, 47, 27, 25, 28, 53] utilize feed-forward network to alleviate this problem. For example, LTD [48] adopts DCT to encode temporal information and adopts GCN to learn graph connectivity automatically. HRI [47] introduces an attention-based feed-forward network to extract motion attention between sub-sequences. DMGNN [27] introduces a multiscale graph to comprehensively model the internal relations of human motion. EqMotion[53] proposes an efficient equivariant motion prediction model to maintain motion equivariance and interaction invariance. However, these methods primarily concern situations involving only one person, lacking the consideration of interaction in real-world scenarios.

Multi-person human motion prediction. Recent studies have emphasized multi-person forecasting problems. Due to the existence of interactive behavior, people's motions are likely to be affected by others. To handle this problem, JSC [2] models global motion and local body joint movements separately with a shared GRU encoder to incorporate scene and social contexts. TRiPOD [3] adopts GAT which considers the person and objects as graph nodes to capture the interaction between them. Nowadays many works [46, 17, 40] start to explore the application of Transformers in this task because of its powerful learning ability. MRT [46] pays attention to the change and correlation of human pose in time and uses a global encoder to capture the overall interaction between humans of each frame. SoMoFormer [40] carefully models the interaction between all joints, but simply uses the Transformer's attention mechanism and lacks a detailed design for relation modeling. DuMMF[54] focuses on stochastic predicting and proposes a dual-level framework to tackle local individual motion and global social interactions. We notice that although achieving a superior result, Transformer-based methods only use the attention mechanism to implicitly infer the joints' interaction, missing the possible explicit relation information between joints. To overcome this problem, we introduce the relation information into the Transformer architecture and propose Joint-Relation Transformer that can tackle and utilize both the joint information and relation information.

Graph transformer. Transformer [39] has made great success in many fields and shown great potential in modeling graph data [18, 34, 5]. Many methods have been proposed to incorporate graph structure into Transformer, which can be divided into three categories. The first one is to adopt GNN as an auxiliary module where the attention block and the GNN block are always separate [37, 56, 33, 29, 50]. The second one uses the improved positional embedding derived

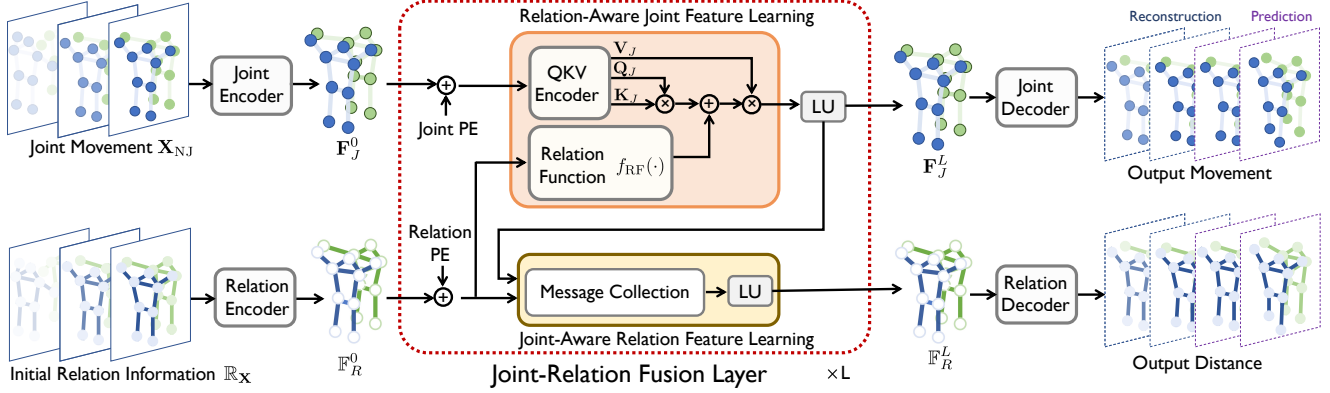


Figure 2. Architecture of proposed Joint-Relation Transformer. The network contains three stages: i) encoder modules to extract features, ii) a fusion module to fuse the joint and relation features, and iii) decoder modules to output final results. LU: local update module.

from graph structure to have a better perception of graph structure information [9, 19, 22, 55]. The third one focus on the attention matrix design and uses graph structure to get a more representative matrix [21, 57, 8]. Our method is more related to the third kind where we design a special attention mechanism to calculate relation-aware attention score.

3. Problem Formulation

Multi-person motion prediction aims to predict the future positions of joints for multiple individuals based on their historical movements. Mathematically, given a scene with N persons, each has J skeleton joints, let $\mathbf{X}_{NJ} = [X_1, X_2, \dots, X_{NJ}] \in \mathbb{R}^{NJ \times (T_h \times 3)}$ be the observed sequence in the scene over T_h history timestamps, where $X_{nj} \in \mathbb{R}^{T_h \times 3}$ refers to the observed joint sequence of the j -th joint on n -th person in 3D world coordinate. Similarly, we define the future motions as $\mathbf{Y}_{NJ} = [Y_1, Y_2, \dots, Y_{NJ}] \in \mathbb{R}^{NJ \times (T_f \times 3)}$, and $Y_{nj} \in \mathbb{R}^{T_f \times 3}$ represents the future joint sequence for the j -th joint of n -th person. The goal is to train a computational model $\mathcal{F}(\cdot)$ that infers the future motions of multiple individuals by the observed motion $\hat{\mathbf{Y}}_{NJ} = \mathcal{F}(\mathbf{X}_{NJ})$ to approximate the ground-truth future motion \mathbf{Y}_{NJ} . Unless otherwise specified, we will use \mathbf{X}/\mathbf{Y} for ground-truth matrices, $\mathbb{D}_{\mathbf{X}}/\mathbb{D}_{\mathbf{Y}}$ for tensors in past and future, and $\hat{\mathbf{X}}/\hat{\mathbf{Y}}$ for generated matrices.

This task is challenging for several reasons: i) the movement of a person is bounded by physical constraints, *e.g.* two joints with a bone connection should maintain the constant relative distance; and ii) the behavior of the joints is influenced by both the internal structure of the person’s joints, as well as inter-personal joints. To alleviate this dilemma, we introduce the relation information and fuse the information with the proposed Joint-Relation Transformer.

4. Architecture

In this section, we present a two-stream architecture for multi-person motion prediction, that simultaneously infers

the position of joints and their relations/distance, as shown in Fig. 2. Mathematically, let \mathbf{X}_{NJ} be the observed joint information, the **inference procedure** can be formulated as:

$$\hat{\mathbf{X}}_{NJ}, \hat{\mathbf{Y}}_{NJ}, \hat{\mathbb{D}}_{\mathbf{X}}, \hat{\mathbb{D}}_{\mathbf{Y}} = \Phi_{\text{decode}}(\cdot) \circ \Phi_{\text{fuse}}(\cdot) \circ \Phi_{\text{encode}}(\mathbf{X}_{NJ}),$$

in particular, the historical sequences \mathbf{X}_{NJ} are firstly processed with an encoder module $\Phi_{\text{encode}}(\cdot)$ (Sec. 4.1) to obtain the initial joints and relation representation. Then the output representation is passed into a fusion module $\Phi_{\text{fuse}}(\cdot)$, to fuse and update the joints and relation representation (Sec. 4.2). Lastly, the fused features are processed by the decoder module $\Phi_{\text{decode}}(\cdot)$, outputting the joints’ positions $\hat{\mathbf{X}}_{NJ}, \hat{\mathbf{Y}}_{NJ}$ and distance between joints $\hat{\mathbb{D}}_{\mathbf{X}}, \hat{\mathbb{D}}_{\mathbf{Y}}$ (Sec. 4.3). At training time, the model is optimized to infer motions for both historical and future timestamps, potentially alleviating the catastrophic forgetting issue; while at inference time, we only take the predictions for future timestamps (as detailed in Sec. 4.4).

4.1. Encoder Module

In the literature, Transformer-based motion prediction models have primarily considered using the 3D world coordinates of joints as input, like MRT [46] and SoMoFormer [40], that requires the model to implicitly learn the complex dependencies between joints. Here, we propose to enrich the motion representation by augmenting it with temporal differentiation and **explicit joint relations**. In particular, we independently process the historical motion sequences by a joint encoder which considers temporal differentiation, and a relation encoder, which considers explicit joint relations, as detailed in the following.

Joint Encoder. Instead of only using the joint position information in world coordinates as input, here we augment the joint information by taking the temporal differentiation between current and one step before: $\Delta X_{nj}^t = X_{nj}^t - X_{nj}^{t-1}$, *i.e.*, the velocity information and concatenating them.

Till here, each joint is effectively encoded by a vector of $T_h \times 6$ dimensions. We then employ a 2-layer MLP encoder

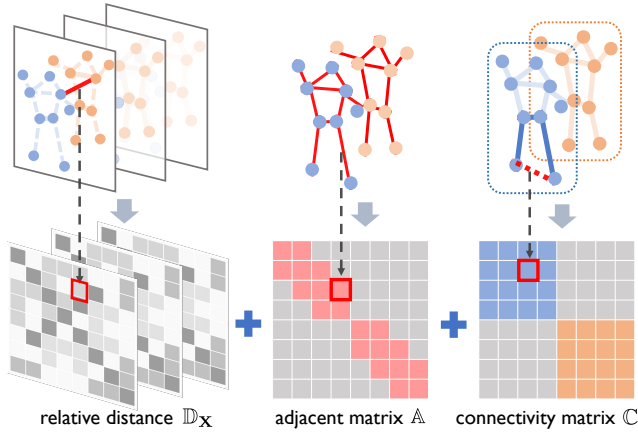


Figure 3. Three key components in the relation information including relative distance matrix sequence, bone adjacent matrix, and connectivity matrix. The upper is a schematic diagram of several relations, and the lower is the corresponding relation information.

to project each joint feature into a higher dimension. We denote the output of the joint encoder as $\mathbf{F}_J^0 \in \mathbb{R}^{NJ \times D}$ with the feature dimension D .

Relation Encoder. In addition to individually encoding joints, we also propose to encode the relation between joints, from three perspectives: (i) to capture the relative distance between joints, we compute the distance information $\mathbb{D}_X \in \mathbb{R}^{NJ \times NJ \times T_h}$ that consists of T_h distance-aware matrices for each frame of the input sequence, where $(\mathbb{D}_X)_{ij}^t = e^{-\|X_i^t - X_j^t\|_2}$ is calculated as the negative exponent of the distance between the i -th and the j -th joint at the t -th frame, X_i and X_j represents the corresponding joint's position; (ii) to reflect the skeleton structure where joints connected by bone show stronger associations, we compute the adjacent matrix $\mathbb{A} \in \mathbb{R}^{NJ \times NJ \times 1}$ where $\mathbb{A}_{ij} = 1$ represents there is a bone connection between joint i and j , and 0 otherwise; (iii) to model the connectivity between joints since those belonging to the same individual tend to exhibit similar overall movement patterns, we introduce a connectivity matrix $\mathbb{C} \in \mathbb{R}^{NJ \times NJ \times 1}$, where $\mathbb{C}_{ij} = 1$ if there is a path defined on bones between joint i and j . We refer to the combination of the last two items \mathbb{A} and \mathbb{C} as physical constraints between joints.

Therefore, the initial relation information can be defined as a tensor, *i.e.*, $\mathbb{R}_X = [\mathbb{D}_X, \mathbb{A}, \mathbb{C}] \in \mathbb{R}^{NJ \times NJ \times (T_h + 2)}$, incorporating both historical distances information and physical constraints between each pair of joints, as shown in Fig. 3. Till here, the relation tensor is encoded by a 1×1 convolution encoder, that fuses the $T_h + 2$ channels and projects it to the same dimension as \mathbf{F}_J . The output relation feature is denoted as $\mathbb{F}_R^0 \in \mathbb{R}^{NJ \times NJ \times D}$.

4.2. Fusion Module

After separately encoding the joints and relations, their output features are further passed into a fusion module with




Figure 4. Example of joint embedding and relation embedding.

positional information injected, and followed by several novel Joint-Relation Fusion Layers, as detailed below.

4.2.1 Positional Embedding

We add the positional embedding to the output features from the joint and relation encoder respectively. Specifically, for joint representation, we use $(N + J)$ learnable embeddings to indicate the person and joint identities respectively; while for relation representation, the position embeddings are constructed by adding the positional embeddings of two related joints, as shown in Fig. 4. As a consequence, the resulting features from two-stream encoders become position-aware now.

4.2.2 Joint-Relation Fusion Layer

In this section, we describe two novel architectural designs that encourage the communication between joint and relation branches, namely, **relation-aware joint feature learning** and **joint-aware relation feature learning**. Specifically, the former enables to update the joints' features with the attention mechanism by querying information from the relation branch, and the latter further updates the relation features with message collection and local updates.

Relation-Aware Joint Feature Learning. To explicitly incorporate the relation information while updating joints' representation, we design a novel relation-aware attention for joint feature learning, formulated as:

$$\begin{aligned}
 \mathbf{F}_J^{l+1} &= f_{\text{RA}}^l(\mathbf{F}_J^l, \mathbb{F}_R^l), \quad \text{omit } l \text{ below for brevity,} \\
 &= \text{MHA}(f_{\text{QKV-J}}(\mathbf{F}_J), f_{\text{RF}}(\mathbb{F}_R)), \\
 &= \text{MH}(\text{softmax}(\frac{\mathbf{Q}_J(\mathbf{K}_J)^\top + f_{\text{RF}}(\mathbb{F}_R)}{\sqrt{F_K}})\mathbf{V}_J).
 \end{aligned} \tag{1}$$

Specifically, $\mathbf{F}_J^l \in \mathbb{R}^{NJ \times D}$ and $\mathbb{F}_R^l \in \mathbb{R}^{NJ \times NJ \times D}$ are the input joint and relation feature with added the positional embedding at the l -th layer. The relation function $f_{\text{RF}}(\mathbb{F}_R) = \mathbb{F}_R \mathbf{W}_l + \sum (\mathbb{F}_R \mathbf{W}_q^1 \odot \mathbb{F}_R \mathbf{W}_q^2)$ is composed with linear projection term and quadratic projection term where $\mathbf{W}_l \in \mathbb{R}^{D \times 1}$, $\mathbf{W}_q^1, \mathbf{W}_q^2 \in \mathbb{R}^{D \times D'}$ are learnable parameters and we perform the sum operation among the last dimension. $\text{MHA}(\cdot)$ refers to a variant of multi-head attention. $\text{MH}(\cdot)$ denotes the operation of computing the

weighted average of values at each head, and F_K is the feature size of the key \mathbf{K}_J .

In Eq. (1), we first generate the query/key/value for joint through the corresponding joint QKV encoder $f_{\text{QKV-J}}(\cdot)$ and calculate the relation score through the relation function $f_{\text{RF}}(\cdot)$ with the multi-head of D_H . We then fuse the relation and joint feature in the attention calculation, which adds the relation score to the initial attention score from the joint’s query/key. We finally output the updated joint feature through the $\text{softmax}(\cdot)$ and weighted sum operations. Note that here queries, keys, and values share the same feature size D_K . We adopt the quadratic term in the relation function for better representation ability.

We further perform local update that first employs a residual connection to the joint feature and then applies a feed-forward layer with a residual connection and a layer normalization to get the updated feature.

Joint-Aware Relation Feature Learning. With the fused and updated features through relation-aware attention, we elaborate the procedure for updating the relation feature by collecting all messages and performing the local update. Specifically, given $\mathbb{F}_R^l \in \mathbb{R}^{NJ \times NJ \times D}$ and $\mathbf{F}_J^{l+1} \in \mathbb{R}^{NJ \times D}$, denoting the relation tensor and updated features for joints, respectively. The operating procedure of the joint-aware relation feature learning can be formulated as:

$$\mathbb{F}_J^{l+1} = \text{Broadcast}(\mathbf{F}_J^{l+1}) \in \mathbb{R}^{NJ \times NJ \times D}, \quad (2a)$$

$$\mathbb{M}^{l+1} = [\mathbb{F}_J^{l+1}, (\mathbb{F}_J^{l+1})^\top, \mathbb{F}_R^l, (\mathbb{F}_R^l)^\top] \in \mathbb{R}^{NJ \times NJ \times D_M}, \quad (2b)$$

$$\mathbb{F}_R^{l+1} = \mathbb{F}_R^l + f_{\text{LU1}}(\text{Norm}(\mathbb{M}^{l+1})) \in \mathbb{R}^{NJ \times NJ \times D}, \quad (2c)$$

$$\mathbb{F}_R^{l+1} = \mathbb{F}_R^{l+1} + f_{\text{LU2}}(\text{Norm}(\mathbb{F}_R^{l+1})) \in \mathbb{R}^{NJ \times NJ \times D}, \quad (2d)$$

where $D_M = 4 \times D$ is the hidden dimension of the collected message \mathbb{M}^{l+1} . In detail, Step (2a) broadcasts the joint feature to relations; Step (2b) collects the message \mathbb{M}^{l+1} for the current relation, which contains the joint features within the linked joints and bi-directional to-be update relation feature; Step (2c) performs the first round of local update with the normalization operation $\text{Norm}(\cdot)$ and local update layer $f_{\text{LU1}}(\cdot)$; and Step (2d) performs the second round.

Note that, in equations, (i) both update layers $f_{\text{LU1}}(\cdot)$ and $f_{\text{LU2}}(\cdot)$ are implemented through MLP, (ii) the transpose operation, e.g. $(\mathbb{F}_R^l)^\top$, is performed on the first two dimensions to collect the corresponding bi-directional features; and (iii) we adopt the $\text{LayerNorm}(\cdot)$ as the normalization operation.

4.3. Decoder Module

Through the fusion module, the joint and relation features are well-fused in high dimensions. We now elaborate on the decoder module which projects the feature back to joint motion and distance between joints.

Joint Decoder. In particular, we adopt the joint decoder $\mathcal{D}_J(\cdot)$ to decode the fused joint feature $\mathbf{F}_J^L \in \mathbb{R}^{NJ \times D}$ via a 3-layer MLP. It outputs the reconstructed joint movement $\hat{\mathbf{X}}_{\text{NJ}} \in \mathbb{R}^{NJ \times (T_h \times 3)}$ over T_h frames and predicted movement $\hat{\mathbf{Y}}_{\text{NJ}} \in \mathbb{R}^{NJ \times (T_f \times 3)}$ over T_f future frames at once.

Relation Decoder. In order to better model the interaction from the perspective of distance changes between joints, we employ a 1×1 convolution decoder $\mathcal{D}_R(\cdot)$ to decode the fused relation feature $\mathbb{F}_R^L \in \mathbb{R}^{NJ \times NJ \times D}$. In specific, the decoder fuses the D feature channels and projects it to $T_h + T_f$ dimension where the first T_h dimension is the reconstructed distance $\hat{\mathbb{D}}_{\mathbf{X}} \in \mathbb{R}^{NJ \times NJ \times T_h}$ and the last T_f represents the predicted distance $\hat{\mathbb{D}}_{\mathbf{Y}} \in \mathbb{R}^{NJ \times NJ \times T_f}$.

4.4. Training Objective

To train the proposed Joint-Relation Model, we here adopt three types of supervision terms, including i) joint supervision on reconstructed and predicted joint motion, ii) relation supervision on reconstructed and predicted distance sequence, and iii) Transformer deep supervision on the input of each fusion layer, as detailed below.

Joint Supervision. Let $\hat{\mathbf{X}}_{\text{NJ}}$ and $\hat{\mathbf{Y}}_{\text{NJ}}$ be the reconstructed and predicted joint information and $\mathbf{X}_{\text{NJ}}/\mathbf{Y}_{\text{NJ}}$ be the corresponding ground-truth, we supervise them through

$$\mathcal{L}_J(\hat{\mathbf{X}}_{\text{NJ}}, \hat{\mathbf{Y}}_{\text{NJ}}) = \|\mathbf{X}_{\text{NJ}} - \hat{\mathbf{X}}_{\text{NJ}}\|_2 + \lambda_J \|\mathbf{Y}_{\text{NJ}} - \hat{\mathbf{Y}}_{\text{NJ}}\|_2,$$

where $\lambda_J \in \mathbb{R}$ is a weight hyperparameter to balance two terms and $\|\cdot\|_2$ takes L2-norm on the 3D coordinates and average other dimensions.

Relation Supervision. Let $\hat{\mathbb{D}}_{\mathbf{X}}$ and $\hat{\mathbb{D}}_{\mathbf{Y}}$ be the reconstructed and predicted distance and $\mathbb{D}_{\mathbf{X}}/\mathbb{D}_{\mathbf{Y}}$ be the corresponding ground-truth, we formulate the supervision as

$$\mathcal{L}_R(\hat{\mathbb{D}}_{\mathbf{X}}, \hat{\mathbb{D}}_{\mathbf{Y}}) = \|\mathbb{D}_{\mathbf{X}} - \hat{\mathbb{D}}_{\mathbf{X}}\|_1 + \lambda_R \|\mathbb{D}_{\mathbf{Y}} - \hat{\mathbb{D}}_{\mathbf{Y}}\|_1,$$

where $\lambda_R \in \mathbb{R}$ is the weight hyperparameter and $\|\cdot\|_1$ represents the average L1-norm between the ground truth distance and the predicted value.

Deep Supervision. To avoid overfitting caused by deep Transformer fusion layers, we use the same joint/relation decoder on the input of each fusion layer and calculate the corresponding losses:

$$\mathcal{L}_{\text{DS}} = \sum_{l=0}^{L-1} \left(\mathcal{L}_J(\mathcal{D}_J(\mathbf{F}_J^l)) + \mathcal{L}_R(\mathcal{D}_R(\mathbb{F}_R^l)) \right). \quad (3)$$

Total Loss. $\mathcal{L} = \mathcal{L}_J(\hat{\mathbf{X}}_{\text{NJ}}, \hat{\mathbf{Y}}_{\text{NJ}}) + \mathcal{L}_R(\hat{\mathbb{D}}_{\mathbf{X}}, \hat{\mathbb{D}}_{\mathbf{Y}}) + \mathcal{L}_{\text{DS}}$.

5. Experimental Setup

5.1. Datasets

We evaluate our method on three multi-person motion datasets, including 3DPW [41], CMU-Mocap [1], and

Table 1. Experimental results in VIM on the 3DPW-SoMoF (left) and 3DPW-SoMoF/RC (right) test sets. The best results are highlighted in bold. Our method outperforms most previous state-of-the-art methods on 3DPW-SoMoF and achieves the best performance on the 3DPW-SoMoF/RC test sets.

Methods	3DPW-SoMoF						3DPW-SoMoF/RC					
	AVG	100	240	500	640	900	AVG	100	240	500	640	900
Zero Velocity	86.7	29.4	53.6	94.5	112.7	143.1	67.2	23.9	42.9	73.1	86.2	109.7
LTD [48] ²⁰¹⁹	76.7	22.0	41.1	81.0	100.2	139.7	62.2	20.2	37.2	68.6	81.3	103.5
TRiPOD [3] ²⁰²¹	84.2	31.0	50.8	84.7	104.1	150.4	-	-	-	-	-	-
DViTA [35] ²⁰²¹	65.7	19.5	36.9	68.3	85.5	118.2	57.4	15.0	32.3	62.7	76.3	100.6
MRT [46] ²⁰²¹	59.2	21.8	39.1	65.1	75.9	94.1	52.3	20.8	36.4	58.2	66.6	79.4
FutureMotion [44] ²⁰²¹	49.4	9.5	22.9	50.9	66.2	97.4	-	-	-	-	-	-
SoMoFormer [40] ²⁰²²	46.3	9.1	21.3	47.5	61.6	91.9	43.8	11.0	24.8	49.1	59.2	74.8
Ours	47.2	9.5	22.1	48.7	62.8	92.8	39.5	9.5	21.7	44.1	53.4	68.8

MuPoTS-3D [32].

3DPW. 3D Poses in the Wild Dataset (3DPW) is a large-scale 3D motion dataset collected by moving mobile phone cameras with pose estimation and optimization [41]. In this paper, we use the SoMoF benchmark [2, 3] splits (3DPW-SoMoF), *i.e.*, the sequences that contain two persons, and predict future 900ms (14 frames) motion using the historical 1030ms (16 frames) motion.

3DPW-SoMoF/RC. We find the camera movement in 3DPW dataset causes a serious unnatural drift on persons, affecting the modeling of multi-person interaction. Here, we subtract the estimated camera velocity for better interaction modeling and generate a new dataset, termed as 3DPW-SoMoF/RC. See more details in supplementary materials.

CMU-Mocap. The Carnegie Mellon University Motion Capture Database (CMU-Mocap) contains a large number of single-person scenes and only limited set of scenes with two persons [1]. We use the training set and test set given in the paper [46], where each scene contains 3 persons, and are obtained by sampling from single-person scenes with multi-person scenes and mixing them together. We aim to predict future 3000ms (45 frames) motion using the historical 1000ms (15 frames) motion.

MuPoTS-3D. Multiperson Pose Test Set in 3D (MuPoTS-3D) consists of over 8000 frames collected from 20 sequences with 8 subjects [32]. Following previous works [46, 40], we evaluate our model’s performance with the same segment length as CMU-Mocap on the test set.

5.2. Metrics

VIM. Visibility-Ignored Metric (VIM) is proposed in SoMoF benchmark [2, 3] to measure the displacement on the joint vector with the dimension of $J \times 3$. To be specific, the VIM for t -th frame is calculated as

$$\text{VIM}@t = \frac{1}{N} \sum_{n=1}^N \sqrt{\sum_{j=1}^J (Y_{nj}^t - \hat{Y}_{nj}^t)^2}.$$

MPJPE. Mean Per Joint Position Error (MPJPE) is another commonly used metric in the fields of pose estimation and motion prediction, which calculates the average Euclidean distance between the predicted value and the ground truth of all joints, that is,

$$\text{MPJPE} = \frac{1}{T_f} \frac{1}{N} \frac{1}{J} \sum_{t=1}^{T_f} \sum_{n=1}^N \sum_{j=1}^J \|Y_{nj}^t - \hat{Y}_{nj}^t\|_2.$$

We use this metric on CMU-Mocap and MuPoTS-3D.

5.3. Implementation Details

For 3DPW-SoMoF and 3DPW-SoMoF/RC datasets, we first pre-train the model on the AMASS [31] dataset following previous works [44, 40], which provides massive motion sequences. We use the *CMU* subset as the training set and the *BioMotionLab_NTroje* for test. We randomly sample single-person sequences from the dataset and mix them to get a synthetic set. While finetuning the model, we extend the 3DPW-SoMoF and 3DPW-SoMoF/RC datasets by sampling with overlap every two frames from the original 3DPW dataset. We perform three kinds of data augmentation including i) random rotation: the entire scene is rotated by a random angle within $[0, 2\pi]$ along the vertical axis; ii) person permutation: the person order in a scene is randomly permuted; and iii) sequence reverse: the entire sequence is temporally reversed and the last T_h frames of the original sequence are taken as input. Here all the input sequences are normalized by subtracting the mean joint position of the first person in the first frame.

Training Details. Our model has $L = 4$ joint-relation fusion layers with $D_H = 8$ attention heads, the feature dimension D is set to 128. On 3DPW-SoMoF and 3DPW-SoMoF/RC datasets, we first pre-train for 100 epochs with an initial learning rate of 1×10^{-3} and decay by 0.8 every 10 epochs. When fine-tuning, our learning rate is set to 1×10^{-4} with a 0.8 decay every 10 epochs. The batch

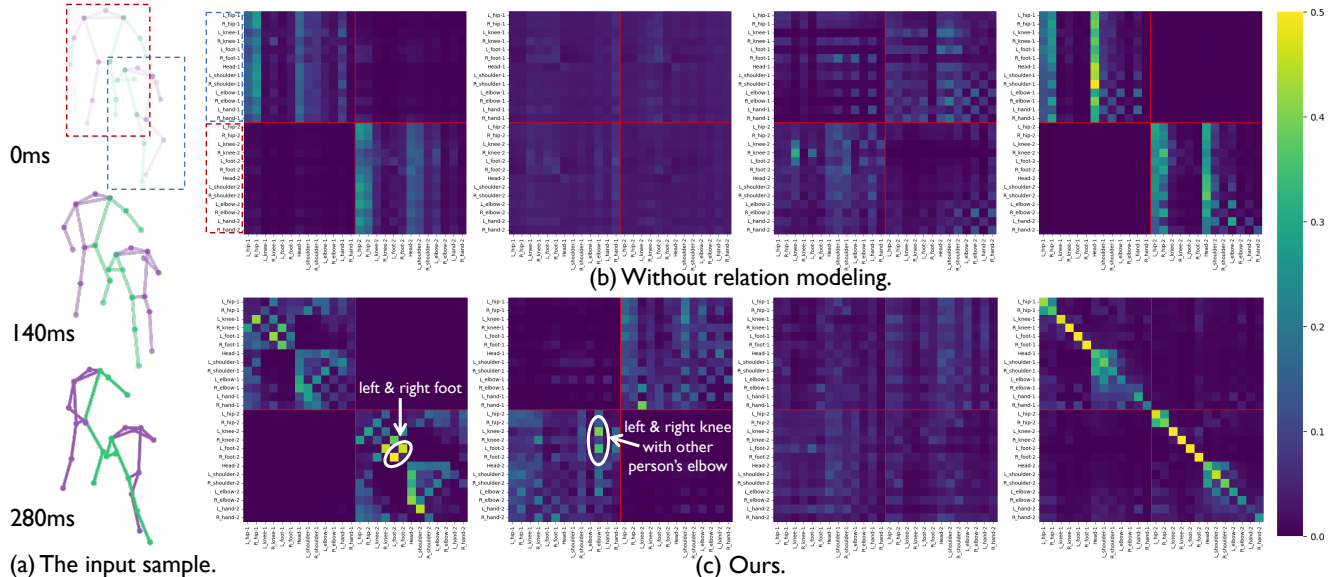


Figure 5. Attention visualization in the first attention layer with different heads. (a) shows the input sequence; (b) and (c) shows the plain/relation-aware attention matrices in the first layer. With the help of relation information, proposed relation-aware attention generates more reasonable attention allocations: (from left to right) intra-body attention, inter-body attention, global attention, and self-attention.

Table 2. Experimental results in MPJPE on CMU-Mocap (left) and MuPoTS-3D (right) test sets. The best results are highlighted in bold. Our method outperforms most state-of-the-art methods.

Methods	CMU-Mocap			MuPoTS-3D		
	1s	2s	3s	1s	2s	3s
LTD [48] ²⁰¹⁹	13.7	21.9	32.6	11.9	18.1	23.4
HRI [47] ²⁰²⁰	14.9	26.0	30.7	9.4	16.8	22.9
MRT [46] ²⁰²¹	9.6	15.7	21.8	8.9	15.9	22.2
SoMoFormer [40] ²⁰²²	10.7	17.2	22.5	10.1	17.7	24.2
Ours	8.3	13.9	18.5	8.9	15.5	21.3

size is set to 128 for both pre-train and finetune. We use $\lambda_J = \lambda_R = 10$. The whole network is implemented using Pytorch with AdamW optimizer.

5.4. Quantitative results

Results on 3DPW-SoMoF and 3DPW-SoMoF/RC. We report the experimental results on the 3DPW-SoMoF and 3DPW-SoMoF/RC datasets in Tab. 1. For a fair comparison, we use the same VIM criterion at multiple future frames. On standard SoMoF benchmark (left section of Tab. 1), our method still outperforms most all previous methods on the 3DPW-SoMoF dataset, despite of serious drift problems; on SoMoF/RC benchmark with drifting removed (right section of Tab. 1), our method reduces the VIM at AVG/900ms from 43.8/74.8 to 39.5/68.8, compared to the current state-of-the-art method SoMoFormer [40], achieving 9.8%/13.4% improvement, reflecting the effectiveness of the proposed method.

Results on CMU-Mocap and MuPots-3D. We also compare the result on CMU-Mocap and MuPots-3D dataset between our method and MRT [46], SoMoFormer [40] as well as two recent single-person motion prediction methods including HRI [47] and LTD [48]. Both methods are trained on the synthesized dataset given by MRT and tested on the corresponding CMU-Mocap test set and MuPots-3D test set. We report the MPJPE result on predicting 1, 2, and 3s motion as the MRT does in Tab. 2. We observe a significant performance improvement on CMU-Mocap and our method also achieves most of state-of-the-art results on MuPots-3D. This demonstrates the strong ability of our model to predict multi-person motion and strong generalization.

5.5. Qualitative Results

Visualization of attention weights. To verify the effectiveness of the proposed relation-aware attention, we visualize the learned attention matrices in the first layer by the relation-aware attention and the plain attention over joints; see Fig. 5. We see that our model has learned to explicitly generate diverse attention matrices including inter-person attention, intra-person attention, global attention, and self-attention for different heads.

Visualization of prediction result. We provide the qualitative comparison between our method with other recent methods including MRT [46] and SoMoFormer [40]. We visualize the predicted motion sequences from the 3DPW-SoMoF/RC test set, see Fig. 6. Compared with MRT’s predictions which tend to be static, and SoMoFormer’s predictions where the human skeletons appear unnaturally dis-

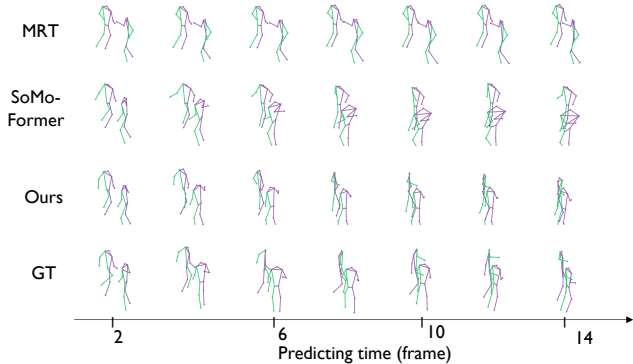


Figure 6. Visualization comparison on 3DPW-SoMoF/RC dataset. We compare the prediction by our method and two previous methods. Our method generates a more precise motion prediction.

torted, our method generates predictions that are not only more vivid and natural but also structurally correct. The visualization results validate the performance of our model.

5.6. Ablation Study

Relation information and relation supervision. To verify: i) the effect of the introduced relation information; ii) the rationality and effect of distance information and physical constraints information included in our designed relational information; and iii) the role of additional relation supervision, in this section, we study the performance of the model on the 3DPW-SoMoF/RC test set with different relation information inputs with or without relation supervision. Tab. 3 represents the VIM result with different settings on 3DPW-SoMoF/RC. We can conclude that: i) incorporating additional relation information, *e.g.*, distance information or physical constraints or both, can significantly improve model performance, proving the effectiveness of relation information in better modeling the joints interaction and movements; ii) both the distance relation information and physical constraints relation information contribute to the improvement of model performance, while the best performance is achieved when both are considered; iii) the relation supervision is also helpful since it can help model capture the interaction information hidden in distance changes.

Interaction modeling. To demonstrate the effectiveness of interaction modeling, we compare our model with other three settings: (i) single person model that uses the general self-attention to only capture the relationship between the intra-person joints; (ii) single person with relation model where the interaction between people is masked and we only provide and supervise the distance of joints inside the human body; (iii) simple multi-person model which replaces the joint-relation fusion layer in our model with the general self-attention layer. The results are shown in Tab. 4. It is obvious that the interaction between persons

Table 3. Ablation of relation information and relation supervision in VIM on 3DPW-SoMoF/RC test sets. Each notation is defined as *Dist.*: the relation information only include distance information \mathbb{D}_x , *Phys.*: physical constraints include the adjacent matrix \mathbb{A} and the connectivity matrix \mathbb{C} . With both kinds of relation information and relation supervision, the model performs best.

Relation		Supervision	3DPW-SoMoF/RC					
<i>Dist.</i>	<i>Phys.</i>		AVG	100	240	500	600	900
✓			41.7	10.3	23.4	46.8	56.7	71.5
	✓		41.1	9.8	22.8	46.2	56.0	71.1
		✓	40.8	9.9	22.6	45.3	55.4	70.7
✓		✓	39.9	9.6	21.8	43.9	54.0	70.1
	✓	✓	40.1	9.8	22.1	44.6	54.1	69.7
✓	✓		40.3	9.8	22.4	45.0	54.6	69.8
✓	✓	✓	39.5	9.5	21.7	44.1	53.4	68.8

Table 4. Ablation of interaction modeling in VIM on 3DPW-SoMoF/RC test sets. With the interaction modeling and relation information, the model achieves the best prediction performance.

Joint	Relation	AVG	100	240	500	600	900
single-person		46.2	10.1	23.7	49.8	62.0	85.4
single-person	✓	45.1	9.6	22.6	48.0	60.1	85.1
multi-person		41.7	10.3	23.4	46.8	56.7	71.5
multi-person	✓	39.5	9.5	21.7	44.1	53.4	68.8

contributes the improvement of motion prediction performance, while with the modeling and utilization of relational information, our model can better capture the interaction information to predict future motion.

6. Conclusion

This paper proposes the Joint-Relation Transformer, a two-stream Transformer-based architecture for multi-person motion prediction, which introduces the relation information and designs a novel relation-aware attention to inject the relation information into joint movement. Extensive experiments show that our method achieves state-of-the-art performance on three datasets and qualitative results show the effectiveness of the learned attention matrices.

Limitation and future work. This work considers the deterministic multi-person motion prediction where the model only predicts once. A possible future work is to explore the stochastic multi-person motion prediction where the model is required to model the diverse future distribution. Also, we will further explore the multi-scale structure among multi-person to make a more precise prediction.

Acknowledgement. This research is supported by NSFC under Grant 62171276 and the Science and Technology Commission of Shanghai Municipal under Grant 21511100900 and 22DZ2229005.

References

- [1] Cmu graphics lab motion capture database. <http://mocap.cs.cmu.edu/>. 2, 5, 6
- [2] Vida Adeli, Ehsan Adeli, Ian Reid, Juan Carlos Niebles, and Hamid Rezatofighi. Socially and contextually aware human motion and pose forecasting. *IEEE Robotics and Automation Letters*, 5(4):6033–6040, 2020. 2, 6
- [3] Vida Adeli, Mahsa Ehsanpour, Ian Reid, Juan Carlos Niebles, Silvio Savarese, Ehsan Adeli, and Hamid Rezatofighi. Tripod: Human trajectory and pose dynamics forecasting in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13390–13400, 2021. 1, 2, 6
- [4] Judith Butepage, Michael J Black, Danica Kragic, and Hedvig Kjellstrom. Deep representation learning for human motion prediction and classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6158–6166, 2017. 1, 2
- [5] Lowik Chanussot, Abhishek Das, Siddharth Goyal, Thibaut Lavril, Muhammed Shuaibi, Morgane Riviere, Kevin Tran, Javier Heras-Domingo, Caleb Ho, Weihua Hu, et al. Open catalyst 2020 (oc20) dataset and community challenges. *ACS Catalysis*, 11(10):6059–6072, 2021. 2
- [6] Changan Chen, Yuejiang Liu, Sven Kreiss, and Alexandre Alahi. Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning. In *2019 international conference on robotics and automation (ICRA)*, pages 6015–6022. IEEE, 2019. 1
- [7] Hsu-kuang Chiu, Ehsan Adeli, Borui Wang, De-An Huang, and Juan Carlos Niebles. Action-agnostic human pose forecasting. In *2019 IEEE winter conference on applications of computer vision (WACV)*, pages 1423–1432. IEEE, 2019. 1, 2
- [8] Cameron Diao and Ricky Loynd. Relational attention: Generalizing transformers for graph-structured tasks. *arXiv preprint arXiv:2210.05062*, 2022. 3
- [9] Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs. *arXiv preprint arXiv:2012.09699*, 2020. 3
- [10] Shaoheng Fang, Zi Wang, Yiqi Zhong, Junhao Ge, and Siheng Chen. Tbp-former: Learning temporal bird’s-eye-view pyramid for joint perception and prediction in vision-centric autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1368–1378, 2023. 1
- [11] 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. 1, 2
- [12] Robert X Gao, Lihui Wang, Peng Wang, Jianjing Zhang, and Hongyi Liu. Human motion recognition and prediction for robot control. In *Advanced Human-Robot Collaboration in Manufacturing*, pages 261–282. Springer, 2021. 1
- [13] Utkarsh Gaur, Yingying Zhu, Bi Song, and A Roy-Chowdhury. A “string of feature graphs” model for recognition of complex activities in natural videos. In *2011 International conference on computer vision*, pages 2595–2602. IEEE, 2011. 1
- [14] Partha Ghosh, Jie Song, Emre Aksan, and Otmar Hilliges. Learning human motion models for long-term predictions. In *2017 International Conference on 3D Vision (3DV)*, pages 458–466. IEEE, 2017. 2
- [15] Haifeng Gong, Jack Sim, Maxim Likhachev, and Jianbo Shi. Multi-hypothesis motion planning for visual object tracking. In *2011 International Conference on Computer Vision*, pages 619–626. IEEE, 2011. 1
- [16] Anand Gopalakrishnan, Ankur Mali, Dan Kifer, Lee Giles, and Alexander G Ororbia. A neural temporal model for human motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12116–12125, 2019. 2
- [17] Wen Guo, Xiaoyu Bie, Xavier Alameda-Pineda, and Francesc Moreno-Noguer. Multi-person extreme motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13053–13064, 2022. 2
- [18] Weihua Hu, Matthias Fey, Hongyu Ren, Maho Nakata, Yuxiao Dong, and Jure Leskovec. Ogb-lsc: A large-scale challenge for machine learning on graphs. *arXiv preprint arXiv:2103.09430*, 2021. 2
- [19] Md Shamim Hussain, Mohammed J Zaki, and Dharmashankar Subramanian. Edge-augmented graph transformers: Global self-attention is enough for graphs. *arXiv preprint arXiv:2108.03348*, 2021. 3
- [20] 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
- [21] Ling Min Serena Khoo, Hai Leong Chieu, Zhong Qian, and Jing Jiang. Interpretable rumor detection in microblogs by attending to user interactions. In *Proceedings of the AAAI conference on artificial intelligence*, 2020. 3
- [22] Devin Kreuzer, Dominique Beaini, Will Hamilton, Vincent Létourneau, and Prudencio Tossou. Rethinking graph transformers with spectral attention. *Advances in Neural Information Processing Systems*, 34:21618–21629, 2021. 3
- [23] 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
- [24] 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. 1, 2
- [25] Maosen Li, Siheng Chen, Xu Chen, Ya Zhang, Yanfeng Wang, and Qi Tian. Symbiotic graph neural networks for 3d skeleton-based human action recognition and motion prediction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6):3316–3333, 2021. 2
- [26] Maosen Li, Siheng Chen, Zihui Liu, Zijing Zhang, Lingxi Xie, Qi Tian, and Ya Zhang. Skeleton graph scattering networks for 3d skeleton-based human motion prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 854–864, 2021. 2

- [27] 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
- [28] Maosen Li, Siheng Chen, Yangheng Zhao, Ya Zhang, Yanfeng Wang, and Qi Tian. Multiscale spatio-temporal graph neural networks for 3d skeleton-based motion prediction. *IEEE Transactions on Image Processing*, 30:7760–7775, 2021. 2
- [29] Kevin Lin, Lijuan Wang, and Zicheng Liu. Mesh graphormer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12939–12948, 2021. 2
- [30] Zhenguang Liu, Shuang Wu, Shuyuan Jin, Qi Liu, Shijian Lu, Roger Zimmermann, and Li Cheng. Towards natural and accurate future motion prediction of humans and animals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10004–10012, 2019. 2
- [31] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *International Conference on Computer Vision*, pages 5442–5451, Oct. 2019. 6
- [32] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Srinath Sridhar, Gerard Pons-Moll, and Christian Theobalt. Single-shot multi-person 3d pose estimation from monocular rgb. In *3D Vision (3DV), 2018 Sixth International Conference on*. IEEE, sep 2018. 2, 6
- [33] Grégoire Mialon, Dexiong Chen, Margot Selosse, and Julien Mairal. Graphit: Encoding graph structure in transformers. *arXiv preprint arXiv:2106.05667*, 2021. 2
- [34] Erxue Min, Yu Rong, Tingyang Xu, Yatao Bian, Peilin Zhao, Junzhou Huang, Da Luo, Kangyi Lin, and Sophia Ananiadou. Masked transformer for neighbourhood-aware click-through rate prediction. *arXiv preprint arXiv:2201.13311*, 2022. 2
- [35] Behnam Parsaefard, Saeed Saadatnejad, Yuejiang Liu, Taylor Mordan, and Alexandre Alahi. Learning decoupled representations for human pose forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 2294–2303, October 2021. 6
- [36] Dario Pavllo, David Grangier, and Michael Auli. Quaternet: A quaternion-based recurrent model for human motion. *arXiv preprint arXiv:1805.06485*, 2018. 2
- [37] Yu Rong, Yatao Bian, Tingyang Xu, Weiyang Xie, Ying Wei, Wenbing Huang, and Junzhou Huang. Self-supervised graph transformer on large-scale molecular data. *Advances in Neural Information Processing Systems*, 33:12559–12571, 2020. 2
- [38] Bohan Tang, Yiqi Zhong, Chenxin Xu, Wei-Tao Wu, Ulrich Neumann, Ya Zhang, Siheng Chen, and Yanfeng Wang. Collaborative uncertainty benefits multi-agent multi-modal trajectory forecasting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 1
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2
- [40] Edward Vendrow, Satyajit Kumar, Ehsan Adeli, and Hamid Reza Tofighi. Somoformer: Multi-person pose forecasting with transformers. *arXiv preprint arXiv:2208.14023*, 2022. 1, 2, 3, 6, 7
- [41] Timo von Marcard, Roberto Henschel, Michael Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *European Conference on Computer Vision (ECCV)*, sep 2018. 5, 6
- [42] Tuan-Hung Vu, Sebastien Ambellouis, Jacques Boonaert, and Abdelmalik Taleb-Ahmed. Anomaly detection in surveillance videos by future appearance-motion prediction. In *VISIGRAPP (5: VISAPP)*, pages 484–490, 2020. 1
- [43] Fabien B Wagner, Jean-Baptiste Mignardot, Camille G Le Goff-Mignardot, Robin Demesmaeker, Salif Komi, Marco Capogrosso, Andreas Rowald, Ismael Seáñez, Miroslav Caban, Elvira Pirondini, et al. Targeted neurotechnology restores walking in humans with spinal cord injury. *Nature*, 563(7729):65–71, 2018. 1
- [44] Chenxi Wang, Yunfeng Wang, Zixuan Huang, and Zhiwen Chen. Simple baseline for single human motion forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2260–2265, 2021. 6
- [45] Jack Wang, Aaron Hertzmann, and David J Fleet. Gaussian process dynamical models. *Advances in neural information processing systems*, 18, 2005. 2
- [46] Jiashun Wang, Huazhe Xu, Medhini Narasimhan, and Xiaolong Wang. Multi-person 3d motion prediction with multi-range transformers. *Advances in Neural Information Processing Systems*, 34:6036–6049, 2021. 1, 2, 3, 6, 7
- [47] Mao Wei, Liu Miaomiao, and Salzmann Mathieu. History repeats itself: Human motion prediction via motion attention. In *ECCV*, 2020. 1, 2, 7
- [48] Mao Wei, Liu Miaomiao, Salzmann Mathieu, and Li Hongdong. Learning trajectory dependencies for human motion prediction. In *ICCV*, 2019. 1, 2, 6, 7
- [49] Di Wu and Ling Shao. Leveraging hierarchical parametric networks for skeletal joints based action segmentation and recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 724–731, 2014. 2
- [50] Zhanghao Wu, Paras Jain, Matthew Wright, Azalia Mirhoseini, Joseph E Gonzalez, and Ion Stoica. Representing long-range context for graph neural networks with global attention. *Advances in Neural Information Processing Systems*, 34:13266–13279, 2021. 2
- [51] 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*, pages 6498–6507, 2022. 1
- [52] Chenxin Xu, Weibo Mao, Wenjun Zhang, and Siheng Chen. Remember intentions: Retrospective-memory-based trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6488–6497, 2022. 1

- [53] 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](#)
- [54] Sirui Xu, Yu-Xiong Wang, and Liang-Yan Gui. Stochastic multi-person 3d motion forecasting. In *ICLR*, 2023. [2](#)
- [55] Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. Do transformers really perform badly for graph representation? *Advances in Neural Information Processing Systems*, 34:28877–28888, 2021. [3](#)
- [56] Jiawei Zhang, Haopeng Zhang, Congying Xia, and Li Sun. Graph-bert: Only attention is needed for learning graph representations. *arXiv preprint arXiv:2001.05140*, 2020. [2](#)
- [57] Jianan Zhao, Chaozhuo Li, Qianlong Wen, Yiqi Wang, Yuming Liu, Hao Sun, Xing Xie, and Yanfang Ye. Gophormer: Ego-graph transformer for node classification. *arXiv preprint arXiv:2110.13094*, 2021. [3](#)