

Spatio-Temporal Gating-Adjacency GCN for Human Motion Prediction

Chongyang Zhong^{1,2}, Lei Hu^{1,2}, Zihao Zhang^{1,2}, Yongjing Ye^{1,2}, Shihong Xia^{1,2*}

¹Institute of Computing Technology, Chinese Academy of Sciences; ²University of Chinese Academy of Sciences
 {zhongchongyang, hulei19z, zhangzihao, yeyongjing, xsh}@ict.ac.cn

Abstract

Predicting future motion based on historical motion sequence is a fundamental problem in computer vision, and it has wide applications in autonomous driving and robotics. Some recent works have shown that Graph Convolutional Networks(GCN) are instrumental in modeling the relationship between different joints. However, considering the variants and diverse action types in human motion data, the cross-dependency of the spatio-temporal relationships will be difficult to depict due to the decoupled modeling strategy, which may also exacerbate the problem of insufficient generalization. Therefore, we propose the Spatio-Temporal Gating-Adjacency GCN(GAGCN) to learn the complex spatio-temporal dependencies over diverse action types. Specifically, we adopt gating networks to enhance the generalization of GCN via the trainable adaptive adjacency matrix obtained by blending the candidate spatio-temporal adjacency matrices. Moreover, GAGCN addresses the cross-dependency of space and time by balancing the weights of spatio-temporal modeling and fusing the decoupled spatio-temporal features. Extensive experiments on Human 3.6M, AMASS, and 3DPW demonstrate that GAGCN achieves state-of-the-art performance in both short-term and long-term predictions.

1. Introduction

The aim of human motion prediction is to predict the motion trend of the skeleton-based human body in the future period from a given historical motion sequence, which is a significant computer vision task with many potential applications, such as autonomous driving, human-robotics interaction, target tracking, and motion planning.

The skeleton-based human motion sequence is a structured time series, which means that the movement of a single joint is affected by the coupling of spatial connections with other joints and the temporal trajectory tendency. We call these complex spatio-temporal relationships as cross-




Figure 1. The illustration of our method. Given the historical input human motion sequence, we try to predict the future motion sequence by enhancing, balancing, and fusing two key factors, i.e. the joint dependencies and the temporal correlations.

dependency. The challenges lie in motion prediction are mainly two-fold. First, earlier literature based on Recurrent Neural Network(RNN), such as LSTM and GRU suggests that predicting the long-term sequence will meet the inherent error accumulation problem [5, 7–10, 14, 27, 32, 34]. Though subsequent convolution-based approaches [3, 11, 17] for sequence-to-sequence prediction reduce the error of long-term prediction to some extent, the error accumulation is still a problem to be solved. Second, it is difficult to model the spatio-temporal relationships since the skeleton-based human motion is very complex and diverse. Rather than using the basic motion representation(joint angle, position, and velocity) directly or extracting the spatial features of the motion with a simple fully connected layer, recent studies try to use GCN to depict the spatio-temporal relationships [6, 18–20, 22, 24–26, 28].

Though GCN-based works are instrumental for solving the long-term prediction problem to some extent, there are two issues to be explored: 1. The inter-joint and inter-frame relationships will change with the motion variance and action types, therefore a stable adjacency matrix will lead to

*Corresponding author.

inherent poor generalization on multi-action motions; 2. Direct concatenation of the decoupled spatial and temporal features can not fully explore the cross-dependency of the spatio-temporal relationships.

In this paper, we propose the Spatio-Temporal Gating-Adjacency GCN(GAGCN) to learn the complex spatio-temporal dependencies over diverse action types. To solve the above two issues, our key idea mainly consists of two parts, namely, **the enhancing strategy** and **the balancing and fusing strategy**(shown in Fig. 1). First, given different historical motion sequences, the gating network in our GAGCN output corresponding blending coefficients which are then used to blend the trainable candidate adjacency matrices. The inter-joint and inter-frame relationships of different motions are learned by the adaptive blended adjacency matrix dynamically, which enhances the generalization of our model on multi-action motions. Second, the proposed GAGCN can be utilized to balance the weight of spatial and temporal modeling by scaling the number of candidate matrices. And the spatio-temporal features are fused to mine the hidden cross-dependency of spatio-temporal relationships from the historical motion sequence.

Extensive experiments are conducted on Human3.6M [12], AMASS [23] and 3DPW [33]. We demonstrate that our method outperforms state-of-the-art methods in both short-term and long-term motion predictions. The main contributions of our work can be summarized as follows:

1. To the best of our knowledge, we are the first to use the gating network to enhance the generalization of GCN on human motion prediction. The adaptive adjacency matrix obtained by blending candidate matrices helps to enhance the scalability of our network across multi-action motions.
2. We capture the cross-dependency of space and time by balancing and fusing the decoupled joint dependencies and temporal correlations to learn the more expressive embedding features.
3. We carry out extensive experiments on Human3.6M, AMASS and 3DPW both quantitatively and qualitatively to demonstrate that the results of our method outperform state-of-the-art works.

2. Related Works

Human Motion Prediction Traditional works on motion prediction attempt to use traditional statistical methods such as hidden Markov models [2] and Gaussian process hidden variable models [36], which have limitations in dealing with the high-dimensional dynamics of human motion and yield unsatisfactory results. With the development of

deep neural networks, exciting progress has been made in motion prediction. Some works use RNN to model the temporal correlations of human motion [5, 7–10, 14, 27, 32, 34]. However, these frame-by-frame methods perform poorly on long-term motion prediction due to their inherent error accumulation problem, and RNN-based networks suffer from first-frame discontinuity. To address these issues, researchers have attempted to improve the prediction results of RNN-based networks using sequence-to-sequence residual models [27], generative adversarial learning [10], and imitation learning [34]. In contrast to frame-by-frame framework, sequence-to-sequence method can effectively reduce cumulative error in long-term prediction, which includes convolution-based [3, 11, 17] and attention-based mechanisms [4, 24, 26]. The convolution-based approach treats the historical sequence as an entirety and extracts motion features in spatial or temporal dimensions, while the attention-based approach uses the attention model to learn the joint-to-joint and frame-to-frame dependencies.

Recently, graph convolutional networks(GCN) [16] have achieved state-of-the-art results in motion prediction [6, 18–20, 22, 24–26, 28]. Researchers use GCN with trainable adjacency matrices to model the joint dependencies of human motion. These methods learn the spatial properties of human motion by dividing the spatial properties into skeletal connections and implicit non-physical connections between individual joints [6], providing semantic prior knowledge to the network [22], dividing the human body into multi-scale [19].

Although the above works have made encouraging progress, most works treat temporal and spatial modeling in a decoupled manner and directly concatenate them while the spatial dependencies of motion are often coupled with the global temporal trajectory. To address this problem, GAGCN is proposed to learn the cross-dependency of space and time by balancing the weight of spatio-temporal modeling and fusing spatio-temporal features, which helps us to capture the spatio-temporal relationships simultaneously.

Spatio-Temporal Modeling for Human Motion To the best of our knowledge, the first work to simultaneously model spatio-temporal relationships is SRNN [14]. They use a graph model to represent the human body, where the joint nodes and edge nodes are composed of RNN, thereby being the first to achieve long-term motion prediction. Another work that combines spatio-temporal modeling more closely is STGCN [37], where they encode both the spatial connection of human joints in a single frame and the temporal connection of the same joint between frames into a single adjacency matrix of GCN. Although their work has made impressive progress in action recognition, it is limited by the constant adjacency matrix. Recently, a newly proposed method named Space-Time-Separable GCN [28] performs spatio-temporal modeling by factorizing the train-

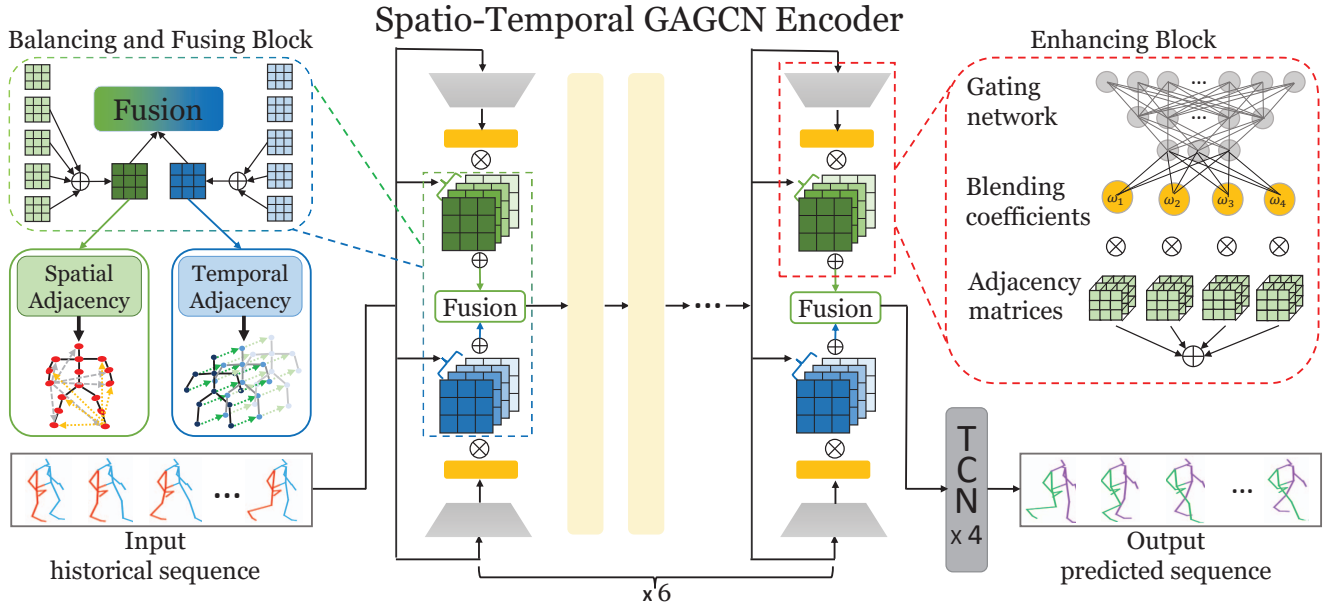


Figure 2. The overview of the proposed GAGCN network. We use Spatio-Temporal Gating-Adjacency GCN (GAGCN) as the encoder to learn the spatio-temporal dependence of the historical motion sequence, and then use TCN as a decoder. We first feed the feature from the previous layer into a spatial gating network and a temporal gating network respectively, obtaining the blending coefficients $\{\omega_s^i\}$ and $\{\omega_t^i\}$. Then we blend the spatial(temporal) adjacency matrix using the blending coefficients to create the adaptive spatial(temporal) adjacency matrix. Finally, we fuse the spatial and temporal dependencies with the Kronecker product to output the features for the next layer.

able adjacency matrix into temporal and spatial to achieve state-of-the-art performance in motion prediction.

Nonetheless, considering the variants and diverse action types in human motion data, a stable adjacency matrix can not effectively capture the changing dependencies between joints and between frames, resulting in the poor generalization of GCN. Mixture of Expert(MoE) [13, 15] is a traditional machine learning method that uses blending coefficients generated by a gating network to blend multiple experts. For human motion, the gating network acts as a motion classifier to automatically calculate the probability that the input motion belongs to each class of motion and blends the results of the relevant experts to obtain the optimal output, which greatly improves the generalization of multiple human motion models [21, 29–31, 38].

Inspired by MoE, we apply the gating network on the adjacency matrices to enhance the generalization of GCN. We first adopt several candidate adjacency matrices as experts in MoE, then we use the gating network to learn the adaptive adjacency matrix by blending the candidate adjacency matrices according to different inputs. The adaptive adjacency matrix can capture the dynamic relationships in human motion, which is helpful to generalize across diverse action types.

3. Our Method

Problem Formulation The purpose of skeleton-based human motion prediction is to predict the future pose se-

quence given the historical pose sequence. We denote the historical pose sequence as $X_{1:T} = \{x_1, x_2, \dots, x_T\}$ with T frames, and the predicted motion sequence of future t time steps as $X_{T+1:T+t} = \{x_{T+1}, x_{T+2}, \dots, x_{T+t}\}$, where x_i is usually represented as 3D coordinates or joint angle of the N body joints.

Overview As is shown in Fig. 2, we adopt the encoder-decoder structure to make motion prediction. To better retrieve the cross spatio-temporal dependency of the historical motion sequence, we propose Gating-Adjacency GCN (GAGCN) as the encoder which consists of three parts. First, the features from the previous layer are fed into a spatial gating network and a temporal gating network respectively to get the blending coefficients $\{\omega_s^i\}$ and $\{\omega_t^i\}$. Then we blend the spatial adjacency matrix and temporal adjacency matrix by using the estimated blending coefficients to obtain the adaptive adjacency matrix. Finally, we fuse the spatial and temporal dependencies with the Kronecker product to output the features for the next layer. For the decoder, given the latent motion representation after passing through 6 GAGCN layers, we use the Temporal Convolutional Networks(TCN) to predict the future sequence.

3.1. Review of GCN

In recent years, GCN-based networks have been widely utilized for the modeling of spatio-temporal dependencies of the structural time series and made inspiring progress, which provides a means of human motion prediction.

Specifically, we represent the skeleton-based pose as a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, in which \mathcal{V} is the joint-node set, \mathcal{E} is edge set. The joint node features are 3D coordinates or joint angles and the edge is related to the adjacency matrix $A \in \mathbb{R}^{N \times N}$.

The state-of-the-art works use the trainable adjacency matrix to replace the constant adjacency matrix, which can model not only skeletal connections but also the implicit dependencies of joints without natural connections, making GCN more powerful for learning spatial dependencies. A single layer with the trainable adjacency matrix can be expressed as:

$$H^{l+1} = f(H^l; A, W) = \sigma(AH^lW^l) \quad (1)$$

where $A, H^l \in \mathbb{R}^{N \times F^l}$ and $W^l \in \mathbb{R}^{F^l \times F^{l+1}}$ are the trainable adjacency matrix, input feature and the trainable transformation matrix, respectively.

3.2. Spatio-Temporal Gating-Adjacency GCN

Most prior works perform decoupled modeling and direct concatenation for spatio-temporal relationships without considering their cross-dependency, which makes it difficult to accurately depict such complex spatio-temporal relationships. Additionally, since the inter-joint and inter-frame relationships will change with the motion variance and action types, a stable adjacency matrix will lead to inherent poor generalization on multi-action motions. Thus, we propose Spatio-Temporal Gating-Adjacency GCN(GAGCN) to cope with these issues.

Gating Adjacency As shown in prior works [6, 18–20, 22, 24–26, 28], a stable trainable adjacency matrix can handle spatio-temporal dependencies to some extent. However, the relationship will be difficult to depict when it comes to the variants and diverse action types in human motion data. Our "Enhancing Block"(shown in the right part of Fig. 2) is motivated by the observation that the spatio-temporal relationships are changing with the diverse variances and action types. Therefore, we aim to find the adaptive spatio-temporal relationships to cope with multi-action motion prediction.

The Mixture of Experts(MoE) is a classic machine learning method, which is proven to be able to enhance the generalization of human motion models [21, 29–31, 38]. The gating network is regarded as a motion classifier to automatically calculates the probability of which motion class the input motion belongs to. The results of the relevant experts are blended to obtain the adaptive output. Therefore, the generalization of the human motion model is largely enhanced. Inspired by MoE, we apply the gating network on GCN to learn the blended adjacency matrix, which is adaptive to the diverse motion variances and action types.

Different from the traditional MoE-based methods which apply weighted sum on all of each expert's network param-

eters, we only use the gating network to blend the adjacency matrix. This can make our network lightweight and ensure that only the feature learning process is affected while keeping the feature transferring process unaffected. Specifically, given the features from the previous layer, the gating network in our GAGCN output several blending coefficient parameters, which can be denoted as follows:

$$\{\omega^i\} = \text{Gating}(H) = \text{softmax}(FC(H)) \quad (2)$$

where FC denotes the 3 fully connected layers, softmax is the activation function, H is the input feature, $\{\omega^i\}$ is the set of blending coefficients. Then the blending coefficients are used to blend the candidate trainable adjacency matrices to get the adaptive adjacency matrix:

$$A = \sum_i A^i \cdot \omega^i \quad (3)$$

where $\{A^i\}$ is the set of trainable adjacency matrices and A is the adaptive adjacency matrix.

Spatio-Temporal Modeling Previous literature usually learn spatial and temporal dependencies in a decoupled manner and directly concatenate them, therefore the cross-dependency of the spatio-temporal information is still not fully explored. To address this issue, we propose the balancing and fusion strategy shown in the left dotted box in Fig. 2. The key idea is to adjust the number of candidate adjacency matrices to control and balance the weight between the spatial and temporal modeling, and then fuse spatio-temporal features with the Kronecker product.

Specifically, our balancing strategy is designed as follows: At first, we divide the adjacency matrix A into A_s and A_t , as shown in the left of Fig. 2. Following [35], we treat all channels of a joint as a node instead of regarding each channel of a joint as a node, which can significantly reduce the size of the adjacency matrix and maintain the correlation of different channels of the same node. $A_s \in \mathbb{R}^{N \times N}$ represents the inter-dependencies between joint and joint, whether they have skeletal connections or not. Meanwhile, $A_t \in \mathbb{R}^{T \times T}$ is trained to learn the frame-to-frame dependencies in the historical sequence.

Then, based on Equ. 2 and Equ. 3, we adopt two gating networks to learn the blending coefficients for spatial and temporal respectively, then blend the candidate adjacency matrices to obtain the adaptive adjacency matrix. We further update these two equations into the following form:

$$\{\omega_k^i\} = \text{Gating}_k(H), \quad \mathcal{A}_k = \sum_i A_k^i \cdot \omega_k^i, \quad (4)$$

where \mathcal{A}_k is the adaptive adjacency matrix, and the subscript $k \in \{s, t\}$ indicate "spatial" s or "temporal" t . It is worth noting that the number of candidate adjacency matrices $q \in \{n, m\}$ can be adjusted, which indicates the complexity of spatial and temporal modeling. The proposed

GAGCN can be utilized to balance the weight of spatial and temporal modeling like a scale by adjusting the number of candidate matrices. For example, we use $n = 4, m = 3$ on Human 3.6M and $n = 6, m = 4$ on AMASS.

As for the fusing strategy, a single layer of GAGCN can be formulated as follows:

$$H^{l+1} = \sigma((\mathcal{A}_s^l \otimes \mathcal{A}_t^l)H^l W^l) \quad (5)$$

where $H^l \in \mathbb{R}^{w^l \times N \times T}$, $W^l \in \mathbb{R}^{w^l \times w^{l+1}}$, \mathcal{A}_s^l and \mathcal{A}_t^l are the input feature, trainable transformation matrix, adaptive spatial adjacency matrix and adaptive temporal adjacency matrix of layer l , respectively. \otimes denote the Kronecker product. The temporal feature and spatial feature are fused by the Kronecker product to ensure that we can find the hidden cross-dependency of spatio-temporal relationships from the historical motion sequence.

The fused features are fed into the next layer for further learning. Through 6 GAGCN layers, we extract flexible and implicit dependencies between joints and frames represented as the spatio-temporal features. Finally, the features are passed into TCN decoders to predict future sequence, which is confirmed to make better performance and less error accumulation than RNN [1].

3.3. Training

Our training process is end-to-end and supervised. With the help of the highly expressive spatio-temporal features extracted by the GAGCN encoder, our network uses a relatively simple loss function to get state-of-the-art results.

For 3D joint coordinates representation, we use MPJPE loss:

$$L_{MPJPE} = \frac{1}{N \cdot t} \sum_{i=1}^t \sum_{j=1}^N \|\tilde{p}_{ij} - p_{ij}\|_2 \quad (6)$$

where p_{ij} represents the predicted 3D coordinates of j_{th} joint in i_{th} frame, and \tilde{p}_{ij} is the corresponding ground truth.

For the angle-based representation, we use MAE loss:

$$L_{MAE} = \frac{1}{N \cdot t} \sum_{i=1}^t \sum_{j=1}^N |\tilde{x}_{ij} - x_{ij}| \quad (7)$$

where x_{ij} represents the predicted joint angle in exponential map of j_{th} joint in i_{th} frame, and \tilde{x}_{ij} is the corresponding ground truth.

4. Experiments

In this section, we evaluate the proposed motion prediction method. First, we will show the details of the used benchmark dataset and baselines in Sec. 4.1. The quantitative comparison results with the state-of-the-art method will be given in Sec. 4.2. Then, we will analyze the main components of our method in Sec. 4.3. Finally, we will show the qualitative evaluation in Sec. 4.4. The implementation details are shown in the supplementary material.

4.1. Datasets and Baselines

The datasets used in our experiments include Human 3.6M [12], AMASS [23], and 3DPW [33]. We will introduce these 3 datasets as follows:

Human 3.6M Human 3.6M is the most used dataset in the field of motion prediction. Human 3.6M has 3.6 million 3D poses, consisting of 15 motion categories from 7 subjects. We down-sample the frame rate to 25Hz. Following [24, 27], we use subjects 1,6,7,8,9 for training, subject 11 for validation and subject 5 for testing.

AMASS The Archive of Motion Capture as Surface Shapes(AMASS) dataset is a recently published human motion dataset, which gathers 18 existing mocap datasets, such as CMU, KIT, and BMLrub. We down-sample the frame rate to 25Hz as for Human 3.6M. Then, following [24], we select 8 datasets from AMASS for training, 4 datasets for validation and 1 dataset(BMLrub) for testing.

3DPW The 3D Pose in the Wild dataset consists of both indoor and outdoor actions, which contains 51,000 frames captured at 30Hz. We down-sample the frame rate to 25Hz as for Human 3.6M. We only use 3DPW to test the generalization of the models trained on AMASS.

Metrics and Baselines Our model can be trained on both 3D coordinates representation and angle-based representation. Thus, we evaluate the results on both 3D coordinates errors and angle errors. Following [24], We adopt the MPJPE metrics for 3D coordinates representation and MAE angle error metrics for angle-based representation. We compare our approach with Res-GRU [27], ConSeq2Seq [17], LTD-10-25 [25], HRI [24], and STSGCN [28] on Human 3.6M and LTD-10-25 [25], HRI [24], and STSGCN [28] on AMASS and 3DPW. We adapt the code and the pre-trained models released by the authors to evaluate their results. Note that HRI [24] takes the past 50 frames as input to predict the future 25 frames while others takes the past 10 frames as input to predict the future 25 frames.

4.2. Comparisons with the State-of-the-art Methods

Human 3.6M Because of the ambiguity in the angle-based representation, most recent works use 3D coordinates to measure the accuracy of motion prediction, i.e. MPJPE. As in previous work, we predict future motion for 25 frames(1000ms) based on a 10(400ms) frames historical motion sequence. We select 14 action types from Human 3.6M and randomly select 8 sequences for each motion to calculate the average error. As in Table. 1, we show the comparison of the short-term and long-term prediction of our model and the baselines on Human 3.6M.

Our method outperforms all state-of-the-art models over almost all time horizons. In particular, thanks to the adaptive adjacency matrix, our model makes larger improvements in action types that are difficult to predict, such as "Posing" and "Sitting down". Moreover, our method has

	Walking						Eating						Smoking					
milliseconds	80	160	320	400	560	1000	80	160	320	400	560	1000	80	160	320	400	560	1000
Res-GRU [27]	23.2	40.9	61.0	66.1	71.6	79.1	16.8	31.5	53.5	61.7	74.9	98.0	18.9	34.7	57.5	65.4	78.1	102.1
ConSeq2Seq [17]	17.7	33.5	56.3	63.6	72.2	82.3	11.0	22.4	40.7	48.4	61.3	87.1	11.6	22.8	41.3	48.9	60.0	81.7
LTD-10-25 [25]	12.6	23.6	39.4	44.5	51.8	60.9	7.7	15.8	30.5	37.6	50.0	74.1	8.4	16.8	32.5	39.5	51.3	73.6
HRI [24]	10.0	19.5	34.2	39.8	47.4	58.1	6.4	14.0	28.7	36.2	50.0	75.7	7.0	14.9	29.9	36.4	47.6	69.5
STSGCN [28]	10.7	16.9	29.1	32.9	40.6	51.8	6.8	11.3	22.6	25.4	33.9	52.4	7.2	11.6	22.3	25.8	33.6	50.0
Ours	10.3	16.1	28.8	32.4	39.9	51.1	6.4	11.5	21.7	25.2	31.8	51.4	7.1	11.8	21.7	24.3	31.1	48.7
	Discussion						Directions						Greeting					
milliseconds	80	160	320	400	560	1000	80	160	320	400	560	1000	80	160	320	400	560	1000
Res-GRU [27]	25.7	47.8	80.0	91.3	109.5	131.8	21.6	41.3	72.1	84.1	101.1	129.1	31.2	58.4	96.3	108.8	126.1	153.9
ConSeq2Seq [17]	17.1	34.5	64.8	77.6	98.1	129.3	13.5	29.0	57.6	69.7	86.6	115.8	22.0	45.0	82.0	96.0	116.9	147.3
LTD-10-25 [25]	12.2	25.8	53.9	66.7	87.6	118.6	9.2	20.6	46.9	58.8	76.1	108.8	16.7	33.9	67.5	81.6	104.3	140.2
HRI [24]	10.2	23.4	52.1	65.4	86.6	119.8	7.4	18.4	44.5	56.5	73.9	106.5	13.7	30.1	63.8	78.1	101.9	138.8
STSGCN [28]	9.8	16.8	33.4	40.2	53.4	78.8	7.4	13.5	29.2	34.7	47.6	71.0	12.4	21.8	42.1	49.2	64.8	91.6
Ours	9.7	17.1	31.4	38.9	53.1	76.9	7.3	12.8	30.3	34.5	45.8	69.9	11.8	20.1	40.5	48.4	62.3	87.7
	Phoning						Posing						Purchases					
milliseconds	80	160	320	400	560	1000	80	160	320	400	560	1000	80	160	320	400	560	1000
Res-GRU [27]	21.1	38.9	66.0	76.4	94.0	126.4	29.3	56.1	98.3	114.3	140.3	183.2	28.7	52.4	86.9	100.7	122.1	154.0
ConSeq2Seq [17]	13.5	26.6	49.9	59.9	77.1	114.0	16.9	36.7	75.7	92.9	122.5	187.4	20.3	41.8	76.5	89.9	111.3	151.5
LTD-10-25 [25]	10.2	20.2	40.9	50.9	68.7	105.1	12.5	27.5	62.5	79.6	109.9	171.7	15.5	32.3	63.6	77.3	99.4	135.9
HRI [24]	8.6	18.3	39.0	49.2	67.4	105.0	10.2	24.2	58.5	75.8	107.6	178.2	13.0	29.2	60.4	73.9	95.6	134.2
STSGCN [28]	8.2	13.7	26.9	30.9	41.8	66.1	9.9	18.0	38.2	45.6	64.3	106.4	11.9	21.3	42.0	48.7	63.7	93.5
Ours	8.8	13.5	25.5	28.7	41.1	66.0	10.1	17.0	35.5	45.1	63.3	99.1	11.9	20.7	41.8	47.6	62.1	85.1
	Sitting						Sitting Down						Taking Photo					
milliseconds	80	160	320	400	560	1000	80	160	320	400	560	1000	80	160	320	400	560	1000
Res-GRU [27]	23.8	44.7	78.0	91.2	113.7	152.6	31.7	58.3	96.7	112.0	138.8	187.4	21.9	41.4	74.0	87.6	110.6	153.9
ConSeq2Seq [17]	13.5	27.0	52.0	63.1	82.4	120.7	20.7	40.6	70.4	82.7	106.5	150.3	12.7	26.0	52.1	63.6	84.4	128.1
LTD-10-25 [25]	10.4	21.4	45.4	57.3	78.5	118.8	17.0	33.4	61.6	74.4	99.5	144.1	9.9	20.5	43.8	55.2	76.8	120.2
HRI [24]	9.3	20.1	44.3	56.0	76.4	115.9	14.9	30.7	59.1	72.0	97.0	143.6	8.3	18.4	40.7	51.5	72.1	115.9
STSGCN [28]	9.1	15.1	29.9	35.0	47.7	75.2	14.4	23.7	41.9	47.9	63.3	94.3	8.2	14.2	29.7	33.6	47.0	76.9
Ours	9.3	14.4	29.6	38.5	45.4	71.1	14.1	24.8	40.0	47.4	62.8	84.1	8.5	13.9	28.8	35.1	45.2	70.0
	Waiting						Walking Dog						Average					
milliseconds	80	160	320	400	560	1000	80	160	320	400	560	1000	80	160	320	400	560	1000
Res-GRU [27]	23.8	44.2	75.8	87.7	105.4	135.4	36.4	64.8	99.1	110.6	128.7	164.5	25.3	46.8	78.2	89.9	108.2	139.4
ConSeq2Seq [17]	14.6	29.7	58.1	69.7	87.3	117.7	27.7	53.6	90.7	103.3	122.4	162.4	16.6	33.5	62.0	73.5	92.1	126.8
LTD-10-25 [25]	10.5	21.6	45.9	57.1	75.1	106.9	22.9	43.5	74.5	86.4	105.8	142.2	12.6	25.5	50.6	61.9	81.1	115.8
HRI [24]	8.7	19.2	43.4	54.9	74.5	108.2	20.1	40.3	73.3	86.3	108.2	146.9	10.4	22.1	46.5	57.5	76.6	112.3
STSGCN [28]	8.6	14.7	29.6	35.2	47.3	72.0	17.6	29.4	52.6	59.6	74.7	102.6	10.2	17.3	33.5	38.9	51.7	77.3
Ours	8.5	14.1	29.8	33.8	45.9	69.3	17.0	28.8	50.1	59.4	70.1	91.3	10.1	16.9	32.5	38.5	50.0	72.9

Table 1. MPJPE error comparison for both short-term and long-term predictions on 14 action types in Human 3.6M. The best results are shown in bold. Our method outperforms all baselines on average over all time horizons. It is worth noting that our model makes larger improvements in action types that are difficult to predict, such as "Posing" and "Sitting down". Moreover, our method has significant advantages in long-term (1000ms) motion prediction.

significant advantages in long-term (1000ms) motion prediction. Nonetheless, there are very few time points where our method does not perform best. These time points are in short term with small prediction errors for all methods, thus it is reasonable to have a marginal error. The bottom right of Table. 1 is the average error of all action types, where our method performs better than all the comparative methods for whole time horizons.

Additionally, we demonstrate the average angle errors on Human 3.6M in Table. 2 with the same setting for MPJPE metrics. The results illustrate that our method also achieves state-of-the-art prediction in angle-based representation.

AMASS & 3DPW We demonstrate the short-term and long-term prediction results on AMASS-BMLrub in Table. 3. We train the model on 8 datasets from AMASS and

use BMLrub for testing. AMASS has much more subjects and motion sequences than Human 3.6M, which is more suitable to test the generalization of the model. Our method outperforms all the baselines on AMASS, which proves that our model can indeed enhance the generalization of GCN.

The model trained on AMASS is further tested on 3DPW, and the results are shown in Table. 4. The significantly better results compared with other methods provide another strong evidence of our model's generalization across different datasets.

4.3. Ablation Study

We perform ablation studies to evaluate the effect of two key component in our method, i.e. enhancing block, balancing block. The effect of fusion block can be found in the

Human 3.6M-average						
milliseconds	80	160	320	400	560	1000
Res-GRU	0.36	0.67	1.02	1.15	-	-
conSeq2Seq	0.38	0.68	1.01	1.13	1.35	1.82
LTD-10-25	0.30	0.54	0.86	0.97	1.15	1.59
HRI	0.27	0.52	0.82	0.94	1.14	1.57
STSGCN	0.24	0.39	0.59	0.66	0.79	1.09
Ours	0.24	0.38	0.54	0.65	0.74	1.02

Table 2. Average MAE angle error comparison on Human 3.6M (note that Res-GRU [27] has no long-term prediction results). The best results are shown in bold. Our method achieves state-of-the-art prediction in angle-based representation

AMASS-BMLrub-average						
milliseconds	80	160	320	400	560	1000
LTD-10-25	11.0	20.7	37.8	45.3	57.2	75.2
HRI	11.3	20.7	35.7	42.0	51.7	67.2
STSGCN	10.0	12.5	21.8	24.5	31.9	45.5
Ours	10.0	11.9	20.1	24.0	30.4	43.1

Table 3. Average MPJPE error comparison on AMASS-BMLrub. The best results are shown in bold. Our method outperforms all the baselines, which proves that our model can indeed enhance the generalization of GCN across datasets.

3DPW-average						
milliseconds	80	160	320	400	560	1000
LTD-10-25	12.6	23.2	39.7	46.6	57.9	75.5
HRI	12.6	23.1	39.0	45.4	56.0	73.7
STSGCN	8.6	12.8	21.0	24.5	30.4	42.3
Ours	8.4	11.9	18.7	23.6	29.1	39.9

Table 4. Average MPJPE error comparison on 3DPW. The best results are shown in bold. The significantly better results than other methods provide another strong evidence of our model’s generalization.

supplementary material.

Effect of Enhancing Block We have shown the generalization of our method across different datasets in 4.2, and then we will further show the generalization of our method across different action types. The results are shown in Table. 5. We explore the generalization of our model by testing on unseen action types (Walking Together). The results in the second row of the table are significantly better than the first row, indicating that GAGCN helps to predict unseen action types. And the results in the second and third rows are very close, indicating that GAGCN performs accurate predictions on unseen action types.

Effect of Balancing Block To demonstrate the effect of balancing block, we set up three contrast experiments against our method (shown in Table. 6): 1. Contrast experiment 1 illustrates the necessity of using gating network both

Human 3.6M-Walking Together							
Model	Motion	80	160	320	400	560	1000
S&TGCN	unseen	10.8	20.7	38.1	42.7	53.1	69.8
GAGCN	unseen	8.9	14.0	26.8	31.1	38.0	51.6
GAGCN	seen	8.8	13.8	26.2	29.9	37.8	50.4

Table 5. Ablation study for the effect of enhancing block. "GAGCN" denotes our proposed model and "S&TGCN" denotes GCN with stable spatial and temporal adjacency matrix, and other experimental settings are the same for both models. "seen" and "unseen" denote that whether the action type (Walking Together) is seen during training or not. The results show that our method can enhance the generalization across unseen action types.

Human 3.6M-average							
milliseconds		80	160	320	400	560	1000
Our method	S_4, T_3	10.1	16.9	32.5	38.5	50.0	72.9
CE1	S_4, T_1	12.5	19.9	38.4	51.3	68.6	93.9
	S_1, T_3	13.1	22.3	40.9	54.1	67.1	91.1
CE2	S_8, T_6	11.4	18.1	33.6	42.5	53.7	76.9
CE3	S_3, T_4	10.3	16.9	33.1	39.2	52.1	75.3

Table 6. Ablation for the effect of balancing block. S and T denote spatial and temporal adjacency matrix, and the subscripts indicate the number of matrices. "CE" denotes the contrast experiment. The best results are shown in bold.

spatially and temporally. The better results of our method indicate that applying gating network in time and space simultaneously helps to model spatio-temporal dependencies more effectively. 2. Contrast experiment 2 shows that more candidate matrices are not always better. We empirically classify the motions in human 3.6M into roughly four categories of similar motions, thus four spatial candidate matrices are used. Too many candidate matrices will increase the complexity of the network and cause under-fitting. 3. Contrast experiment 3 says that the weight of spatio-temporal modeling affects the accuracy of the prediction results indeed, and that is why we balance them by adjusting the number of candidate matrices.

4.4. Qualitative Evaluation

Visualization of Predicted Sequence We visualize the predicted sequence on Human 3.6M and compare them with Ground Truth in Fig. 3. For periodic motions such as "Walking" and "Walking Together", our predictions are almost identical to Ground Truth over the entire time horizons. Meanwhile, for more complex non-periodic motions like "Discussions" and "Posing", our predictions match the Ground Truth well in the short term, and the long term predictions are also quite close to GT despite some acceptable errors on the left leg of "Discussions" and the right arm of "Posing". Non-periodic motion prediction is a more chal-



Figure 3. Visualization of predicted sequences against Ground Truth sequences for 80, 160, 320, 560, 720, 880, 1000ms. We demonstrate the prediction of "Walking", "Walking Together", "Discussions", and "Posing", where the green and purple lines indicate prediction and the red and blue lines indicate the corresponding Ground Truth.

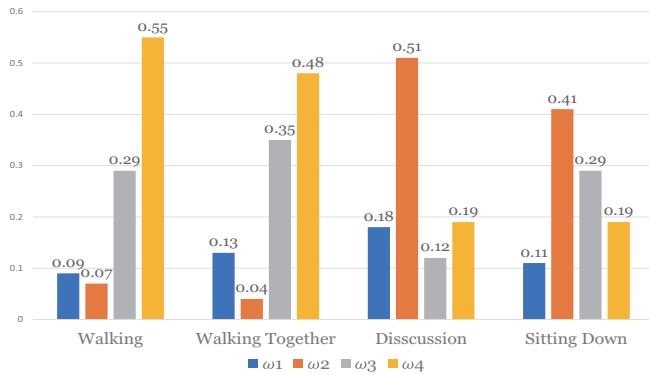


Figure 4. Visualization of average spatial blending coefficients for 4 action types. $\omega_1, \omega_2, \omega_3, \omega_4$ denote the 4 blending coefficients, respectively. Different action types (like "Walking", "Discussions" and "Sitting Down") have different coefficients distribution while the coefficients of similar motions are similarly distributed (like "Walking" and "Walking Together").

lenging problem, especially when the subjects in the testing dataset perform in different ways compared with subjects in the training dataset. The visualization of predicted sequences on AMASS is shown in the supplementary material.

Visualization of Spatial Blending Coefficients Moreover, we randomly select 16 sequences from a single action type to compute the average spatial blending coefficients (visualization of temporal blending coefficients can be found in supplementary material). Then we do the same operation on several action types and visualize them (seeing Fig. 4). We can see that there is a clear difference in the blending coefficients distribution for different action types. The blending coefficients for "Walking Together" are de-

rived from the partial train model in 4.3. Since "Walking Together" and "Walking" are similar periodic motions, their blending coefficients are similarly distributed with higher ω_3 and ω_4 values. That is why our model can achieve accurate prediction results for "Walking Together" without seeing it before. Non-periodic motions like "Discussion" and "Sitting Down" have higher ω_2 values, but their coefficient distributions are very different. Given different inputs, GAGCN can generate the corresponding blending coefficients, which helps to learn the adaptive adjacency matrix for diverse action types. The visualization of adaptive adjacency matrix can be found in the supplementary material.

5. Conclusion and Future Work

In this paper, we propose a novel method called GAGCN to solve motion prediction for multi-action motions. We use the gating network to learn adaptive adjacency matrix by blending candidate adjacency matrices, which effectively enhance the generalization on multi-action motions. Meanwhile, GAGCN can balance the spatio-temporal modeling by adjusting the number of candidate matrices. Combined with the fusion of spatio-temporal features, we can extract the cross-dependency of spatial and temporal relationships to achieve state-of-the-art results on several widely used benchmark datasets. In the future, we will study how to automatically balance the weight for spatio-temporal modeling instead of manually adjusting them and explore a more efficient approach to enhance the generalization of GCN.

Acknowledgements This work was supported by the National Key R&D Program of Science and Technology for Winter Olympics (No.2020YFF0304701) and the National Natural Science Foundation of China (No.61772499).

References

- [1] Shaojie Bai, J Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*, 2018. 5
- [2] Matthew Brand and Aaron Hertzmann. Style machines. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 183–192, 2000. 2
- [3] 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
- [4] 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 *European Conference on Computer Vision*, pages 226–242. Springer, 2020. 2
- [5] 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
- [6] 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, 4
- [7] 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
- [8] 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. 1, 2
- [9] 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. 1, 2
- [10] 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
- [11] Alejandro Hernandez, Jurgen Gall, and Francesc Moreno-Noguer. Human motion prediction via spatio-temporal inpainting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7134–7143, 2019. 1, 2
- [12] 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, 5
- [13] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991. 3
- [14] 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. 1, 2
- [15] Michael I Jordan and Robert A Jacobs. Hierarchical mixtures of experts and the em algorithm. *Neural computation*, 6(2):181–214, 1994. 3
- [16] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. 2
- [17] 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, 5, 6
- [18] 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*, 2021. 1, 2, 4
- [19] 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, 4
- [20] 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. 1, 2, 4
- [21] Hung Yu Ling, Fabio Zinno, George Cheng, and Michiel Van De Panne. Character controllers using motion vaes. *ACM Transactions on Graphics (TOG)*, 39(4):40–1, 2020. 3, 4
- [22] 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. 1, 2, 4
- [23] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5442–5451, 2019. 2, 5
- [24] Wei Mao, Miaomiao Liu, and Mathieu Salzmann. History repeats itself: Human motion prediction via motion attention. In *European Conference on Computer Vision*, pages 474–489. Springer, 2020. 1, 2, 4, 5, 6
- [25] 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, 4, 5, 6

- [26] Wei Mao, Miaomiao Liu, Mathieu Salzmann, and Hongdong Li. Multi-level motion attention for human motion prediction. *International Journal of Computer Vision*, pages 1–23, 2021. [1](#), [2](#), [4](#)
- [27] 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](#), [5](#), [6](#), [7](#)
- [28] 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. [1](#), [2](#), [4](#), [5](#), [6](#)
- [29] Sebastian Starke, He Zhang, Taku Komura, and Jun Saito. Neural state machine for character-scene interactions. *ACM Trans. Graph.*, 38(6):209–1, 2019. [3](#), [4](#)
- [30] Sebastian Starke, Yiwei Zhao, Taku Komura, and Kazi Zaman. Local motion phases for learning multi-contact character movements. *ACM Transactions on Graphics (TOG)*, 39(4):54–1, 2020. [3](#), [4](#)
- [31] Sebastian Starke, Yiwei Zhao, Fabio Zinno, and Taku Komura. Neural animation layering for synthesizing martial arts movements. *ACM Transactions on Graphics (TOG)*, 40(4):1–16, 2021. [3](#), [4](#)
- [32] Yongyi Tang, Lin Ma, Wei Liu, and Wei-Shi Zheng. Long-term human motion prediction by modeling motion context and enhancing motion dynamics. In *IJCAI*, 2018. [1](#), [2](#)
- [33] 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](#), [5](#)
- [34] Borui Wang, Ehsan Adeli, Hsu-kuang Chiu, De-An Huang, and Juan Carlos Niebles. Imitation learning for human pose prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7124–7133, 2019. [1](#), [2](#)
- [35] 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. [4](#)
- [36] Jack M Wang, David J Fleet, and Aaron Hertzmann. Gaussian process dynamical models for human motion. *IEEE transactions on pattern analysis and machine intelligence*, 30(2):283–298, 2007. [2](#)
- [37] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Thirty-second AAAI conference on artificial intelligence*, 2018. [2](#)
- [38] He Zhang, Sebastian Starke, Taku Komura, and Jun Saito. Mode-adaptive neural networks for quadruped motion control. *ACM Transactions on Graphics (TOG)*, 37(4):1–11, 2018. [3](#), [4](#)