

HDBN: A NOVEL HYBRID DUAL-BRANCH NETWORK FOR ROBUST SKELETON-BASED ACTION RECOGNITION

Jinfu Liu^{*1}, Baiqiao Yin^{*1}, Jiaying Lin¹, Jiajun Wen¹, Yue Li¹, Mengyuan Liu^{†2}

¹School of Intelligent Systems Engineering, Sun Yat-sen University

²National Key Laboratory of General Artificial Intelligence, Peking University, Shenzhen Graduate School

* means co-first authors with equal contributions.

Corresponding author is Mengyuan Liu (e-mail: liumengyuan@pku.edu.cn).

ABSTRACT

Skeleton-based action recognition has gained considerable traction thanks to its utilization of succinct and robust skeletal representations. Nonetheless, current methodologies often lean towards utilizing a solitary backbone to model skeleton modality, which can be limited by inherent flaws in the network backbone. To address this and fully leverage the complementary characteristics of various network architectures, we propose a novel Hybrid Dual-Branch Network (HDBN) for robust skeleton-based action recognition, which benefits from the graph convolutional network’s proficiency in handling graph-structured data and the powerful modeling capabilities of Transformers for global information. In detail, our proposed HDBN is divided into two trunk branches: MixGCN and MixFormer. The two branches utilize GCNs and Transformers to model both 2D and 3D skeletal modalities respectively. Our proposed HDBN emerged as one of the top solutions in the Multi-Modal Video Reasoning and Analyzing Competition (MMVRAC) of 2024 ICME Grand Challenge, achieving accuracies of 47.95% and 75.36% on two benchmarks of the UAV-Human dataset by outperforming most existing methods. Our code will be publicly available at: <https://github.com/liujf69/ICMEW2024-Track10>.

Index Terms— Action recognition, graph convolutional network, transformer

1. INTRODUCTION

Human action recognition is an important task in video understanding, which also has applications in research fields including human-robot interaction and security surveillance. Due to the robustness against environmental interference and the effective representation of body movements, the skeleton has become the mainstream modality used in human action recognition. In previous studies [1, 2, 3, 4, 5, 6, 7], skeletons were commonly represented in the form of pseudo-images and topological graphs, while being modeled using classic neural networks simultaneously.

Initially, skeletons were transformed into 2D pseudo-images [1, 2] and feature extraction was performed using convolutional neural networks (CNNs), which effectively utilized the modeling capability of convolutional layers for Euclidean data. However, these methods tended to overlook the connections between skeleton joints. Indeed, the skeleton modality can be perceived as an inherent topological graph, with the graph vertices and edges symbolizing the joints and bones of the human body respectively. Therefore, graph convolutional networks (GCNs) are used to model skeletons in many studies [3, 4, 5, 6, 7, 8, 9, 10]. However, GCNs are often limited by the design of the adjacency matrix, making it challenging to migrate to skeleton inputs with different numbers of joints. Currently, due to the powerful abstraction and modeling capabilities of Transformers for global information, some studies have applied Transformers [11, 12] in action recognition. Similarly, Transformers face the issue of overlooking the design of adjacency matrices for skeletal joints, resulting in recognition accuracy inferior to state-of-the-art GCNs.

In reality, CNNs, GCNs and Transformers each have their own unique advantages and disadvantages when processing skeleton modality, leading to poor performance in fine-grained recognition of certain human actions. A good approach to address the above issues is to simultaneously use multiple network structures, achieving robust skeleton-based action recognition through the complementarity between networks. Although there have been existing solutions [13, 14, 15, 16] that simultaneously use CNN and GCN networks, they often apply to scenarios where both RGB image modality and skeleton modality are used as inputs.

Therefore, in order to better handle scenarios with only skeleton inputs and leverage the complementarity between different backbone networks, we propose a novel Hybrid Dual-Branch Network (HDBN) for robust skeleton-based action recognition. Specifically, our HDBN consists of two trunk branches called MixGCN branch and MixFormer branch, which is shown in Fig.1. In the MixGCN branch, we utilize GCNs to separately process 2D and 3D skeleton inputs, and employ a late-fusion strategy to aggregate the clas-

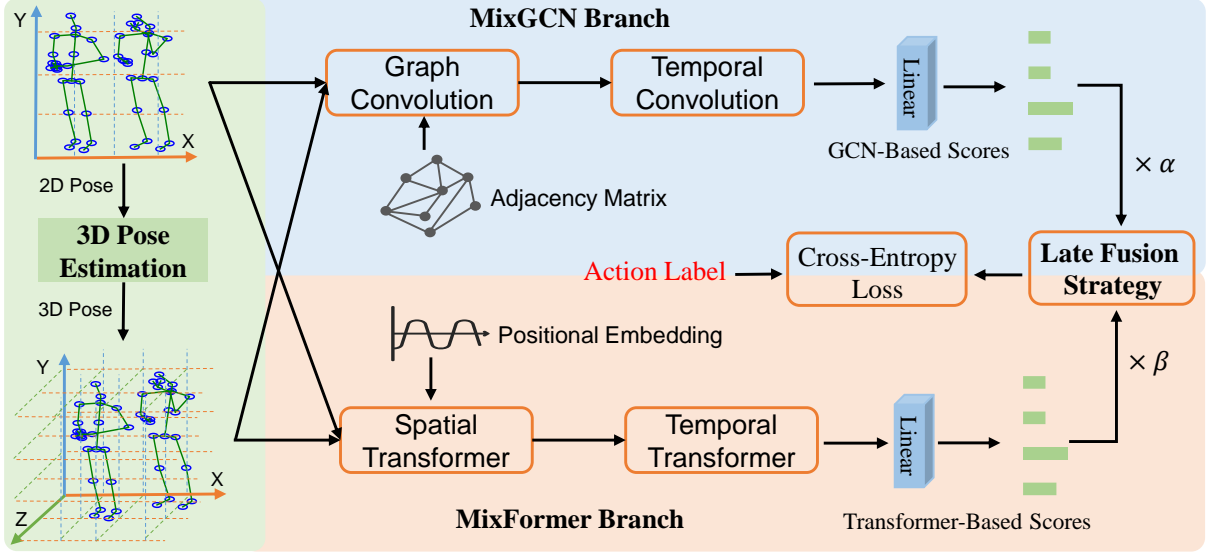


Fig. 1. Framework of our proposed Hybrid Dual-branch Network.

sification results from different skeleton input. Similarly, in the Mixformer branch, we employ Transformer to model the skeleton inputs, fully harnessing Transformer’s abstraction capability for global information. By leveraging the proposed HDBN, we effectively integrate GCNs and TransFormers to achieve better human action recognition.

Our contributions are summarized as follows:

- We advocate for utilizing different network structures to achieve robust skeleton-based action recognition, fully leveraging the structural complementarity between different backbones.
- We propose a new dual-branch framework called Hybrid Dual-Branch Network (HDBN), which effectively combines GCNs and Transformers. In detail, the skeleton data are inputted to GCN and Transformer backbones for modeling high-level features, which will be effectively combined through a late fusion strategy for more robust skeleton-based action recognition.
- Extensive experiments on benchmark UAV-Human dataset verify the effectiveness of our HDBN. In this large-scale action recognition datasets, our model outperforms most existing action recognition methods.

2. RELATED WORK

2.1. CNN-Based and RNN-Based Methods

In the realm of skeleton-based action recognition, deep neural network (DNN) approaches predominantly leverage Recurrent Neural Networks (RNNs) and Convolutional Neural Networks (CNNs). RNNs typically incorporate outputs from previous time steps as inputs for the current time step, forming a recursive connection structure that excels in handling sequen-

tial data. On the other hand, CNNs harness the unique advantages inherent in processing Euclidean data by often transforming skeleton sequences into 2D pseudo-images. Several prior studies have harnessed the power of RNNs [17, 18, 19] and CNNs [1, 2] to attain impactful outcomes in action recognition based on skeletal data. However, neither RNN-based nor CNN-based methods have been able to fully capture the intricate dependencies among human joints, thereby limiting advancements in recognition accuracy.

2.2. GCN-Based Methods

GCNs represent a class of networks adept at efficiently handling structured data, particularly graph data. In the domain of skeleton-based action recognition, researchers [3, 4, 5, 6, 7] have harnessed the power of GCNs with promising results. For example, Yan et al. [3] introduced the Spatial-Temporal Graph Convolutional Network (ST-GCN), conceptualizing joints as vertices and predefining edges to capture dependencies informed by prior knowledge. Liu et al. [7] introduced the Temporal Decoupling Graph Convolutional Network (TD-GCN), innovatively leveraging a temporal-dependent adjacency matrix to extract high-level features from skeleton data with remarkable efficacy. Nevertheless, the efficacy of GCNs is frequently constrained by the fixed design of the adjacency matrix, posing challenges when adapting to skeleton inputs with varying numbers of joints.

2.3. Transformer-Based Methods

Transformers inherently excel at capturing and processing global information, which is helpful for improving the classification capabilities in skeleton-based action recognition. Nu-

merous research endeavors have been dedicated to optimizing transformers specifically for action recognition tasks. For instance, to address problems that previous methods are hard to capture interactive relations, Wen et al. [11] proposed an Interactive Spatiotemporal Token Attention Network (ISTANet) based on the Transformer, which simultaneously model spatial, temporal, and interactive relations. To address problems that current skeleton transformer depends on the self-attention mechanism of the complete channel of the global joint, ignoring the highly discriminative differential correlation within the channel, Xin et al. [12] proposed the Skeleton MixFormer, which effectively represent the physical correlations and temporal interactivity of the compact skeleton data. Nevertheless, Transformer-based approaches often encounter challenges in adequately designing adjacency matrices for skeletal joints, leading to recognition accuracy that falls short of cutting-edge Graph Convolutional Networks.

2.4. Hybrid Dual-Branch Methods

Due to the inherent limitations of a single backbone network, it struggles to fully extract and model information from the skeleton modality. Consequently, some research efforts [13, 14, 15, 16] are dedicated to employing multiple backbone networks (such as the combination of CNNs and GCNs) to achieve more robust action recognition. For instance, Liu et al. [15] utilize both CNN and GCN backbones to process human parsing and human pose modalities separately, which implement a late fusion strategy to combine features from both modalities. Das et al. [13] propose the Video-Pose Network (VPN), which employs both CNN and GCN to model RGB and skeletal modalities, enabling the learning of enhanced spatiotemporal features. However, the aforementioned hybrid dual-branch methods tend to rely on both CNNs and GCNs and require input from two different modalities, which is not conducive to modeling situations where only skeleton is available. Distinct from all the methods mentioned above, our proposed HDBN is a novel dual-branch scheme that combines the state-of-the-art GCN and Transformer backbones for robust skeleton-based action recognition.

3. METHOD

3.1. Multimodal Data Processing

In our proposed Hybrid Dual-branch Network (HDBN), we use robust skeleton modality for human action recognition. Conceptually, the sequence of skeletal structure forms a natural topological graph, wherein joints represent graph vertices and bones serve as edges. The graph can be defined as $G = \{V, L\}$, where $V = \{v_1, v_2, \dots, v_N\}$ and $L = \{l_1, l_2, \dots, l_M\}$ represent N joints and M bones of human body. The joint v_i is defined as $\{x_i, y_i\}$ in 2D pose data, where x_i and y_i locate v_i in two-dimensional Euclidean space.

Here we define four skeleton modalities, namely joint (J), bone (B), joint motion (JM) and bone motion (BM). Given two joints data $v_i = \{x_i, y_i\}$ and $v_j = \{x_j, y_j\}$, a bone data of the skeleton is defined as $e_{v_i, v_j} = (x_i - x_j, y_i - y_j)$. Given two joints data $v_{ti}, v_{(t+1)i}$ from two consecutive frames, the data of joint motion is defined as $m_{ti} = v_{(t+1)i} - v_{ti}$. Similarly, given two bones data $e_{v_{(t+1)i}, v_{(t+1)j}}, e_{v_{ti}, v_{tj}}$ from two consecutive frames, the data of bone motion is defined as $m_{v_{ti}, v_{tj}} = e_{v_{(t+1)i}, v_{(t+1)j}} - e_{v_{ti}, v_{tj}}$.

Typically, 3D skeleton data provides more information compared to 2D skeleton data, which can compensate for the limitations of 2D data. Therefore, our HDBN will utilize both 2D and 3D skeleton data simultaneously. Specifically, we will employ a 3D pose estimation encoder [24] to derive 3D skeleton data from 2D pose data, which can be formulated as:

$$P_{3D} = \text{Encoder}(x_i, y_i) \quad (1)$$

3.2. MixGCN Backbone

Due to its natural advantages in processing graph-structured data, the MixGCN branch in our proposed HDBN utilizes GCNs to model skeleton data. Actually the Graph Convolutional Networks consist of two key parts: graph convolution module and temporal convolution module. The normal graph convolution can be formulated as:

$$H^{l+1} = \sigma \left(D^{-\frac{1}{2}} A D^{-\frac{1}{2}} H^l W^l \right), \quad (2)$$

where H^l is the features at layer l and σ is the activation function. $D \in R^{N \times N}$ is the degree matrix of N joints and W^l is the learnable parameter of the l -th layer. A is the adjacency matrix representing joint connections. Generally, the A can be generated by using static and dynamic ways. In the MixGCN branch, we employ the backbone with three types of dynamic adjacency matrices (i.e. TD-GCN [7], CTR-GCN [6] and MST-GCN [3]) to handle both 2D and 3D skeletal data, fully leveraging the complementarity among different GCNs.

3.3. MixFormer Backbone

Since Transformers possess powerful modeling and abstraction capabilities for global information, the MixFormer branch in our HDBN utilizes Transformer as the core backbone to model the skeleton data. Generally, a standard Transformer consists of modules such as positional encoding, self-attention, and feedforward neural networks. We can utilize the Transformer architecture to encode features of input modality and combine specific downstream layers to handle various downstream tasks. A simplified representation of the encoded features is as follows:

$$F = \text{Norm}[X + \text{Attention}(XW_Q, XW_K, XW_V)W_O] \quad (3)$$

Where W_Q , W_K , W_V , and W_O represent the query, key, value and output weight matrices respectively.

Table 1. Accuracy comparison with state-of-the-art methods on UAV-Human dataset. Symbol \star represents the result reproduced in our work.

Method	Source	CSv1 (%)	CSv2 (%)
DGNN [20]	CVPR'19	29.90	-
ST-GCN [3]	AAAI'18	30.25	56.14
2s-AGCN [4]	CVPR'19	34.84	66.68
HARD-Net [21]	ECCV'20	36.97	-
Shift-GCN [5]	CVPR'20	37.98	67.04
FR-AGCN [22]	Neuroc'22	43.98	69.50
CTR-GCN \star [6]	ICCV'21	45.60	72.50
ACFL-CTR [23]	ACMMM'22	45.30	-
TD-GCN \star [7]	TMM'24	45.43	72.86
Skeleton MixFormer \star [12]	ACMMM'23	47.23	73.47
MixGCN (Ours)	-	46.73	74.06
MixFormer (Ours)	-	47.23	73.47
HDBN (Ours)	-	47.95	75.36




Fig. 2. The confusion matrix of UAV-Human. The more yellow the squares on the diagonal, the more accurate the recognition. (a) UAV-Human dataset on the benchmark of CSv1. (b) UAV-Human dataset on the benchmark of CSv2

Table 2. Accuracy of different 2D skeleton modalities on UAV-Human dataset when using CTR-GCN as the backbone.

Modality				UAV-Human (%)	
J	B	JM	BM	CSv1	CSv2
✓				43.52	69.00
	✓			43.32	68.68
		✓		36.25	57.93
			✓	35.86	58.45
✓	✓	✓	✓	45.60	72.50

Table 3. Accuracy of different 3D skeleton modalities on UAV-Human dataset when using CTR-GCN as the backbone.

Modality				UAV-Human (%)	
J	B	JM	BM	CSv1	CSv2
✓				35.14	64.60
	✓			35.66	63.25
		✓		31.08	55.80
			✓	30.89	54.67
✓	✓	✓	✓	38.85	67.87

3.4. Ensemble

In our proposed HDBM, we utilize a late-fusion strategy to combine the classification scores from the MixGCN branch and the MixFormer branch. In detail, the late-fusion strategy employs a set of weight parameters to linearly combine the classification scores from the two branches, which is formulated as:

$$\mathbf{S} = \alpha \cdot \mathcal{M}_{gcn}(\mathbf{P}) + \beta \cdot \mathcal{M}_{former}(\mathbf{P}), \quad (4)$$

Where \mathcal{M}_{gcn} and \mathcal{M}_{former} represent the MixGCN backbone

and MixFormer backbone respectively, while the α and β represent the weights for linear fusion. The final result \mathbf{S} that integrates two branches will be processed by a softmax layer for final robust action recognition.

In our proposed HDBN, we use Cross-Entropy (CE) loss to train the entire network:

$$\mathcal{L}_{cls} = - \sum_i^N \mathcal{Y}^i \log \mathbf{S}^i, \quad (5)$$

where N is the number of samples in a batch and \mathcal{Y}^i is the one-hot presentation of the true label about action sample i .

Table 4. Accuracy of different 2D skeleton modalities on UAV-Human dataset when using TD-GCN as the backbone.

Modality				UAV-Human (%)	
J	B	JM	BM	CSv1	CSv2
✓				43.21	69.50
	✓			43.33	69.30
		✓		35.74	57.74
			✓	35.56	55.14
✓	✓	✓	✓	45.43	72.86

Table 5. Accuracy of different 2D skeleton modalities on UAV-Human dataset when using MST-GCN as the backbone.

Modality				UAV-Human (%)	
J	B	JM	BM	CSv1	CSv2
✓				41.48	67.48
	✓			41.57	67.30
		✓		33.82	54.43
			✓	34.74	52.13
✓	✓	✓	✓	43.78	70.95

The \mathcal{S}^i is the true classification score of action i output by the linear layer.

4. EXPERIMENTS

4.1. Datasets and Implementation Details

Datasets: UAV-Human [25] is a dataset designed for action recognition, comprising 22,476 video clips across 155 distinct classes. Crafted by a UAV traversing diverse urban and rural landscapes day and night, this dataset captures the essence of human activity. Curated from 119 unique individuals engaging in 155 varied activities across 45 environmental settings, it offers a rich description of human behavior. The original paper suggests two benchmark scenarios (CSv1 and CSv2) for evaluation: utilizing 89 subjects for training and 30 for testing, ensuring robustness and reliability in analysis.

Implementation Details: All experiments are conducted on the GeForce RTX 3090 GPUs. In the MixGCN branch, we employ TD-GCN [7], CTR-GCN [6] and MST-GCN [3] as the foundational components to derive classification scores based on GCN. Model training is executed using SGD over 65 epochs with a batch size of 64. Initially, the learning rate is established at 0.1 and undergoes decay by a factor of 0.1 at epochs 35 and 55, ensuring optimal convergence and performance refinement. Within the MixFormer branch, we integrate Skeleton MixFormer [12] as the core framework to derive classification scores based on Transformers. Employing SGD, we train the model across 90 epochs, utilizing a batch size of 128. The initial learning rate is defined at 0.02, ensuring a robust foundation for effective training and refinement.

Table 6. Accuracy of different 2D skeleton modalities on UAV-Human dataset when using Skeleton MixFormer as the backbone. The definitions of k2 and k2M can be found in [8].

Modality						UAV-Human (%)	
J	B	JM	BM	k2	k2M	CSv1	CSv2
✓						41.43	66.03
	✓					37.40	64.89
		✓				33.41	54.58
			✓			30.24	52.95
				✓		39.21	65.56
					✓	32.60	55.01
✓	✓	✓	✓	✓	✓	47.23	73.47

4.2. Comparisons with the State-of-the-Art

In Table 1, we report the Top-1 accuracy of our HDBN and compare with the existing methods on the UAV-Human dataset. In this action recognition dataset, our HDBN outperforms most existing methods under the recommended evaluation benchmarks. In the UAV-Human dataset, the classification accuracy is 47.95% and 75.36% on the benchmark of CSv1 and CSv2 respectively, which outperforms CTR-GCN by 2.35% and 2.86%. We also present the confusion matrices of our proposed HDBN in Fig. 2.

4.3. Ablation Studies

In Table 2, when using CTR-GCN as the backbone in the MixGCN branch, we present the recognition accuracy of four skeleton modalities in UAV-Human dataset. Meanwhile, in Table 3, we demonstrate the accuracy of estimating 3D poses from 2D poses using MotionBert and utilizing these 3D poses for action recognition. In addition, in Table 4 and Table 5, we also report the recognition accuracy of using TD-GCN and MST-GCN to model 2D skeleton modalities. Finally, in Table 6, we present the recognition accuracy when using Skeleton MixFormer as the backbone in the MixFormer branch.

5. CONCLUSIONS

We propose a novel dual-branch framework called Hybrid Dual-branch Network (HDBN) for robust skeleton-based action recognition, which introduces the Transformers and GCNs to model 2D and 3D skeleton data. Different from previous methods using single backbone or using different backbone to process various modalities, our HDBN is to leverage both Transformers and GCNs with the aim of robust skeleton-based action recognition. The effectiveness of HDBN is verified on the UAV-Human dataset, where our HDBN outperforms most existing methods.

6. ACKNOWLEDGEMENT

This work was supported by National Natural Science Foundation of China (No. 62203476), Natural Science Foundation of Shenzhen (No. JCYJ20230807120801002).

7. REFERENCES

- [1] Mengyuan Liu, Hong Liu, and Chen Chen, “Enhanced skeleton visualization for view invariant human action recognition,” *Pattern Recognition*, vol. 68, pp. 346–362, 2017.
- [2] Qihong Ke, Mohammed Bennamoun, Senjian An, Ferdous Sohel, and Farid Boussaid, “Learning clip representations for skeleton-based 3d action recognition,” *IEEE Transactions on Image Processing*, vol. 27, no. 6, pp. 2842–2855, 2018.
- [3] Sijie Yan, Yuanjun Xiong, and Dahua Lin, “Spatial temporal graph convolutional networks for skeleton-based action recognition,” in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- [4] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu, “Two-stream adaptive graph convolutional networks for skeleton-based action recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12026–12035.
- [5] Ke Cheng, Yifan Zhang, Xiangyu He, Weihang Chen, Jian Cheng, and Hanqing Lu, “Skeleton-based action recognition with shift graph convolutional network,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [6] Yuxin Chen, Ziqi Zhang, Chunfeng Yuan, Bing Li, Ying Deng, and Weiming Hu, “Channel-wise topology refinement graph convolution for skeleton-based action recognition,” in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [7] Jinfu Liu, Xinshun Wang, Can Wang, Yuan Gao, and Mengyuan Liu, “Temporal decoupling graph convolutional network for skeleton-based gesture recognition,” *IEEE Transactions on Multimedia*, pp. 1–13, 2023.
- [8] Hyung-gun Chi, Myoung Hoon Ha, Seunggeun Chi, Sang Wan Lee, Qixing Huang, and Karthik Ramani, “Infogcn: Representation learning for human skeleton-based action recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, New Orleans, USA, 2022, pp. 20186–20196, IEEE.
- [9] Ruwen Bai, Min Li, Bo Meng, Fengfa Li, Miao Jiang, Junxing Ren, and Degang Sun, “Hierarchical graph convolutional skeleton transformer for action recognition,” in *Proceedings of the IEEE International Conference on Multimedia and Expo (ICME)*, IEEE, 2022, pp. 01–06.
- [10] Mengyuan Liu, Fanyang Meng, Chen Chen, and Songtao Wu, “Novel motion patterns matter for practical skeleton-based action recognition,” in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2023, pp. 1701–1709.
- [11] Yuhang Wen, Zixuan Tang, Yunsheng Pang, Beichen Ding, and Mengyuan Liu, “Interactive spatiotemporal token attention network for skeleton-based general interactive action recognition,” in *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 7886–7892.
- [12] Wentian Xin, Qiguang Miao, Yi Liu, Ruyi Liu, Chi-Man Pun, and Cheng Shi, “Skeleton mixformer: Multivariate topology representation for skeleton-based action recognition,” in *Proceedings of the ACM International Conference on Multimedia (ACM MM)*, 2023, pp. 2211–2220.
- [13] Srijan Das, Saurav Sharma, Rui Dai, Francois Bremond, and Monique Thonnat, “Vpn: Learning video-pose embedding for activities of daily living,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, Glasgow, USA, 2020, pp. 72–90, Springer.
- [14] Runwei Ding, Yuhang Wen, Jinfu Liu, Nan Dai, Fanyang Meng, and Mengyuan Liu, “Integrating human parsing and pose network for human action recognition,” in *Proceedings of the CAAI International Conference on Artificial Intelligence (CICAI)*, Fuzhou, China, 2023, pp. 182–194, CAAI.
- [15] Jinfu Liu, Runwei Ding, Yuhang Wen, Nan Dai, Fanyang Meng, Shen Zhao, and Mengyuan Liu, “Explore human parsing modality for action recognition,” *CAAI Transactions on Intelligence Technology*, 2024.
- [16] Bruce X.B. Yu, Yan Liu, Xiang Zhang, Sheng-hua Zhong, and Keith C.C. Chan, “Mmnet: A model-based multimodal network for human action recognition in rgb-d videos,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 3, pp. 3522–3538, 2023.
- [17] Jun Liu, Gang Wang, Ling-Yu Duan, Kamila Abdiyeva, and Alex C. Kot, “Skeleton-based human action recognition with global context-aware attention lstm networks,” *IEEE Transactions on Image Processing*, vol. 27, no. 4, pp. 1586–1599, 2018.
- [18] Sijie Song, Cuiling Lan, Junliang Xing, Wenjun Zeng, and Jiaying Liu, “Spatio-temporal attention-based lstm networks for 3d action recognition and detection,” *IEEE Transactions on Image Processing*, vol. 27, no. 7, pp. 3459–3471, 2018.
- [19] Jun Liu, Amir Shahroudy, Dong Xu, and Gang Wang, “Spatio-temporal lstm with trust gates for 3d human action recognition,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016, pp. 816–833.
- [20] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu, “Skeleton-based action recognition with directed graph neural networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 7912–7921.
- [21] Tianjiao Li, Jun Liu, Wei Zhang, and Lingyu Duan, “Hard-net: Hardness-aware discrimination network for 3d early activity prediction,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020, pp. 420–436.
- [22] Zesheng Hu, Zihao Pan, Qiang Wang, Lei Yu, and Shumin Fei, “Forward-reverse adaptive graph convolutional networks for skeleton-based action recognition,” *Neurocomputing*, vol. 492, pp. 624–636, 2022.
- [23] Xuanhan Wang, Yan Dai, Lianli Gao, and Jingkuan Song, “Skeleton-based action recognition via adaptive cross-form learning,” in *Proceedings of the ACM International Conference on Multimedia (ACM MM)*, 2022, pp. 1670–1678.
- [24] Wentao Zhu, Xiaoxuan Ma, Zhaoyang Liu, Libin Liu, Wayne Wu, and Yizhou Wang, “Motionbert: A unified perspective on learning human motion representations,” in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2023, pp. 15085–15099.
- [25] Tianjiao Li, Jun Liu, Wei Zhang, Yun Ni, Wenqian Wang, and Zhiheng Li, “UAV-Human: A Large Benchmark for Human Behavior Understanding With Unmanned Aerial Vehicles,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 16266–16275.