

Self-Supervised Action Representation Learning from Partial Spatio-Temporal Skeleton Sequences

Yujie Zhou^{1,4}, Haodong Duan³, Anyi Rao³, Bing Su^{1,2*}, Jiaqi Wang⁴

¹Gaoling School of Artificial Intelligence, Renmin University of China

²Beijing Key Laboratory of Big Data Management and Analysis Methods

³Chinese University of Hong Kong

⁴Shanghai AI Laboratory

yujiezhou@ruc.edu.cn, anyirao@link.cuhk.edu.hk, {dhd.efz, subingats, wjqdev}@gmail.com

Abstract

Self-supervised learning has demonstrated remarkable capability in representation learning for skeleton-based action recognition. Existing methods mainly focus on applying global data augmentation to generate different views of the skeleton sequence for contrastive learning. However, due to the rich action clues in the skeleton sequences, existing methods may only take a global perspective to learn to discriminate different skeletons without thoroughly leveraging the local relationship between different skeleton joints and video frames, which is essential for real-world applications. In this work, we propose a **Partial Spatio-Temporal Learning (PSTL)** framework to exploit the local relationship from a partial skeleton sequences built by a unique spatio-temporal masking strategy. Specifically, we construct a negative-sample-free triplet stream structure that is composed of an anchor stream without any masking, a spatial masking stream with Central Spatial Masking (CSM), and a temporal masking stream with Motion Attention Temporal Masking (MATM). The feature cross-correlation matrix is measured between the anchor stream and the other two masking streams, respectively. (1) Central Spatial Masking discards selected joints from the feature calculation process, where the joints with a higher degree of centrality have a higher possibility of being selected. (2) Motion Attention Temporal Masking leverages the motion of action and remove frames that move faster with a higher possibility. Our method achieves state-of-the-art performance on NTURGB+D 60, NTURGB+D 120 and PKU-MMD under various downstream tasks. Furthermore, to simulate the real-world scenarios, a practical evaluation is performed where some skeleton joints are lost in downstream tasks. In contrast to previous methods that suffer from large performance drops, our PSTL can still achieve remarkable results under this challenging setting, validating the robustness of our method. Our code is available at <https://github.com/YujieOuO/PSTL.git>.

Introduction

In recent years, human action recognition has been extensively studied for various applications in the real world, such as security, human-robot interaction, and virtual reality. With the development of sensor technology (Zhang

2012; Shotton et al. 2011) and human pose estimation methods (Cao et al. 2017; Hua et al. 2022), skeleton data becomes easier to obtain. Its robustness against the variation of appearance and backgrounds has attracted many researchers to explore fully supervised methods for action recognition, such as applying CNN (Ke et al. 2017), RNN (Du, Wang, and Wang 2015) or GCN (Yan, Xiong, and Lin 2018; Shi et al. 2019) on skeleton data.

However, fully supervised learning relies on abundant manually annotated labels, which is time-consuming and expensive. Therefore, self-supervised frameworks, *e.g.*, MoCo (He et al. 2020), SimCLR (Chen et al. 2020), Barlow Twins (Zbontar et al. 2021), and MAE (He et al. 2022) have received a lot of attention. For self-supervised learning on skeleton data, early methods also apply pretext tasks (Lin et al. 2020; Su, Liu, and Shlizerman 2020; Zheng et al. 2018). While nowadays, most researchers focus on contrastive learning and apply MoCo structure to the skeleton sequences (Thoker, Doughty, and Snoek 2021; Li et al. 2021; Rao et al. 2021; Guo et al. 2022).

The key of current contrastive learning-based methods is finding effective global data augmentation and using them to create various views of the skeleton carefully. However, due to the rich action clues stored in the skeleton sequence, this paradigm can readily discriminate different views from a global perspective without thoroughly leveraging the local relationship between different skeleton joints and video frames. As a result, it limits the performance and robustness of the model in real-world scenarios. Moreover, a large batch size or a memory bank is necessary for contrastive learning, which is unsuitable for a small amount of skeleton data. To evaluate the robustness of the current methods in real-world scenarios (such as video surveillance), we formulate a new evaluating setting where only a partial skeleton is available for downstream tasks to simulate that the skeleton data captured by the camera are shaded. In this challenging setting, previous works suffer severe performance drops. We conjecture that if the local spatial and temporal relationship can be explored and captured, it will boost the representation ability and robustness of the model on downstream tasks.

Therefore, we propose a Partial Spatio-Temporal Learning (PSTL) framework to exploit the local relationship from a partial skeleton sequence built by a unique spatio-temporal masking strategy. Specifically, we construct a negative-

*Corresponding author

sample-free triplet stream structure that is composed of an anchor stream, a spatial masking stream with Central Spatial Masking (CSM), and a temporal masking stream with Motion Attention Temporal Masking (MATM).

In each stream, an ordinary augmentation is first applied to improve the diversity of the input samples. (1) In the spatial masking stream, Central Spatial Masking discards selected joints from the feature calculation process. Notably, it is not simply set the selected joint position to zeros, which is unreasonable for skeleton data. Considering that the topology of the human skeleton is a predefined graph, the joints with a higher degree of centrality get stronger connectivity. By filtering out such joints with higher probabilities, the encoder can focus more on these skeletons that are less explored. (2) In the temporal masking stream, considering the frames that change quickly always contain more semantic information about the actions, we propose a Motion Attention Temporal Masking strategy, where frames moves quickly are more likely to be masked. (3) In the anchor stream, the masking policy is not introduced to keep the original semantic information. A shared encoder extracts the triple stream features simultaneously. Inspired by the loss of Barlow Twins (Zbontar et al. 2021), we adopt the cross-correlation matrix to capture the correlation between the anchor stream and the other two masking streams, respectively.

As a negative-sample-free self-supervised method, PSTL encourages the empirical cross-correlation matrix between learned representations to be an identity matrix, and minimizes the redundancy within the representation vector to alleviate feature collapse. In this way, the spatial and temporal relationship can be well captured and the joint robustness can be boosted for downstream tasks. Extensive results on three large benchmark demonstrate the effectiveness of our proposed masking strategy. Our contributions:

- We propose Partial Spatio-Temporal skeleton representation Learning (PSTL), a negative-sample-free self-supervised framework for skeleton-based action representation.
- We are the first to introduce the masking strategy to skeleton-based self-supervised learning and using Central Spatial Masking (CSM) strategy during pre-training to improve the joint robustness of learned representation for downstream tasks.
- A unique Motion Attention Temporal Masking strategy (MATM) is proposed to use the motion of skeleton data to guide temporal masking modules to find and mask key frames of the sequences.
- Extensive experiment results on NTU-60, NTU120, and PKU-MMD datasets show the proposed PSTL achieves the state-of-art performance under diverse downstream tasks, such as linear evaluation, finetune evaluation, and semi-supervised evaluation protocol.

Related Work

Skeleton-based action recognition. To tackle skeleton-based action recognition, early methods mainly focus on hand-crafted features (Vemulapalli, Arrate, and Chellappa

2014; Xia, Chen, and Aggarwal 2012). With the development of the deep neural network, skeleton data is processed as pseudo-images with CNNs (Ke et al. 2017; Kim and Reiter 2017). Others structure it as a long-term sequence and feed it into RNNs (Du, Wang, and Wang 2015; Zhang et al. 2017). Later on, ST-GCN (Yan, Xiong, and Lin 2018) structures the skeleton data as a predefined spatial graph and utilize GCN to integrate the skeleton joint information. After that, many GCN-based methods are proposed, which contain more data stream or add attention mechanisms (Shi et al. 2019; Song et al. 2020; Chen et al. 2021). In this paper, we also utilize the widely used ST-GCN as the backbone.

Self-supervised representation learning. Self-supervised learning aims to acquire a high-level representation from a large number of unlabeled samples. Early methods rely on pretext tasks to generate supervision, such as jigsaw puzzles (Noroozi and Favaro 2016; Wei et al. 2019), colorizing (Larsson, Maire, and Shakhnarovich 2016; Zhang, Isola, and Efros 2016), or rotation prediction (Gidaris, Singh, and Komodakis 2018). Then, discriminative methods like MoCo (He et al. 2020) and SimCLR (Chen et al. 2020) achieve remarkable performance, which are comparable to supervised learning. By distinguishing the positive pairs from the negative ones, the backbone network can acquire the representation with rich discriminant information. But it requires a large batch size or memory bank to store the negative samples. To tackle this issue, negative-sample-free methods are introduced such as BYOL (Grill et al. 2020), SimSiam (Chen and He 2021), and Barlow Twins (Zbontar et al. 2021). They do not require explicit construction of negative sample pairs, but most of them utilize an asymmetric network to avoid feature collapse. Specially, Barlow Twins avoids the reliance on the symmetry-breaking network but minimizes the redundancy within the representation vector to alleviate collapse. More recently, mask auto-encoder (He et al. 2022) is introduced to learn effective visual representations from local complementary information.

Self-supervised skeleton-based action recognition. To apply self-supervised learning to skeleton data, LongT GAN (Zheng et al. 2018) uses a recurrent encoder-decoder GAN to reconstruct the input sequence. Based on this, Predict&Cluster (Su, Liu, and Shlizerman 2020) proposes a decoder to improve the representation ability of the encoder. MS2L (Lin et al. 2020) proposes a multi-task self-supervised learning framework, including motion prediction and jigsaw puzzle. AS-CAL (Rao et al. 2021) uses momentum LSTM to regularize the feature space together with several skeleton augmentation strategy. ISC (Thoker, Doughty, and Snoek 2021) proposes inter-skeleton contrastive method that combines sequence-based and graph-based representation. CrossCLR (Li et al. 2021) introduces SkeletonCLR that applies memory bank to store the negative samples and a cross-view knowledge mining strategy to capture more comprehensive representation. AimCLR (Guo et al. 2022) aims to utilize plenty of augmentation to introduce movement patterns and force the encoder to learn more general representations. However, these contrastive-based methods heavily rely on strong data augmentation strategies, but the redundancy within the spatial joints and temporal frames have

not been considered, which can improve the robustness of 3D visual representation and essential for downstream task. Hence, we introduce PSTL and propose a spatial-temporal mask strategy to find a more general representation.

Method

As shown in Fig 2, Partial Spatio-Temporal skeleton representation Learning (PSTL) takes a triple stream architecture, which is based on our Skeleton Barlow Twins (SkeletonBT), and applies Central Spatial Masking (CSM) on its spatial masking stream and Motion Attention Temporal Masking (MATM) on its temporal masking stream together with an extra anchor stream to keep the original semantic information. Specifically, 1) **SkeletonBT** is a basic negative-sample-free double stream network with ordinary augmentation for skeleton self-supervised learning. 2) **CSM** first adjusts the computing mode of the GCN backbone to ensure the information of masked joints out of the feature calculation process. Then, the joints with higher degree centrality are assigned higher probability to be masked. 3) **MATM** calculates the corresponding motion value for each frame of the action, which serves as the attention weight to select the masked frames.

Skeleton Barlow Twins

Inspired by the recent advance on negative-sample-free self-supervised methods (Zbontar et al. 2021) that avoid negative sample and do not need a large batch size/memory bank, we develop Skeleton Barlow Twins (SkeletonBT) to learn the skeleton-based action representation. It is a symmetric network, where a shared encoder is used to generate features.

Specifically, a 3D human skeleton sequence can be denoted as $s \in \mathbb{R}^{C \times T \times V}$, which has T frames and V joints. C is the channel dimension representing the 3D position. As shown in Fig. 1, the input skeleton s is first applied with an ordinary augmentation \mathcal{T} to get diverse views x and x' . The encoder f is used to extract features: $h = f(x; \theta)$ and $h' = f(x'; \theta)$, where $h, h' \in \mathbb{R}^{c_h}$. Then a projector g maps each feature to the higher dimension space and acquire the embeddings: $z = g(h)$ and $z' = g(h')$, where $z, z' \in \mathbb{R}^{c_z}$. At last, by encouraging the empirical cross-correlation matrix between embeddings of the distorted variations be an identity matrix, the encoder can capture the relationship between two streams. The loss can be formulated as:

$$\mathcal{L} = \sum_i (1 - C_{ii})^2 + \lambda \sum_i \sum_{j \neq i} C_{ij}^2. \quad (1)$$

Here C is the cross-correlation matrix computed between embeddings z and z' along the batch dimension b :

$$C_{ij} = \frac{\sum_b z_{b,i} z'_{b,j}}{\sqrt{\sum_b (z_{b,i})^2} \sqrt{\sum_b (z'_{b,j})^2}}, \quad (2)$$

where i, j represent the embedding dimension.

Specifically, in Eq. (1) the first term encourages the diagonal elements of C to 1, which makes the embedding invariant to the variety augmentation applied. The second term




Figure 1: The framework of Skeleton Barlow Twins.

equates the off-diagonal elements of C to 0, helping the different embedding components to be decoupled with each other, which can minimize the redundancy within the representation. λ serves as a trade-off parameter to balance two terms.

Central Spatial Masking

To boost the joint robustness of learned representation, we propose a Central Spatial Masking (CSM) module that encourages partial skeleton data to generate similar features as the whole skeleton data. The relationship between masked and unmasked joints can be captured by the encoder, which is essential for the hard downstream task.

As discussed above, it is critical to construct partial skeleton data for the encoder. Early methods (Rao et al. 2021) try to directly set the values of masked joints to zero. But it is unreasonable for skeleton data since the joint semantic information is the 3D position that is changed. For CSM, we first adjust the computing mode of the backbone ST-GCN (Yan, Xiong, and Lin 2018). Specifically, for a skeleton sequence s , the masked joints are abandoned and then drop the corresponding rows and columns in the adjacency matrix, forcing masked joints out of the calculation process and the learned feature is only generated by the unmasked ones.

Next, we take the advantage of the degree centrality of the human skeleton graph topology, and observe that the joints with more degrees can acquire richer neighborhood information. Thus, we assign higher probabilities to mask such connective joints which help the encoder capture the relationship between a wider range. As shown in the green box of Fig 2, the human skeleton joint is $V_i, i \in (1, 2, \dots, n)$, where n is the total number of joints. Note that there are four types of joints. The degree of light gray joints located at the margin of the graph is 1. The degree of dark gray joints is 2, which are more connective and their number is the majority in the total graph. The blue one and red one hold degree of 3 and 4, which are the center of the upper half and lower half of the body. Masking such center joints can help the encoder to capture the relationship between a large range of joint information. To this end, our CSM strategy first calculates the degree d_i of each joint V_i , and sets their masked probability:

$$p_i = \frac{d_i}{\sum_{j=1}^n d_j}. \quad (3)$$

Motion Attention Temporal Masking

Compared with the spatial joints, the temporal information is more redundant and the semantic of action is more concentrated. Some key-frames play an pivotal role in representing the semantics of an action. To utilize this feature, our MATM

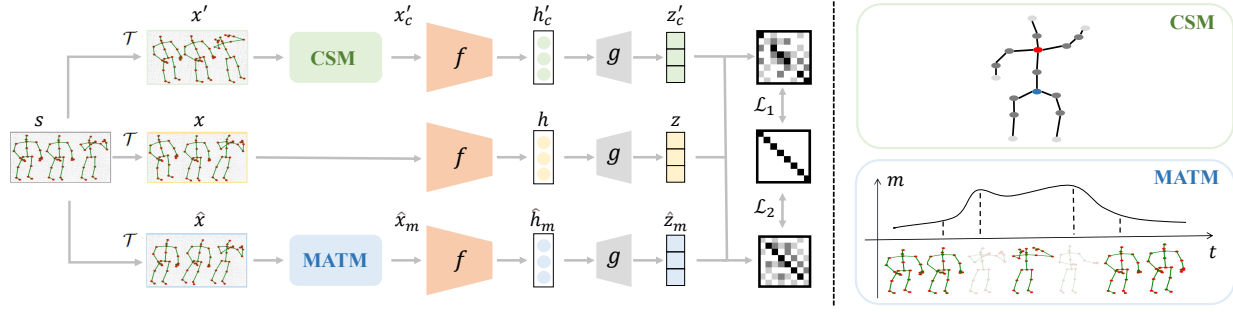


Figure 2: The overall framework of Partial Spatio-Temporal skeleton representation Learning. In Central Spatial Masking (CSM), the degree is red (4), blue (3), dark grey (2) and light grey (1). In Motion Attention Temporal Mask (MATM), m denotes motion density while t is the time code.

aims to find such key-frames and mask them to generate a hard pair for anchor stream and temporal masking stream in our proposed PSTL. By encouraging the empirical cross-correlation matrix between embeddings of two stream be an identity matrix, the encoder can capture the relationship between redundant frames and key-frames effectively.

Specifically, the motion $m \in \mathbb{R}^{C \times T \times V}$ of the sequence s is computed by the temporal displacement between frames: $m_{:,t,:} = x_{:,t+1,:} - x_{:,t,:}$. Then we acquire the overall motion rate of a frame which serves as the attention weight:

$$a_t = \frac{m_t^2}{\sum_{i=1}^T m_i^2}. \quad (4)$$

Then, the top- K attention weights a_{t_1}, \dots, a_{t_K} are selected and the sequence x_{t_1}, \dots, x_{t_K} serves as the key-frames which contain more semantic information about the actions. As shown in the blue box of Fig 2, the bottom one is a skeleton sequence of *put on glasses*, the top one shows the motion density. We can find that some frames in the middle of the sequence acquire higher motion values, which are the key-frames for this action. Thus, we mask these key-frames and encourage the feature from masked sequence close to the anchor feature which contains total semantic information.

Loss Function

Recall the pipeline of the proposed triple stream PSTL shown in Fig. 2. Given a skeleton sequence s , an ordinary augmentation is applied to obtain different views x, x', \hat{x} . For sequence x' , a Central Spatial Masking strategy is used to generate partial skeleton data x'_c . Considering the degree centrality of the skeleton graph, the joints with more degrees are assigned higher masked probability in this module. For sequence \hat{x} , a Motion Attention Temporal Masking strategy is applied. We first compute the motion value of each frame which represents the key fragment in a skeleton action. Then the top- K frames are selected to be masked and to further improve diversity, K random frames are still masked from remain sequence to generate \hat{x}_m . The sequence x in the anchor stream does not introduce any mask strategy to keep the original semantic information.

After that, the shared encoder f is used to extract the features simultaneously. $h = f(x; \theta)$, $h'_c = f(x'; \theta)$, $\hat{h}_m =$

$f(\hat{x}_m; \theta)$, where θ is the parameters learned by the encoder which is also used in downstream tasks. Then a projector g is applied and the features are projected to the embeddings z, z'_c, \hat{z}_m in a high dimension space. *i.e.* $z = g(h)$, $z'_c = g(h'_c)$, $\hat{z}_m = g(\hat{h}_m)$. To capture the relationship between masked joints and unmasked ones, we compute the cross-correlation matrix C' between embeddings z and z'_c . The calculation process refers to Eq. (2). And then \mathcal{L}_1 is applied:

$$\mathcal{L}_1 = \sum_i (1 - C'_{ii})^2 + \lambda \sum_i \sum_{j \neq i} C'^2_{ij}, \quad (5)$$

where the first term encourages the diagonal element of C' to 1. It forces the representation of partial data to be similar to the total one. The second term is designed to prompt the different embedding components to be decoupled with each other, which helps to minimize the redundancy within the representation and avoid it being a constant.

Similarly, the cross-correlation matrix \hat{C} between z and \hat{z}_m is acquired for the \mathcal{L}_2 :

$$\mathcal{L}_2 = \sum_i (1 - \hat{C}_{ii})^2 + \lambda \sum_i \sum_{j \neq i} \hat{C}^2_{ij}, \quad (6)$$

which is utilized to capture the relationship between masked and unmasked frames. Note that the trade-off parameter λ is used to balance the dimension difference between first and second terms, so it keep same on Eq. (5) and Eq. (6).

The total loss \mathcal{L}_p for our PSTL can be formulated as:

$$\mathcal{L}_p = \mathcal{L}_1 + \mathcal{L}_2. \quad (7)$$

Experiments

Datasets

NTU-RGB+D 60. NTU-60 dataset (Shahroudy et al. 2016) is collected by Microsoft Kinect sensors. It contains 56,578 skeleton sequences of 60 action categories, performed by 40 volunteers. The skeleton of each subject is represented by 25 joints, while each video has no more than two subjects. There are two official dataset splits: 1) Cross-Subject (xsu): half of the subjects belong to the training set, and the

rest make up the testing sets; 2) Cross-View (xview): training and testing sets are captured by cameras with different views.

NTU-RGB+D 120. NTU-120 dataset (Liu et al. 2019) is the extended version of the NTU-60, which contains 113,945 skeleton sequences of 120 action categories, performed by 106 volunteers. There are two official dataset splits: 1) Cross-Subject (xsub): the training set is performed by 53 subjects, and testing data is performed by the rest volunteers; 2) Cross-Setup (xset): the training and testing sets are captured by cameras with different setups.

PKU-MMD. PKU-MMD dataset (Liu et al. 2020) is captured via the Kinect v2 sensors from multiple viewpoints. It has almost 20000 action instances in 51 categories collected by 66 subjects. There are two parts of the dataset: 1) Part I contains 21539 sequences; 2) Part II contains 6904 sequences. Part II is more challenging due to the large view variation with more skeleton noises.

Implementation Details

To perform an apple-to-apple comparison with other methods, we follow the same pre-processing methods of CrosS-CLR (Li et al. 2021) and AimCLR (Guo et al. 2022), which resize the skeleton sequences to 50 frames. Similarly, ST-GCN (Yan, Xiong, and Lin 2018) with 16 hidden channels is used as the backbone. For all experiments (both representation learning and downstream tasks), we adopt Adam optimizer and the CosineAnnealing scheduler with 150 epochs. The mini batch size is 128.

Data Augmentation. Before applying the joint and temporal masking strategy proposed, some ordinary augmentations are used to generate different views of the input skeleton sequences. The ordinary augmentations include 3 spatial augmentations and 1 temporal augmentation.

(1) *Shear.* It is a linear transformation for skeleton 3D coordinates. A shear matrix is

$$\mathbf{S} = \begin{bmatrix} 1 & s_{12} & s_{13} \\ s_{21} & 1 & s_{23} \\ s_{31} & s_{32} & 1 \end{bmatrix},$$

where $s_{12}, s_{13}, s_{21}, s_{23}, s_{31}, s_{32}$ are the shear factors and are randomly sampled from $[-\beta \sim \beta]$. β is the amplitude to adjust the augmentation strength. Here we set $\beta = 1$. By directly multiplying the shear matrix with the 3D coordinates, the human pose slants in a random angle.

(2) *Crop.* As a temporal augmentation, Crop aims to improve data diversity by first padding part of frames to the original sequence and then randomly cropping it to the original length. The padding length is γT , T is the original sequence length, and γ is the padding ratio (we set $\gamma = 1/6$).

(3) *Rotate.* This is an efficient spatial augmentation which helps the representation adapt to the variation of spatial perspective. Specifically, it first randomly selects an axis $M \in \{X, Y, Z\}$ as main axis and then a random angle in $[0 \sim \pi/6]$ is selected for it to rotate. Next, the rest two axis both randomly choose the angle in $[0 \sim \pi/180]$.

(4) *Spatial Flip.* It is another spatial augmentation, which swaps the left and right sides of the skeleton data with the probability $p = 0.5$.

Experiment Setting

Self-supervised Pre-training. We use ST-GCN (Yan, Xiong, and Lin 2018) as the encoder, which extracts 256-dim features from skeleton sequences. Then we project features to 6144-dim embeddings. λ in the loss of each stream is set to $2e-4$. A 10-epoch warmup is used for stabilizing the training process. The weight decay is set to $1e-5$. Note that for all datasets and evaluation settings, the number of mask joints is 9 in CSM and $K = 10$ in MATM. To compare with the state-of-the-art, we use multiple skeleton modalities, including joint (**J**), bone (**B**), and motion (**M**). We evaluate the learned representations with various evaluation protocols.

Linear Evaluation. We add a linear classifier on top of the frozen pre-trained encoder, and then train the recognizer on the target skeleton action recognition dataset with initial learning rate 0.01.

Partial Body Evaluation. To encourage the learned representation to adapt to real-life scenarios with skeleton shaded or partial body information lost, we introduce the Partial Body Evaluation protocol. Basically, we follow the settings of Linear Evaluation. However, some joints or body parts may be masked to simulate the joint/part missing in extracted skeletons. We consider two settings: 1) Joint shaded: some joints are masked randomly during evaluation; 2) Body part shaded: we divide skeletons into 5 major parts¹ and randomly drop parts during evaluation.

Finetuning. We append a linear classifier on top of the encoder and then finetune the entire network on the target skeleton action recognition dataset with an initial learning rate 0.005.

Semi-Supervised Learning. This protocol shares the same setting as “Finetuning”, but only uses 1% or 10% training labeled data that is randomly selected.

Ablation Studies of PSTL

We ablate the proposed CSM and MATM strategies on our baseline SkeletonBT respectively to show their effectiveness. The linear evaluation results are shown in Tab. 2, SkeletonBT achieves 68.5% and 71.6% accuracy on xsub and xview. We find that the CSM module improves about 3.3%, 3.8% Top-1 accuracies and the MATM module brings about 5.2%, 6.4% performance gain on xsub and xview, respectively. With the help of both modules, our PSTL achieves 8.8% and 9.4% increments. These results show that our proposed CSM and MATM strategies can encourage the encoder to acquire stronger features which are more suitable for downstream tasks.

Comparisons with SkeletonBT. To further verify the effectiveness of our proposed method, we conduct experiments on three datasets to compare our PSTL with the baseline method SkeletonBT in detail. As shown in Tab. 1, PSTL performs much better than SkeletonBT on all three different streams. For the fusion results, 3s-PSTL substantially improves SkeletonBT, especially for the NTU-120 datasets.

Comparisons with Random Masking Strategy. To demonstrate the effectiveness of our masking strategy, we

¹left/right arms, left/right legs, torso

Method	Stream	NTU-60(%)				PKU(%)				NTU-120(%)			
		xsub		xview		part I		part II		xsub		xset	
		acc.	Δ	acc.	Δ	acc.	Δ	acc.	Δ	acc.	Δ	acc.	Δ
SkeletonBT	J	68.5		71.6		84.7		38.2		53.9		52.7	
PSTL	J	77.3	+8.8	81.8	+10.2	88.4	+3.9	49.3	+11.1	66.2	+12.3	67.7	+15.0
SkeletonBT	M	58.9		58.2		63.2		30.3		43.5		44.7	
PSTL	M	69.9	+11.0	72.6	+14.4	80.4	+17.2	34.7	+4.4	65.4	+21.9	57.7	+13.0
SkeletonBT	B	62.4		64.7		79.1		29.8		52.4		53.6	
PSTL	B	75.4	+13.0	79.2	+14.5	84.8	+5.7	41.3	+11.5	58.7	+6.3	66.6	+13.0
3s-SkeletonBT	J+M+B	69.3		72.2		84.8		44.1		56.3		56.8	
3s-PSTL	J+M+B	79.1	+9.8	82.6	+10.4	89.2	+4.4	52.3	+8.2	69.2	+12.9	70.3	+13.5

Table 1: Linear evaluation results compared with SkeletonBT on NTU-60, PKU-MMD, and NTU-120 datasets. ” Δ ” represents the gain compared to SkeletonBT using the same stream data. J, M and B indicate joint stream, motion stream, and bone stream.

SkeletonBT	w/ CSM	w/ MATM	NTU-60(%)	
			xsub	xview
✓			68.5	71.6
✓	✓		71.8	75.4
✓		✓	73.7	78.0
✓	✓	✓	77.3	81.8

Table 2: Ablation studies of PSTL on NTU-60 dataset.

Method	NTU-60(%)	
	xsub	xview
RSM + RTM	75.6	79.7
RSM + MATM	76.4	81.2
CSM + RTM	76.1	80.4
CSM + MATM	77.3	81.8

Table 3: Compared with random masking strategy on NTU-60 dataset. RSM means random spatial masking and RTM means random temporal masking.

compare our CSM and MATM modules with random spatial masking (RSM) strategy and random temporal masking (RTM) strategy, which selects the mask position randomly. The linear evaluation results are shown in Tab. 3, with the help of our triple stream network, the random masking can capture the relationship between joints and frames and get a good results on downstream tasks. Our CSM and MATM can further strengthen this capacity to get a better accuracy.

Comparison with State-of-the-art

Partial Body Evaluation. To verify the *joint robustness* of extracted features, we conduct the partial body evaluation on NTU-60 xsub datasets. The first level of the evaluation protocols is **joint shaded**, which applies random masks on joints to simulate the situation of joints shade in realistic scenarios. Tab. 4 shows that, without CSM and MATM, the baseline SkeletonBT suffers serious performance loss. With the number of mask joints increasing to 10, the SOTA method AimCLR (Guo et al. 2022) and its baseline methods SkeletonCLR (Li et al. 2021) also suffer about 20.0% and 17.6% performance decrease. Note that the feature extracted by our PSTL acquires strong joint robustness, and

# of masked joints	0	2	4	6	8	10
SkeletonCLR	68.3	65.6	61.9	58.6	54.5	50.7
SkeletonBT	68.5	65.0	61.7	58.6	55.0	50.6
AimCLR	74.3	70.4	67.2	63.0	58.5	54.3
PSTL(ours)	77.3	76.9	76.2	75.7	74.6	73.8

Table 4: Linear evaluation results of joints shaded on NTU-60 xsub dataset.

# of masked body parts	0	1	2	3	4
SkeletonCLR	68.3	63.6	57.7	47.5	33.5
SkeletonBT	68.5	64.0	58.5	50.2	38.4
AimCLR	74.3	69.3	63.9	53.7	37.9
PSTL (ours)	77.3	73.6	68.3	60.0	47.3

Table 5: Linear evaluation results of body part shaded on NTU-60 xsub dataset.

only 3.5% performance decrease under 10 masked joints.

The second level evaluation protocols **body part shaded**, as shown in Tab. 5, is much harder since we mask the total parts of the body. But our PSTL still achieves better results than other methods. To sum up, more skeleton sequence information is lost, and the advantage of our method becomes more obvious. And the results tell that with the help of local relationship between joints and frames, our CSM and MATM strategy can improve the robustness of learned features, which is essential to downstream tasks.

Linear Evaluation. As shown in Tab. 6, for a single stream, PSTL achieves considerable gains on all other methods on NTU60 and NTU120. Compared with the current SOTA method AimCLR, our PSTL improves 3.0% and 2.1% on xsub and Xview respectively. The results indicate that the local relationship between joints and frames can help extract suitable features for downstream tasks. For the performance of the 3-streams, PSTL still achieves comparable results to current methods. Our method also achieves the best results on NTU-120 datasets. For single stream, PSTL outperforms the current SOTA method AimCLR 2.8% and 4.3% on xsub and Xset respectively. And the fusion can also achieve 1.0% and 1.5% increment.

Noting that the PKU-MMD dataset has about half the data

Method	NTU-60(%)		NTU-120(%)	
	xsub	xview	xsub	xset
<i>Single-stream:</i>				
LongT GAN (AAAI 18)	39.1	48.1	-	-
MS2L (ACM MM 20)	52.6	-	-	-
AS-CAL (Inf Sci 21)	58.5	64.8	48.6	49.2
P&C (CVPR 20)	50.7	76.3	42.7	41.7
SeBiReNet (ECCV 20)	-	79.7	-	-
AimCLR (AAAI 22)	74.3	79.7	63.4	63.4
PSTL (ours)	77.3	81.8	66.2	67.7
<i>Three-stream:</i>				
3s-SkeletonBT (ours)	69.3	72.2	56.3	56.8
3s-SkeletonCLR (CVPR 21)	75.0	79.8	60.7	62.6
3s-Colorization (ICCV 21)	75.2	83.1	-	-
3s-CrosSCLR (CVPR 21)	77.8	83.4	67.9	67.1
3s-AimCLR (AAAI 22)	78.9	83.8	68.2	68.8
3s-PSTL (ours)	79.1	83.8	69.2	70.3

Table 6: Linear evaluation results on NTU-60 and NTU-120.

Method	Part I	Part II
<i>Single-stream:</i>		
LongT GAN (AAAI 18)	67.7	26.0
MS2L (ACM MM 20)	64.9	27.6
SkeletonCLR (CVPR 21)	80.9	35.2
ISC (ACM MM 20)	80.9	36.0
AimCLR (AAAI 22)	83.4	36.8
SkeletonBT (ours)	84.7	38.2
PSTL(ours)	88.4	49.3
<i>Three-stream:</i>		
3s-CrosSCLR (CVPR 21)	84.9	21.2
3s-AimCLR (AAAI 22)	87.8	38.5
3s-SkeletonBT (ours)	84.8	44.1
3s-PSTL (ours)	89.2	52.3

Table 7: Linear evaluation results on PKU-MMD.

size of NTU-60, Tab. 7 shows that the advantage of our method is more obvious in a small dataset. It is worth mentioning that our baseline method SkeletonBT can achieve better results than AimCLR on the single stream, which indicates that our negative-sample-free methods without a memory bank is more efficient for small datasets. Similarly, our PSTL achieves remarkable results on PKU-MMD datasets. For single stream, PSTL outperforms AimCLR 5.0% and 12.5% on Part I and Part II. For 3 stream fusion, our 3s-PSTL outperforms 3s-AimCLR 1.4% and 13.8% on both parts.

Finetune Evaluation. We compare PSTL with other self-supervised methods and the supervised method ST-GCN using the same structure and parameters as ours. As shown in Tab. 8, for the single steam, all methods use bone steam since the results are better than the joint stream. Our PSTL achieves better performance than others on NTU-60 and NTU-120 datasets. Notably, only single steam PSTL outperforms supervised 3s-ST-GCN on NTU-120 datasets, which indicates the effectiveness of our method. Similarly, 3s-PSTL outperforms 3s-CrosSCLR and 3s-AimCLR for the fusion results on both datasets.

Semi-Supervised Evaluation. To show the superiority of our approach, we test all methods under the semi-supervised setting, where only 1% or 10% labeled data can be used such that models are more likely to overfit. As shown in Tab. 9,

Method	NTU-60(%)		NTU-120(%)	
	xsub	xview	xsub	xset
<i>Single-stream:</i>				
SkeletonCLR (CVPR 21)	82.2	88.9	73.6	75.3
AimCLR (AAAI 22)	83.0	89.2	77.2	76.0
PSTL (ours)	84.5	92.0	78.6	78.9
<i>Three-stream:</i>				
3s-ST-GCN (AAAI 18) §	85.2	91.4	77.2	77.1
3s-CrosSCLR (CVPR 21)	86.2	92.5	80.5	80.4
3s-AimCLR (AAAI 22)	86.9	92.8	80.1	80.9
3s-PSTL (ours)	87.1	93.9	81.3	82.6

Table 8: Finetune evaluation results on NTU-60 and NTU-120. "§" means the model is trained with a fully-supervised manner.

Method	Part I	Part II
<i>1% labeled data:</i>		
LongT GAN (AAAI 18)	35.8	12.4
MS2L (ACM MM 20)	36.4	13.0
ISC (ACM MM 21)	37.7	-
3s-CrosSCLR (CVPR 21)	49.7	10.2
3s-AimCLR (AAAI 22)	57.5	15.1
3s-PSTL(ours)	62.5	16.9
<i>10% labeled data:</i>		
LongT GAN (AAAI 18)	69.5	25.7
MS2L (ACM MM 20)	70.3	26.1
ISC (ACM MM 21)	72.1	-
3s-CrosSCLR (CVPR 21)	82.9	28.6
3s-AimCLR (AAAI 22)	86.1	33.4
3s-PSTL (ours)	86.9	42.0

Table 9: Semi-supervised evaluation results on PKU-MMD.

our PSTL far exceeds the current SOTA method AimCLR on PKU-MMD datasets. And for fusion results, 3s-PSTL outperforms 3s-AimCLR 8.6% on Part II. It demonstrates that such contrastive methods which need a large memory bank are subjected to serious performance loss with insufficient samples. But our negative-sample-free PSTL can well alleviate this problem and achieve better results.

Conclusion

In this paper, we introduce a triple stream network named Partial Spatio-Temporal skeleton sequences Learning framework for self-supervised skeleton-based action recognition. Firstly, all streams apply an ordinary augmentation for producing a variety of skeleton views. Secondly, one stream introduces a Central Spatial Masking strategy to force the selected joints out of the calculation of GCN backbone. The third stream utilizes a Motion Attention Temporal Masking module to find the keyframes and mask them. At last, the feature cross-correlation matrix is measured to acquire the correlation among each stream. Experiments show that such a masking strategy can boost the robustness of representations of skeleton joints, which is essential to downstream tasks. our PSTL achieves remarkable performance in three large skeleton datasets under various evaluation settings.

Acknowledgements

This work was supported by Shanghai AI Laboratory, the National Natural Science Foundation of China No. 61976206 and No. 61832017, Beijing Outstanding Young Scientist Program NO. BJJWZYJH012019100020098, Beijing Academy of Artificial Intelligence (BAAI), the Fundamental Research Funds for the Central Universities, the Research Funds of Renmin University of China 21XNLG05, and Public Computing Cloud, Renmin University of China.

References

- Cao, Z.; Simon, T.; Wei, S.-E.; and Sheikh, Y. 2017. Real-time Multi-Person 2D Pose Estimation Using Part Affinity Fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Chen, X.; and He, K. 2021. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15750–15758.
- Chen, Z.; Li, S.; Yang, B.; Li, Q.; and Liu, H. 2021. Multi-scale spatial temporal graph convolutional network for skeleton-based action recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 1113–1122.
- Du, Y.; Wang, W.; and Wang, L. 2015. Hierarchical recurrent neural network for skeleton based action recognition. *IEEE Conference on Computer Vision and Pattern Recognition*.
- Gidaris, S.; Singh, P.; and Komodakis, N. 2018. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*.
- Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P.; Buchatskaya, E.; Doersch, C.; Avila Pires, B.; Guo, Z.; Gheshlaghi Azar, M.; et al. 2020. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33: 21271–21284.
- Guo, T.; Liu, H.; Chen, Z.; Liu, M.; Wang, T.; and Ding, R. 2022. Contrastive Learning from Extremely Augmented Skeleton Sequences for Self-supervised Action Recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 762–770.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16000–16009.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9729–9738.
- Hua, G.; Liu, H.; Li, W.; Zhang, Q.; Ding, R.; and Xu, X. 2022. Weakly-supervised 3D Human Pose Estimation with Cross-view U-shaped Graph Convolutional Network. *IEEE Transactions on Multimedia*.
- Ke, Q.; Bennamoun, M.; An, S.; Sohel, A. F.; and Boussaïd, F. 2017. A New Representation of Skeleton Sequences for 3D Action Recognition. *CVPR*.
- Kim, S. T.; and Reiter, A. 2017. Interpretable 3D Human Action Analysis with Temporal Convolutional Networks. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, 1623–1631.
- Larsson, G.; Maire, M.; and Shakhnarovich, G. 2016. Learning representations for automatic colorization. In *European conference on computer vision*, 577–593. Springer.
- Li, L.; Wang, M.; Ni, B.; Wang, H.; Yang, J.; and Zhang, W. 2021. 3D Human Action Representation Learning via Cross-View Consistency Pursuit. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4741–4750.
- Lin, L.; Song, S.; Yang, W.; and Liu, J. 2020. Ms2l: Multi-task self-supervised learning for skeleton based action recognition. In *Proceedings of the 28th ACM International Conference on Multimedia*, 2490–2498.
- Liu, J.; Shahroudy, A.; Perez, M.; Wang, G.; Duan, L.-Y.; and Kot, A. C. 2019. Ntu rgb+ d 120: A large-scale benchmark for 3d human activity understanding. *IEEE transactions on pattern analysis and machine intelligence*, 42(10): 2684–2701.
- Liu, J.; Song, S.; Liu, C.; Li, Y.; and Hu, Y. 2020. A benchmark dataset and comparison study for multi-modal human action analytics. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 16(2): 1–24.
- Noroozi, M.; and Favaro, P. 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision*, 69–84. Springer.
- Rao, H.; Xu, S.; Hu, X.; Cheng, J.; and Hu, B. 2021. Augmented skeleton based contrastive action learning with momentum lstm for unsupervised action recognition. *Information Sciences*, 569: 90–109.
- Shahroudy, A.; Liu, J.; Ng, T.-T.; and Wang, G. 2016. Ntu rgb+ d: A large scale dataset for 3d human activity analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1010–1019.
- Shi, L.; Zhang, Y.; Cheng, J.; and Lu, H. 2019. Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12026–12035.
- Shotton, J.; Fitzgibbon, A.; Cook, M.; Sharp, T.; Finocchio, M.; Moore, R.; Kipman, A.; and Blake, A. 2011. Real-time human pose recognition in parts from single depth images. In *CVPR 2011*, 1297–1304. Ieee.
- Song, Y.-F.; Zhang, Z.; Shan, C.; and Wang, L. 2020. Stronger, faster and more explainable: A graph convolutional baseline for skeleton-based action recognition. In *proceedings of the 28th ACM international conference on multimedia*, 1625–1633.

- Su, K.; Liu, X.; and Shlizerman, E. 2020. Predict & cluster: Unsupervised skeleton based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9631–9640.
- Thoker, F. M.; Doughty, H.; and Snoek, C. G. 2021. Skeleton-contrastive 3D action representation learning. In *Proceedings of the 29th ACM International Conference on Multimedia*, 1655–1663.
- Vemulapalli, R.; Arrate, F.; and Chellappa, R. 2014. Human action recognition by representing 3d skeletons as points in a lie group. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 588–595.
- Wei, C.; Xie, L.; Ren, X.; Xia, Y.; Su, C.; Liu, J.; Tian, Q.; and Yuille, A. L. 2019. Iterative reorganization with weak spatial constraints: Solving arbitrary jigsaw puzzles for unsupervised representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1910–1919.
- Xia, L.; Chen, C.-C.; and Aggarwal, J. K. 2012. View invariant human action recognition using histograms of 3d joints. In *2012 IEEE computer society conference on computer vision and pattern recognition workshops*, 20–27. IEEE.
- Yan, S.; Xiong, Y.; and Lin, D. 2018. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Thirty-second AAAI conference on artificial intelligence*.
- Zbontar, J.; Jing, L.; Misra, I.; LeCun, Y.; and Deny, S. 2021. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, 12310–12320. PMLR.
- Zhang, P.; Lan, C.; Xing, J.; Zeng, W.; Xue, J.; and Zheng, N. 2017. View adaptive recurrent neural networks for high performance human action recognition from skeleton data. In *Proceedings of the IEEE international conference on computer vision*, 2117–2126.
- Zhang, R.; Isola, P.; and Efros, A. A. 2016. Colorful image colorization. In *European conference on computer vision*, 649–666. Springer.
- Zhang, Z. 2012. Microsoft kinect sensor and its effect. *IEEE multimedia*, 19(2): 4–10.
- Zheng, N.; Wen, J.; Liu, R.; Long, L.; Dai, J.; and Gong, Z. 2018. Unsupervised representation learning with long-term dynamics for skeleton based action recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.